

IPv4/IPv6 Routing and Deployment Workshop

July 10-14, 2012, Karachi, Pakistan

In conjunction with

The SANDOG logo is displayed in white, bold, uppercase letters on a black rectangular background. The letters are stylized with a slight gap between the 'O' and 'D'.

APNIC



Presenter

- Champika Wijayatunga
 - Training Unit Manager, APNIC
 - champika@apnic.net
- GZ Kabir
 - Systems Integrator
 - gzkabir@bdcom.com
- Vivek Nigam
 - Internet Resource Analyst, APNIC
 - vivek@apnic.net

Agenda

- Internet Fundamentals
- Internet Resource Management
- IP addressing and IP routing basics
- Introduction to IPv6 and Protocol Architecture
- IPv6 Addressing and Sub-netting
- IPv6 Host Configuration
- IPv4/IPv6 Deployment Plan – Case Study
- IPv4/IPv6 Deployment in IGP – Case Study
- IPv4 to IPv6 Transition Technologies
- IPv4/IPv6 Deployment in EGP – Case Study
- IPv6 DNS


Overview

- History of the Internet
- Evolution of Internet Eco System
- Regional Internet Registry (RIR) System
- Internet Operational Fundamentals
- Managing Internet Resources
- Internet Registry Policies

History of the Internet

In the beginning...

- 1968 - DARPA
 - (Defense Advanced Research Projects Agency) contracts with BBN to create ARPAnet
- 1969 – First four nodes



THE ARPA NETWORK

DEC 1969

4 NODES


Pre 1992

RFC 1020
1987


DDNIC



RFC 1261
1991



RFC 790
1981



“The assignment of numbers is also handled by Jon. If you are developing a protocol or application that will require the use of a link, socket, port, protocol, or network number **please contact Jon to receive a number assignment.**”

Address Architecture - History

- Initially, only 256 networks in the Internet!
- Then, network “classes” introduced:
 - Class A (128 networks x 16M hosts)
 - Class B (16,384 x 65K hosts)
 - Class C (2M x 254 hosts)

Internet Challenges 1992

- Address space depletion
 - IPv4 address space is finite
 - Historically, many wasteful allocations
- Routing chaos
 - Legacy routing structure, router overload
 - CIDR & aggregation are now vital
- Inequitable management
 - Unstructured and wasteful address space distribution

Classless & Classful addressing

Best Current Practice

Classful



Obsolete

- *inefficient*
- *depletion of B space*
- *too many routes from C space*

Classless


Addresses	Prefix	Classful	Net Mask
...
8	/29		255.255.255.248
16	/28		255.255.255.240
32	/27		255.255.255.224
64	/26		255.255.255.192
128	/25		255.255.255.128
256	/24	1 C	255.255.255.0
...
4096	/20	16 C's	255.255.240.0
8192	/19	32 C's	255.255.224
16384	/18	64 C's	255.255.192
32768	/17	128 C's	255.255.128
65536	/16	1 B	255.255.0.0
...

- Network boundaries may occur at *any* bit

Evolution of Internet Eco System

Evolution of Internet Resource Management

- 1993: Development of “CIDR”
 - addressed both technical problems



Address depletion

→ Through more accurate assignment

- variable-length network address


Routing table overload

→ Through address space aggregation

- “supernetting”

Evolution of Address Policy

- Establishment of RIRs
 - Regional open processes
 - Cooperative policy development
 - Industry self-regulatory model
 - bottom up

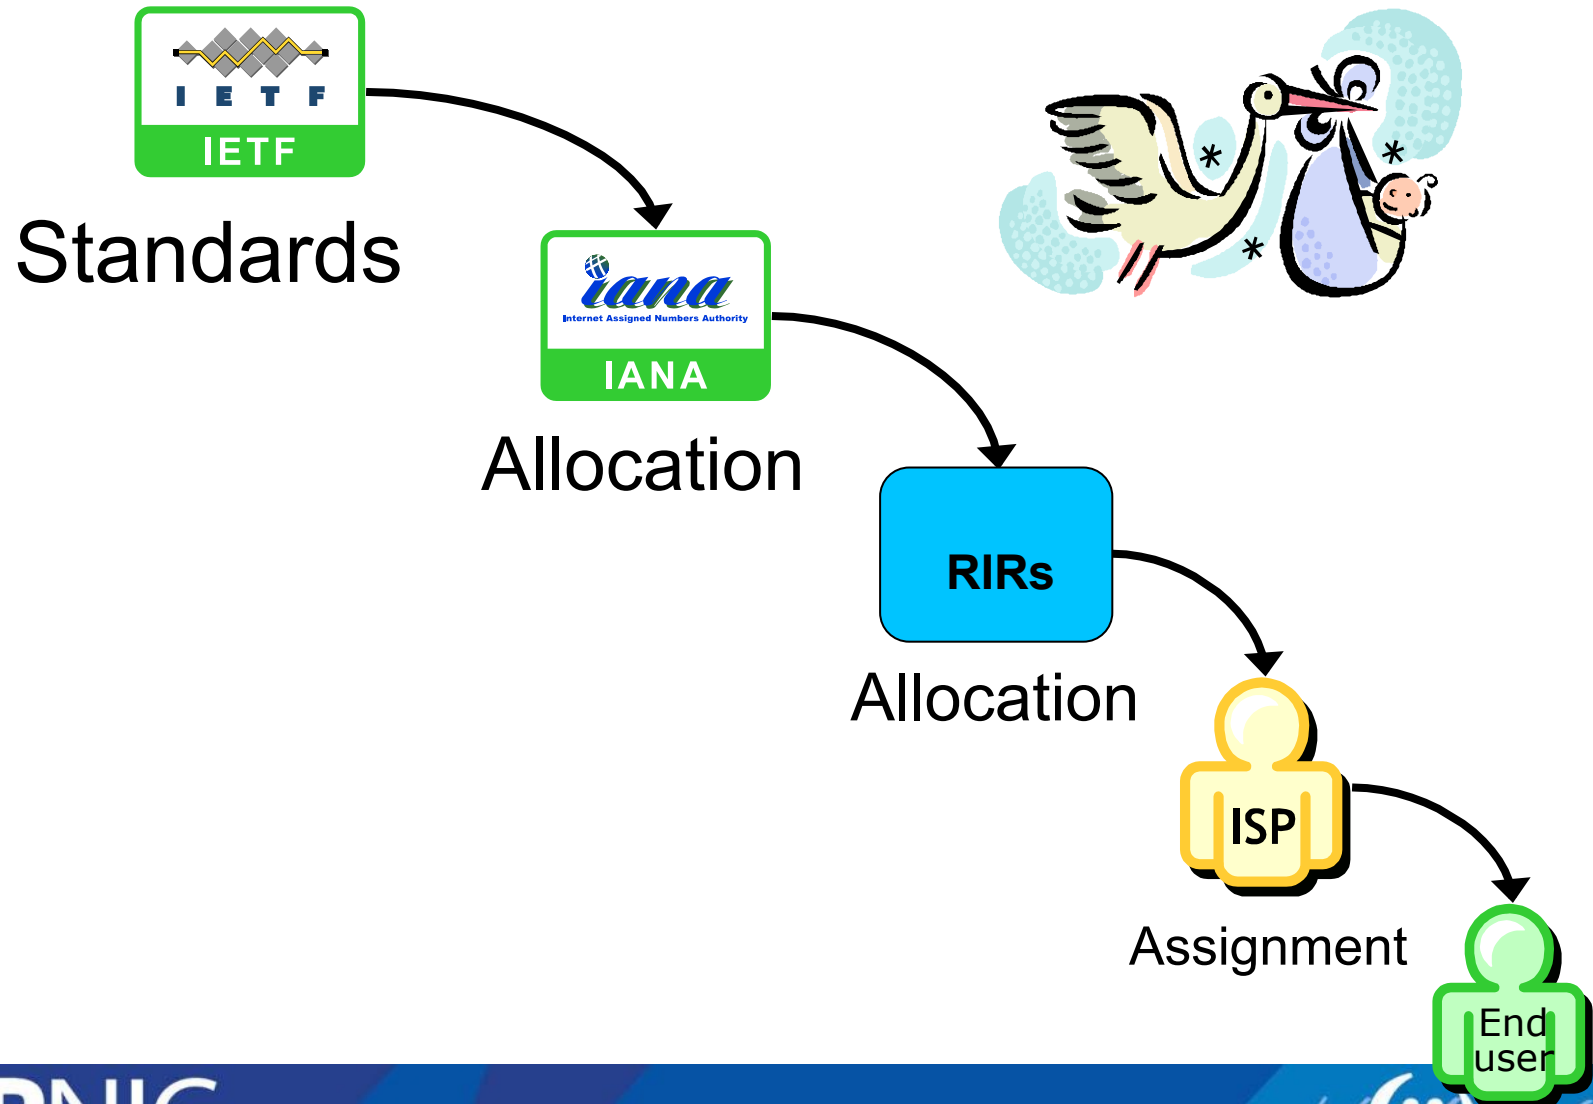


Regional Internet Registry System

Where Are The RIR Regions?



Where do IP addresses come from?

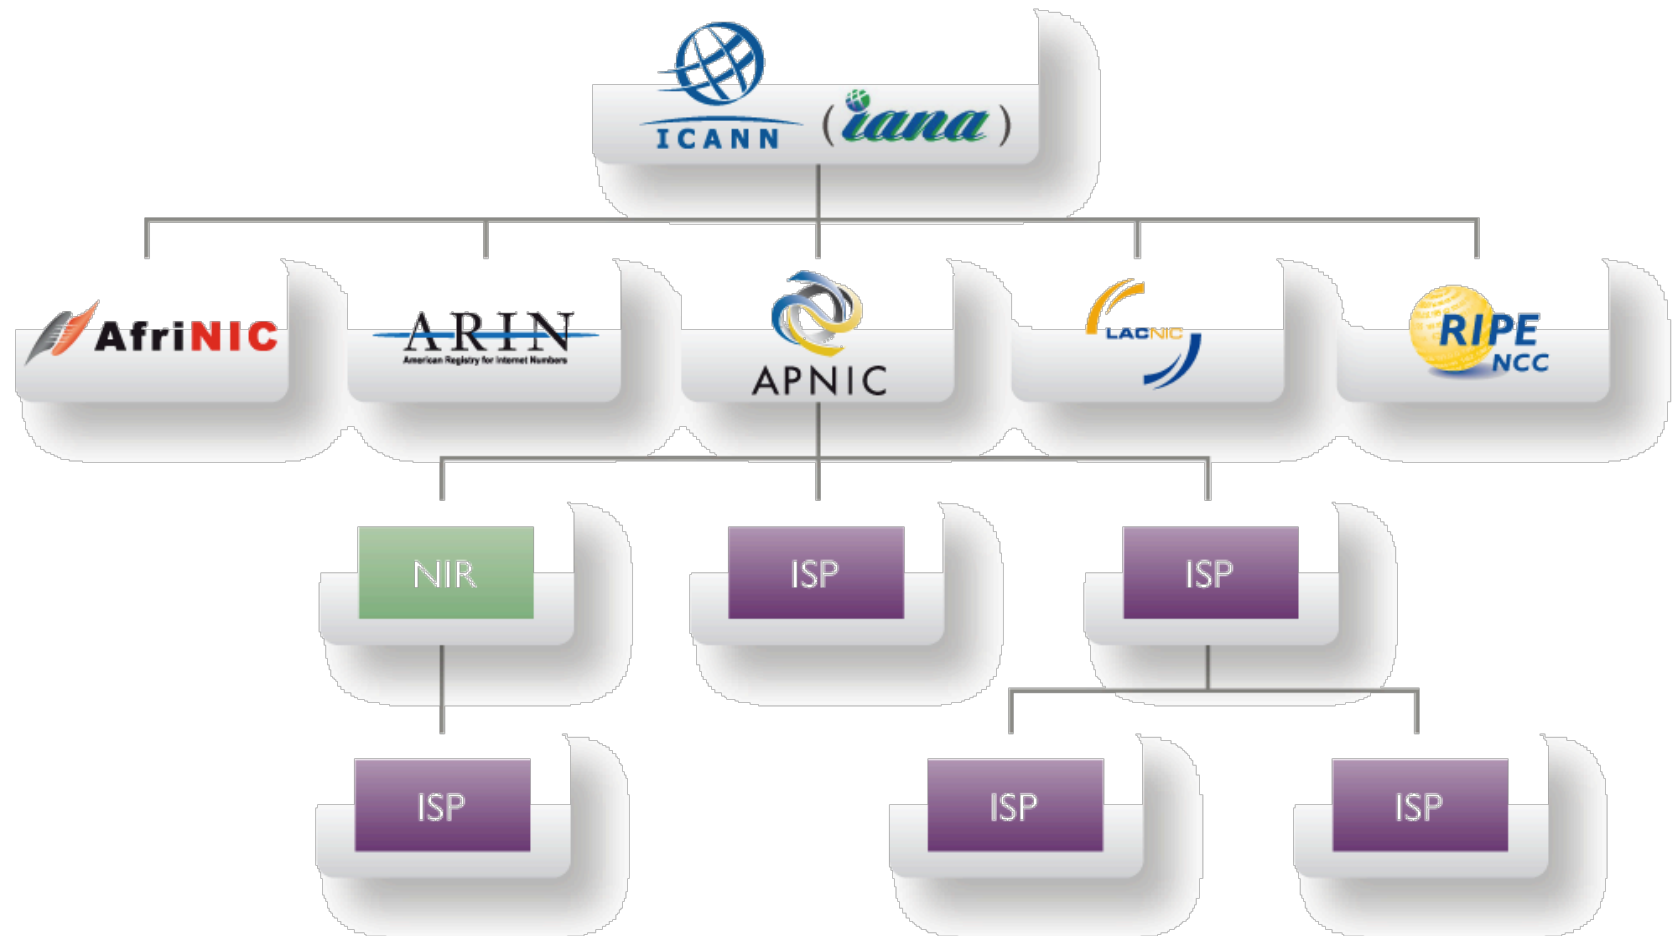


Regional Internet Registries

- RIRs manage, distribute, and register Internet number resources (IPv4 and IPv6 addresses and Autonomous System Numbers) within their respective regions.
 - Ensuring the fair distribution and responsible management
- Five RIRs:
 - AfriNIC, APNIC, ARIN, LACNIC, RIPE NCC



Regional Internet Registry Structure



What are the Goals of the RIRs?

- The Regional Internet Registries have been charged with the following goals for the number resources they are responsible for:
 - Conservation
 - Aggregation
 - Registration

What is role of RIR?

- Provides resource services to their respective regions
 - IPv4, IPv6, ASN
 - Resource registration
 - Authoritative registration server: Whois
 - Provides reverse DNS delegation for the resources allocated to the region

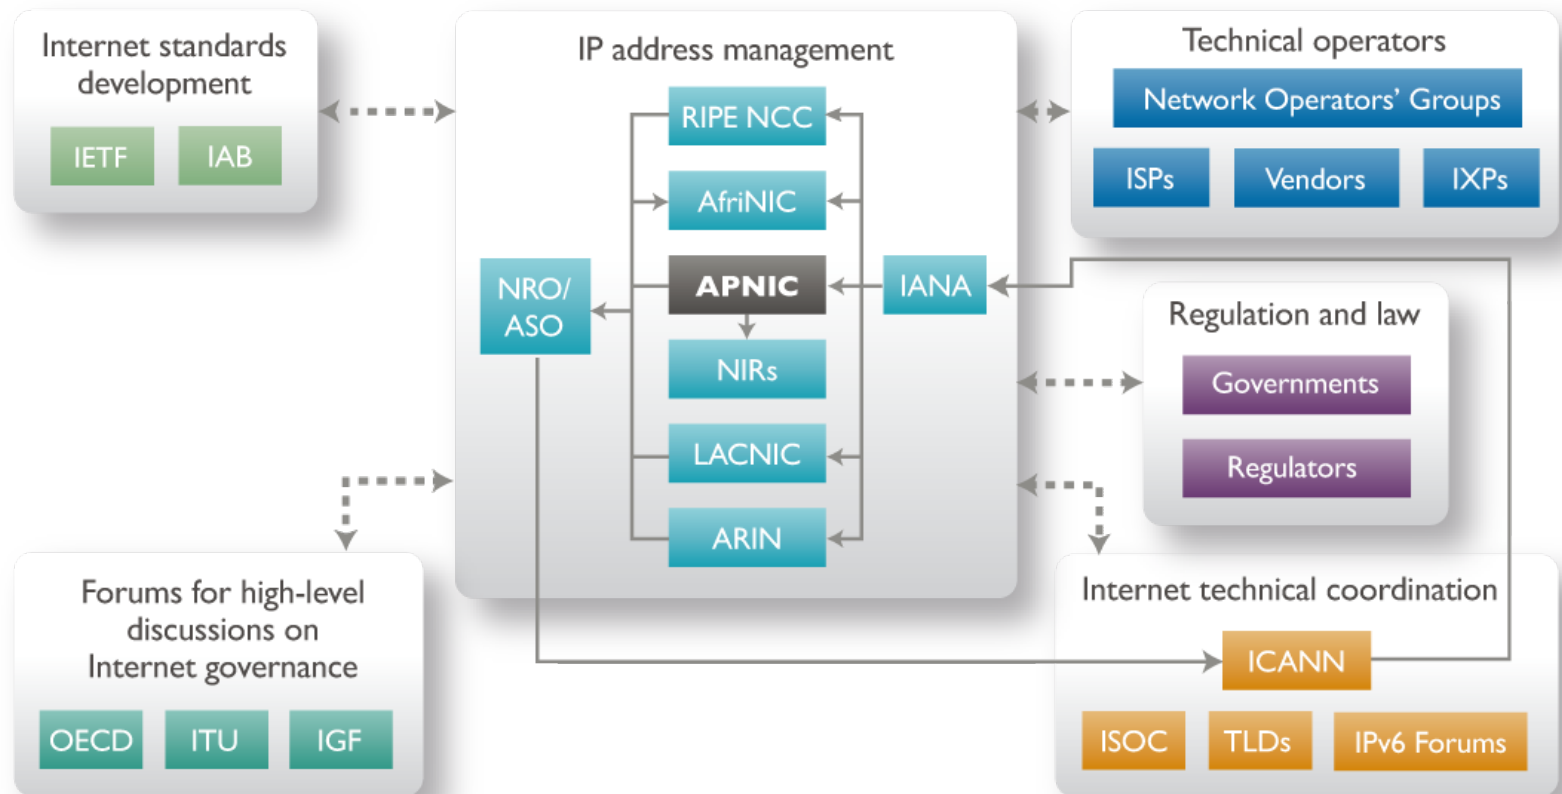


What RIR's Do?

- Facilitate the policy development process
 - Via mailing lists and meetings
- Implements policy changes
 - When the community has discussed and agreed to them
- Information dissemination
- Training services
- Collaboration & Liaison



Internet Eco System Today – a Global Perspective



Getting to know APNIC

What is APNIC?

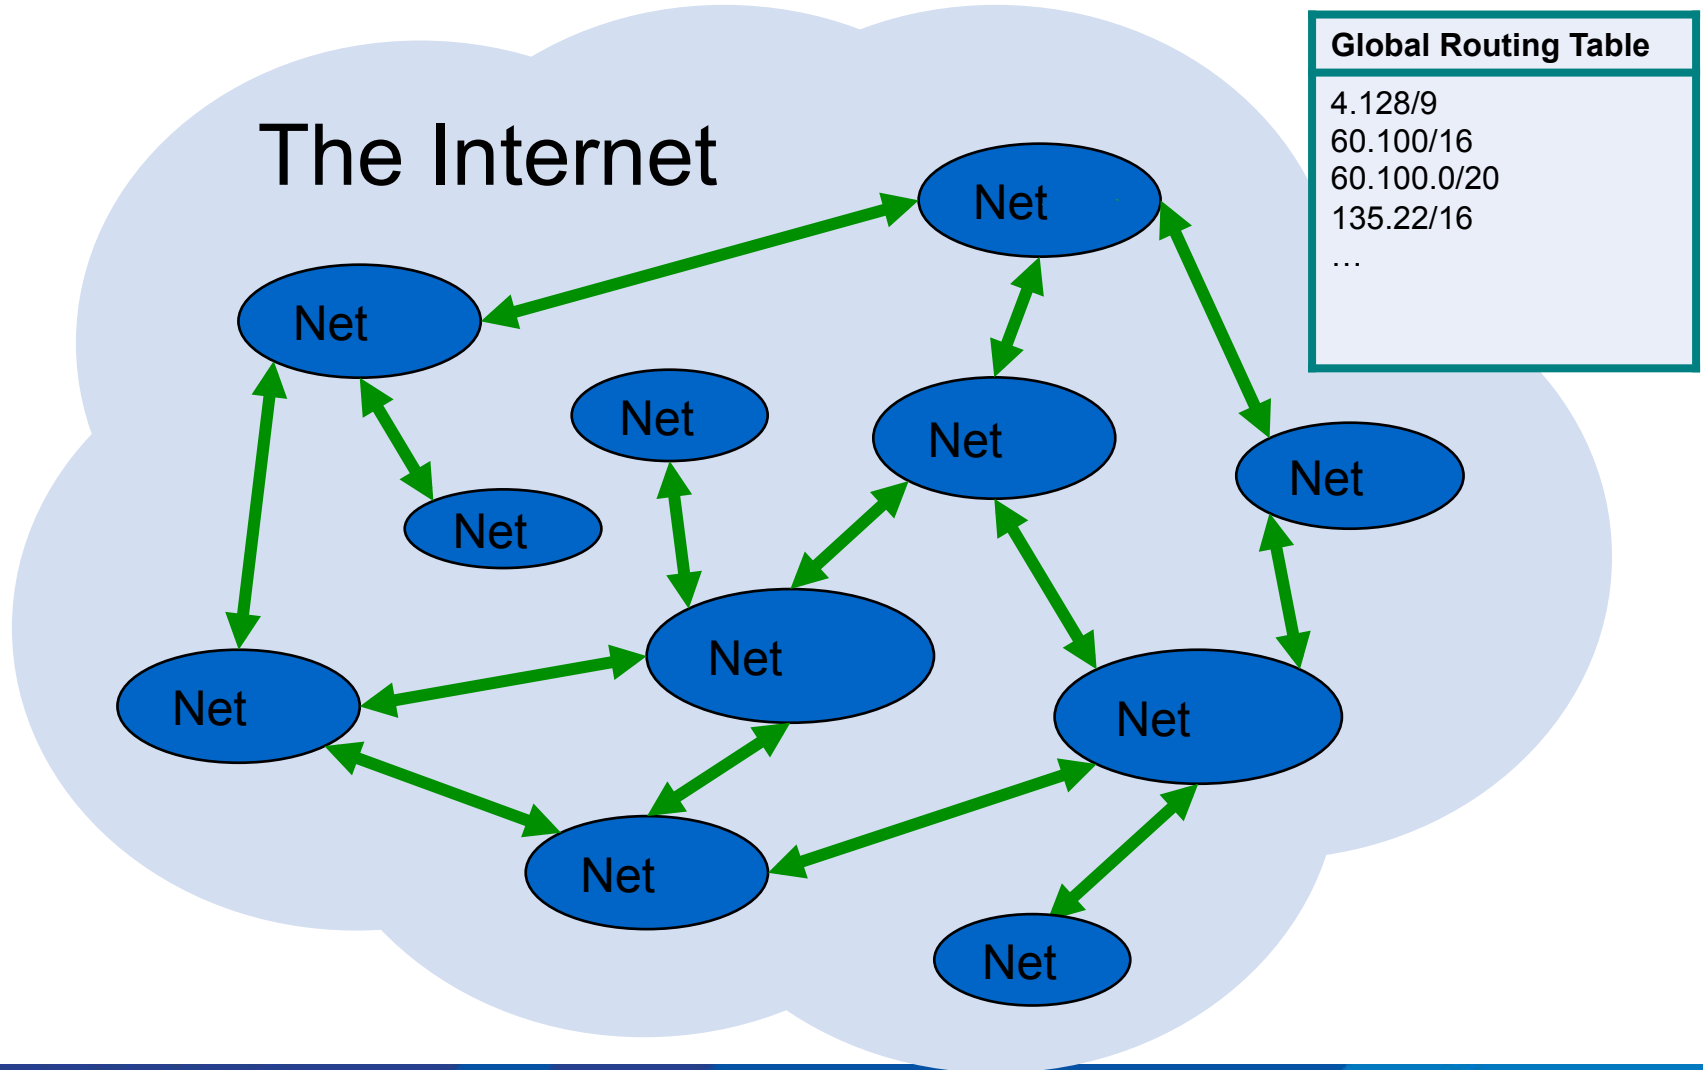
- APNIC is one of 5 Regional Internet Registries (RIRs) around the world.
- APNIC takes care of the Asia Pacific region
- APNIC is a non-profit, membership based organisation
- Policies are proposed and agreed upon by the APNIC community.

Internet Operational Fundamentals

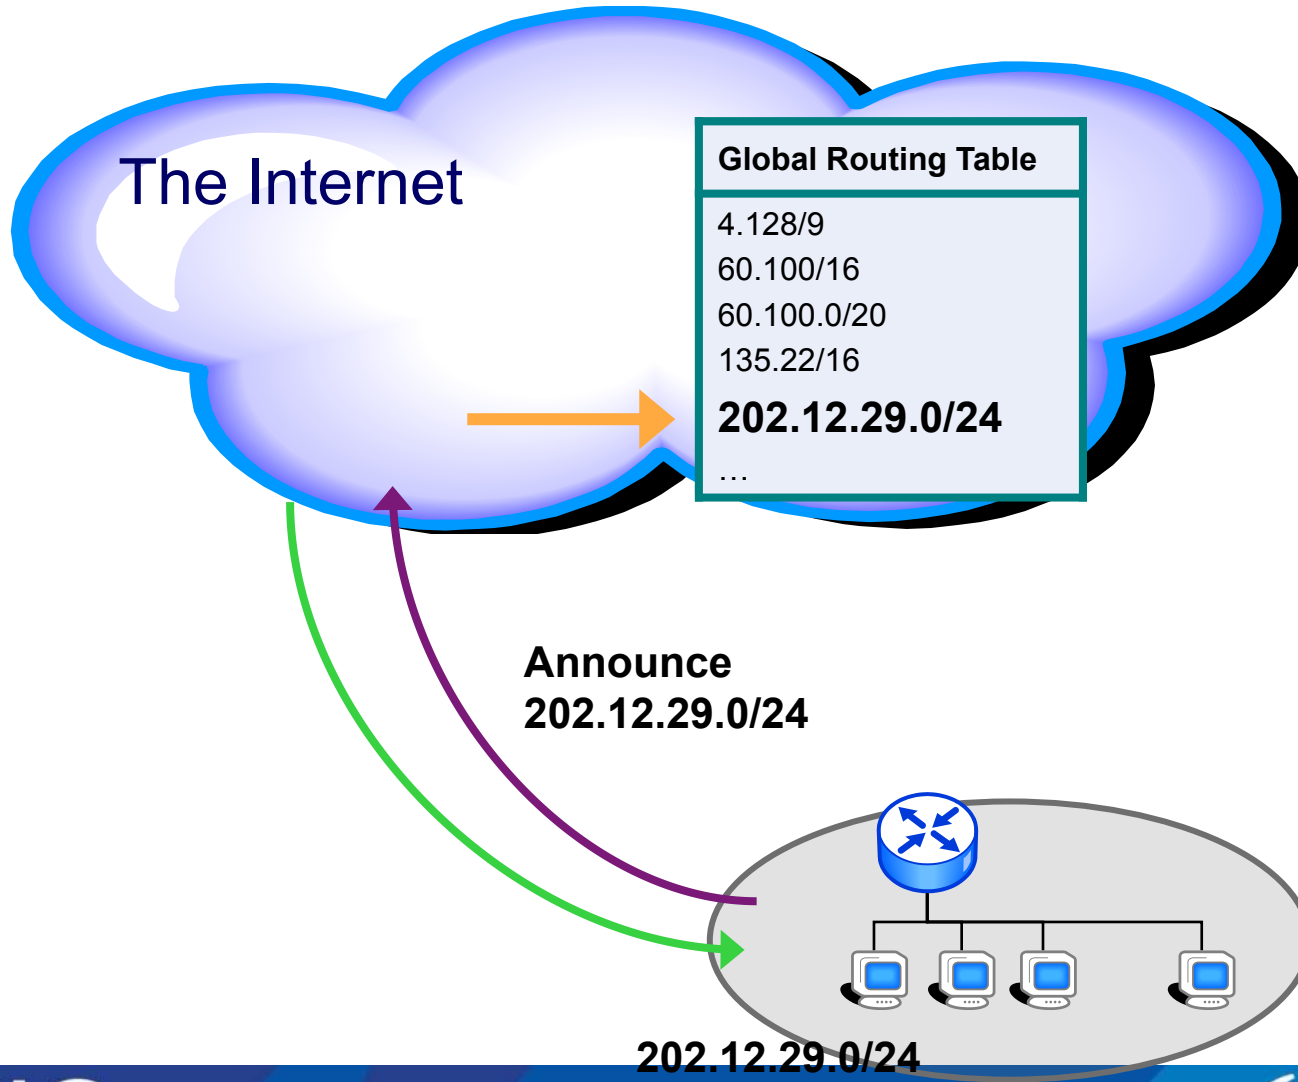
How does the Internet work

- Physical connectivity and reachability
 - Packet switching
- Protocols – common communication and rules
 - TCP/IP
- Addressing – global accessibility
 - IPv4, AS numbers, IPv6
 - IANA - RIRs


Internet Routing



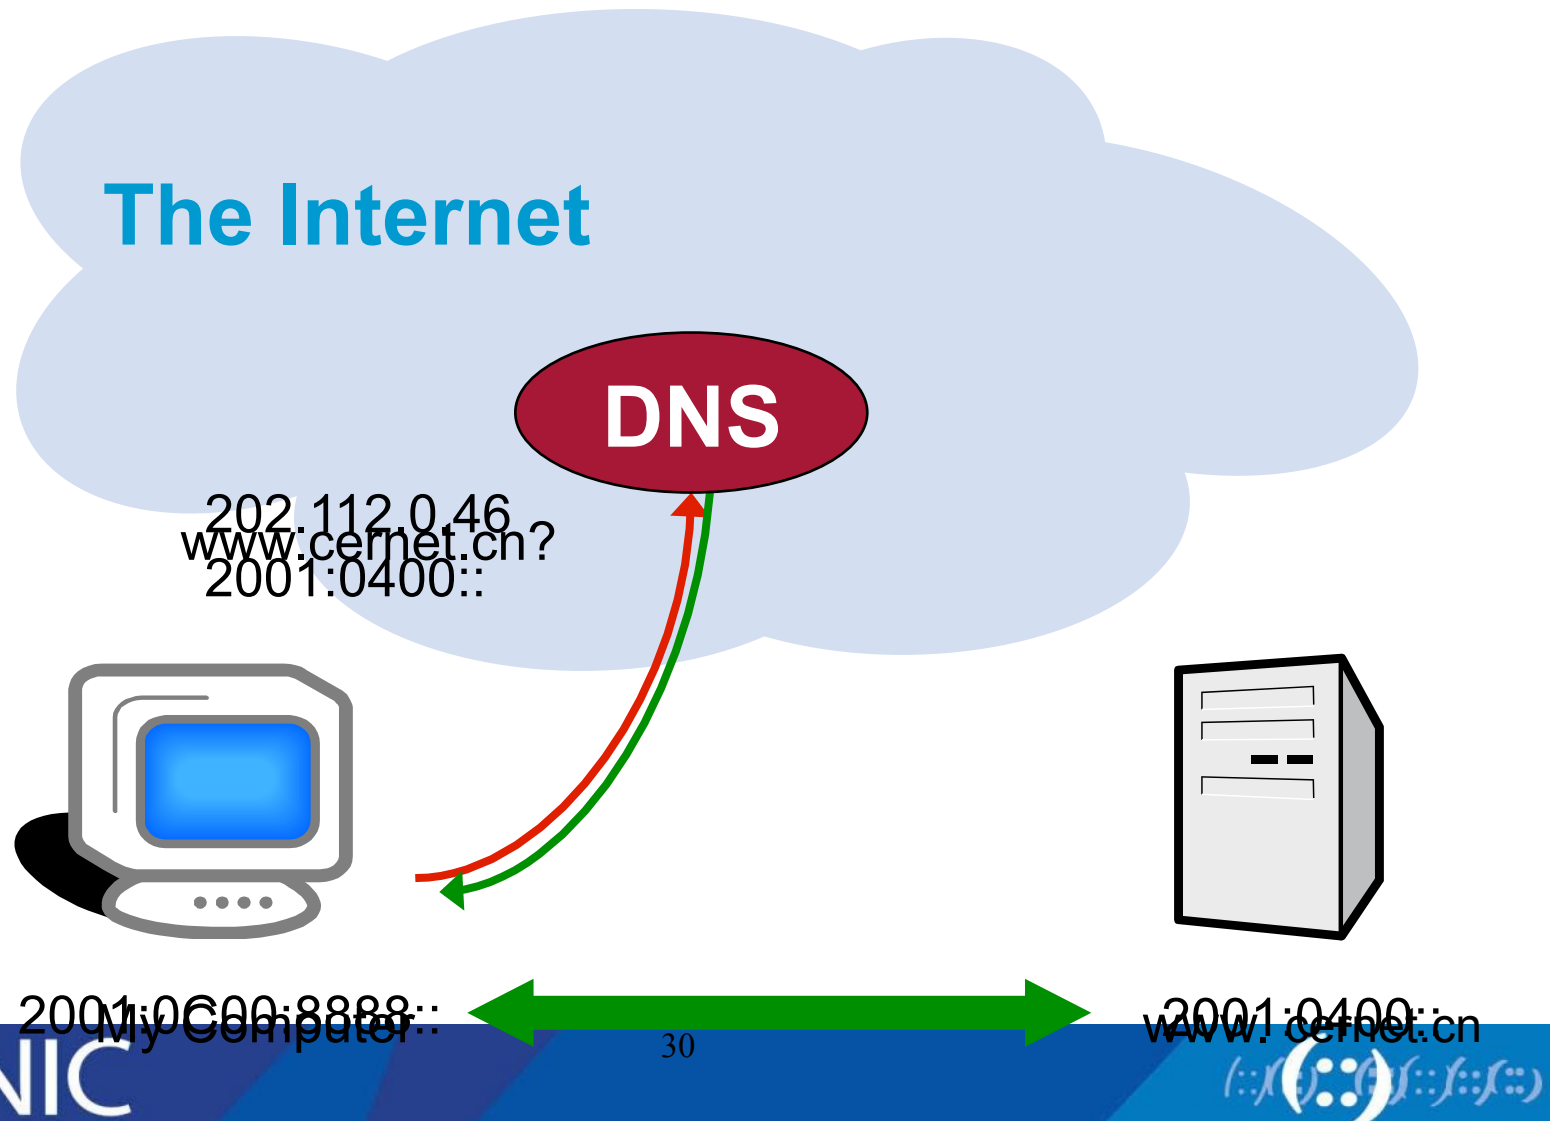
Internet Routing



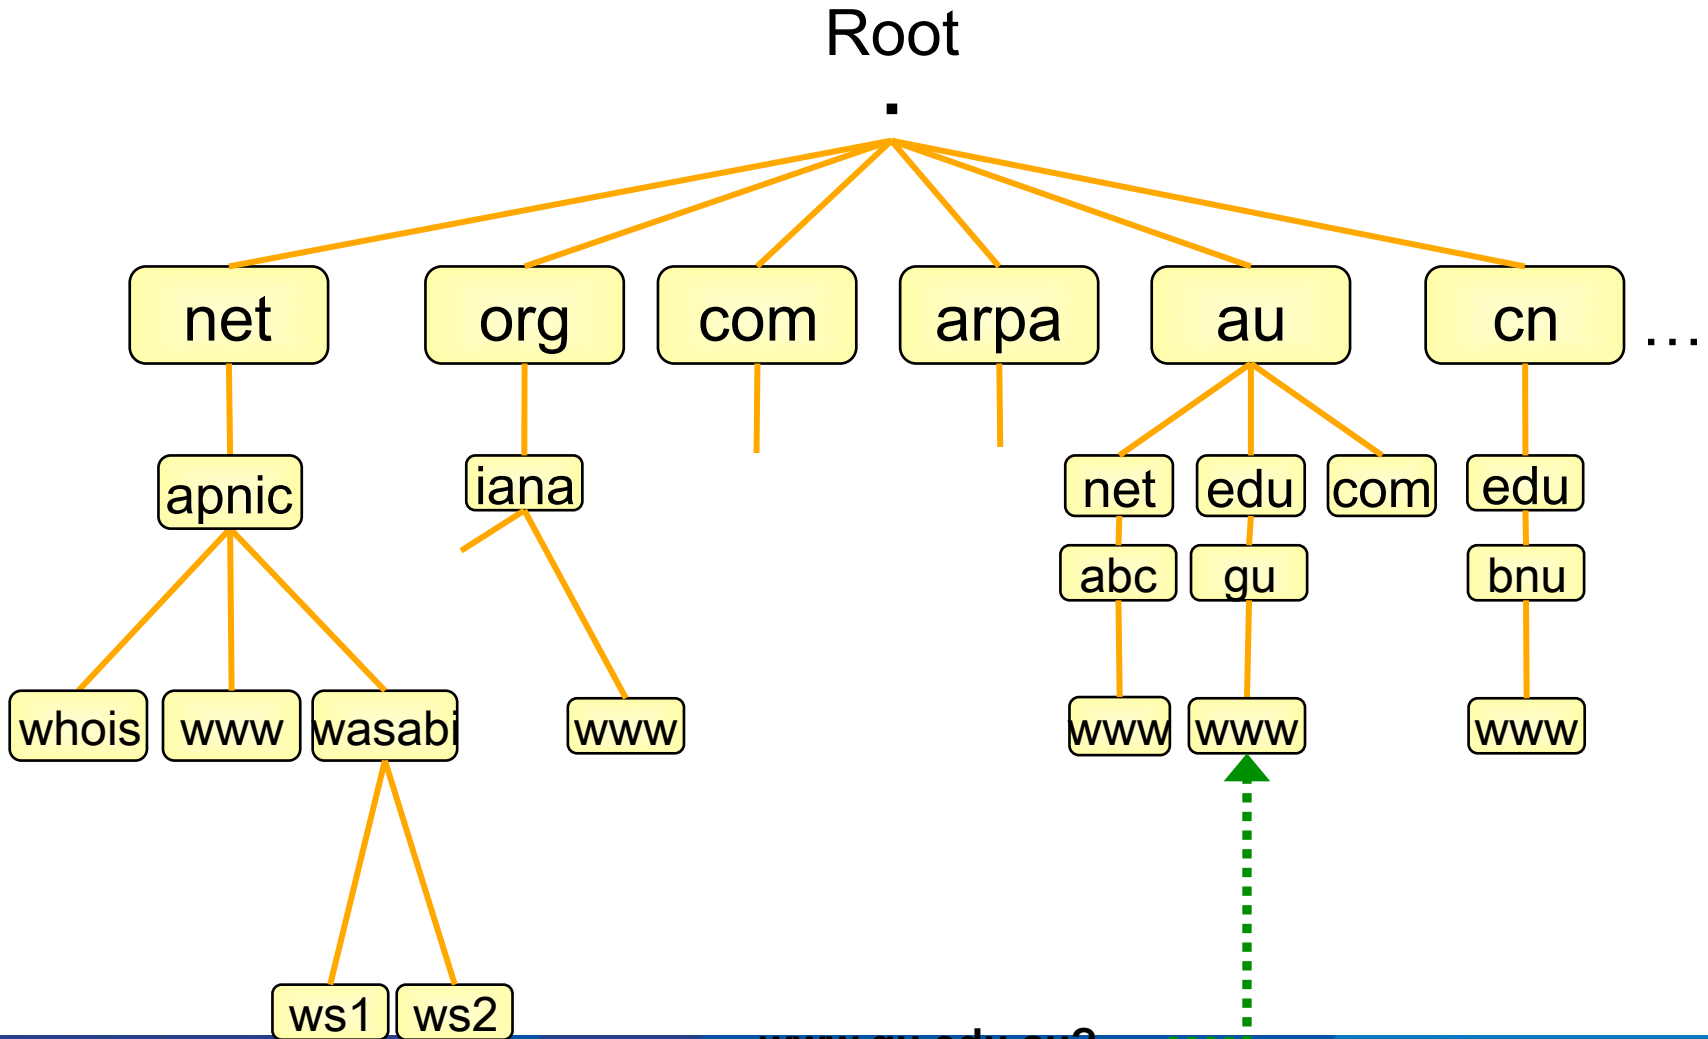
Internet Routing



IP Addresses vs Domain Names



The DNS tree



www.gu.edu.au?

Managing Internet Resources

Internet Resource Management Objectives

Conservation

- Efficient use of resources
- Based on demonstrated need

Aggregation

- Limit routing table growth
- Support provider-based routing


Registration

- Ensure uniqueness
- Facilitate trouble shooting

Uniqueness, fairness and consistency

IPv4 Address Space

STATUS OF 256 /8s IPv4 ADDRESS SPACE




Source: NRO

APNIC



IPv6 Address Space

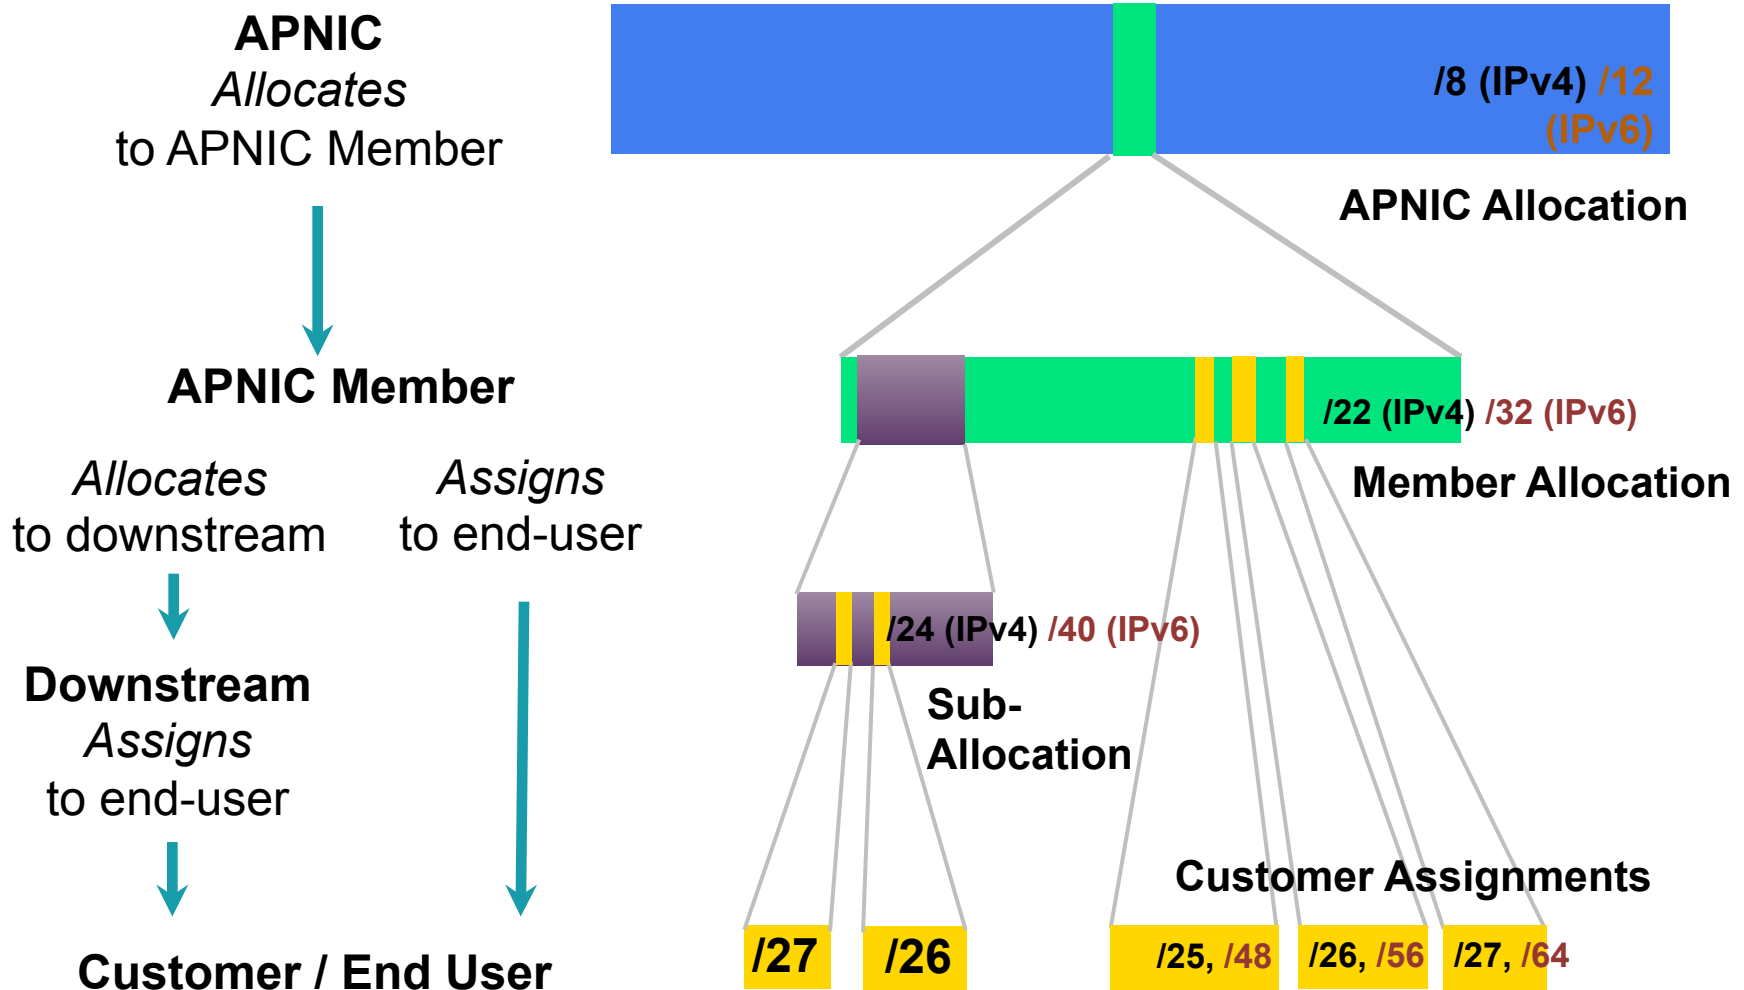


Source - NRO

Allocation and Assignment

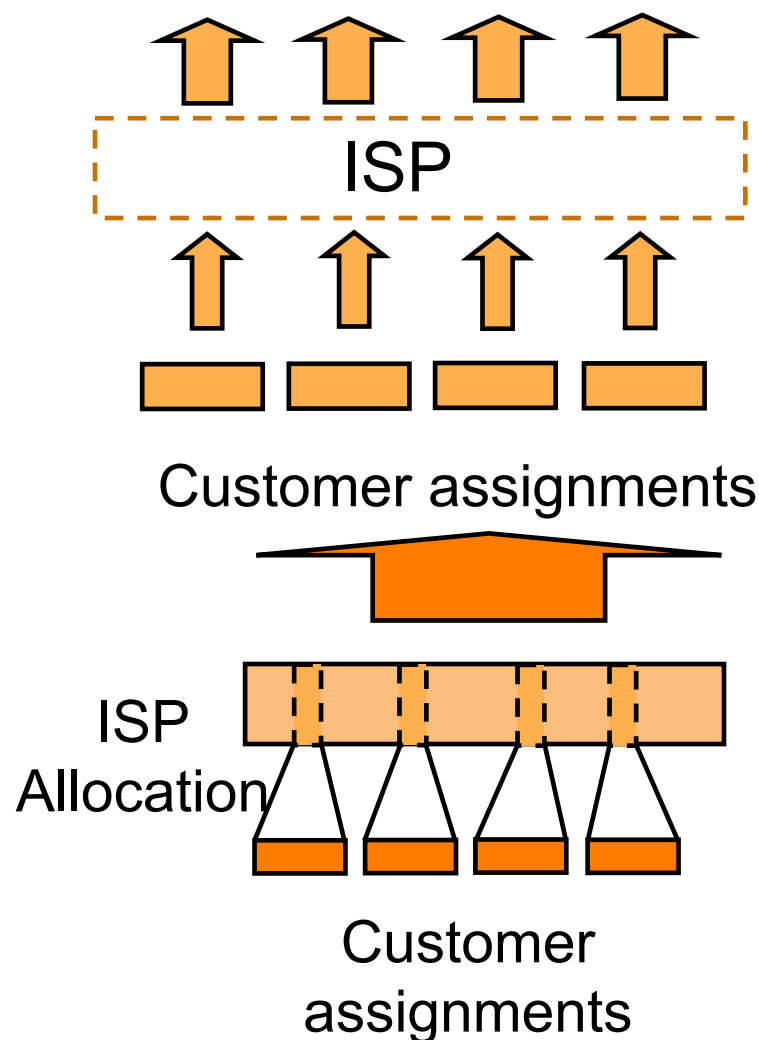
- Allocation
 - “A block of address space held by an IR (or downstream ISP) for subsequent allocation or assignment”
 - Not yet used to address any networks
- Assignment
 - “A block of address space used to address an operational network”
 - May be provided to ISP customers, or used for an ISP’s infrastructure (‘self-assignment’)

Allocation and Assignment




Portable and Non-Portable

- Portable Assignments
 - Customer addresses independent from ISP
 - Keeps addresses when changing ISP
 - Bad for size of routing tables
 - Bad for QoS: routes may be filtered, flap-dampened
- Non-portable Assignments
 - Customer uses ISP's address space
 - Must renumber if changing ISP
 - Only way to effectively scale the Internet
- Portable allocations
 - Allocations made by APNIC/NIRs



Address Management Hierarchy



Describes “portability” of the address space

Internet Resource Management Objectives

Conservation

- Efficient use of resources
- Based on demonstrated need

Aggregation


- Limit routing table growth
- Support provider-based routing

Registration


- Ensure uniqueness
- Facilitate trouble shooting

Uniqueness, fairness and consistency

Growth of the Global Routing Table




Growth of the Global Routing Table – IPv6




Routing Fundamentals

Graphics / Symbols Used



Router
(layer 3, IP datagram forwarding)



Network Access Server
(layer 3, IP datagram forwarding)



Ethernet switch
(layer 2, packet forwarding)

What is a Routing Protocol?

- A set of rules defined to facilitate the exchanges of routing information between routers (Layer 3 device) inside networks
- Builds routing tables dynamically based on updates from its neighbours
 - Allows the router to find the best path in a network that has more than one path to a remote network.
- Maintains connectivity between devices within the network.

Routing Protocol Behaviour

- Updates Layer 3 routing devices, to route the data across the best path
- Learns participating routers advertised routes to discover their neighbors
- Learned routes are stored inside the routing table

What is Routing?

- Routing is the method of delivering an item from one location to another
 - Example Post Mail = delivery is being done via Post Office
- In a router network environment, it forwards traffic to a logical device destination interface.
- Routers perform two functions to deliver the packets to their destination:
 1. **Routing:** Learning the logical topology of the network to store the path inside the routing table to where the traffic should flow
 2. **Switching:** Forwarding the packets from an inbound interface to the outbound interface within the router

Distinction between *Routed* and *Routing* Protocols

- Routed protocols
 - Layer3 datagram that carry the information required in transporting the data across the network
- Routing protocols
 - Handles the updating requirement of the routers within the network for determining the path of the datagram across the network

Routing and Routed Protocols

Routed protocol	Routing protocol
AppleTalk	RTMP, AURP, EIGRP
IPX	RIP, NLSP, EIGRP
Vines	RTP
DecNet IV	DecNet
IP	RIPv2, OSPF, IS-IS, BGP and (Cisco Systems proprietary) EIGRP,

Routing Requirements

- Activation of the protocol suite from such devices participating in the network
- Knowledge of the network destination
 - Must have an available entry in the routing table
 - Must have a valid and current route entry
- Interface presenting the best route path
 - Outbound interface with the lowest metric path

Routing Information

A routing table entry must contain the following information:

- Network field
- Outgoing interface
- Metric field
- Next-hop field

Network Field

- Contains information of entries
 - Networks learned (destination logical network or subnets)
 - Manually (static or default routes)
 - Dynamically (learned from routing protocol as dynamic routes)
- Information recorded is the entry on where to forward traffic to its destination when the datagram is received.

Outgoing Interface Field

- Interface to where the router sends the datagram
- Informs the administrator of interface where the update came through

Metric Field

- Used to determine which path to use if there are multiple paths to the remote network
- Provides the value to select the best path
- But take note of the administrative distance selection process 😊

Routing Protocol Metrics

Routing protocol	Metric
RIPv2	Hop count
EIGRP	Bandwidth, delay, load, reliability, MTU
OSPF	Cost (the higher the bandwidth indicates a lowest cost)
IS-IS	Cost

Administrative Distance

- The method used for selection of route priority of IP routing protocol. The lowest administrative distance is preferred.
 - Manually entered routes are preferred over dynamically learned routes
 - Static routes
 - Default routes
 - Dynamically learned routes depend on the routing protocol metric calculation algorithm. For default metric values, the smallest metric value is preferred.

Administrative Distance Chart (Cisco)

Route sources	Default distance
Connected interface	0
Static route out an interface	0
Static route to a next hop	1
External BGP	20
IGRP	100
OSPF	110
IS-IS	115
RIP v1, v2	120
EGP	140
Internal BGP	200
Unknown	255

Next Hop Field

- Contains the destination address of the next forwarding router
 - Address of the next hop (outgoing interface) usually within the same subnet
 - iBGP (exemption to the rule)
- Identifies the next hop so that the router can create the Layer2 frame with the destination address

Routing Table Sample (Cisco)

Cisco-router#sh ip route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set ??????????????



Sample only

202.41.143.0/24 is variably subnetted, 2 subnets, 2 masks

S 202.41.143.0/24 is directly connected, Null0

C 202.41.143.17/32 is directly connected, Loopback0

O E2 10.110.0.0 [110/5] via 10.119.254.6, 0:01:00, Ethernet2

O 10.67.10.0 [110/128] via 10.119.254.244, 0:02:22, Ethernet2

B 217.170.115.0/24 [20/0] via 12.123.29.249, 5d16h

Routing Table Updates

- Routing table entry accuracy is required to make sure of the following:
 - Table entries are current and correct
 - New networks are inserted into the table
 - Best path is available to reach the destination network
 - Alternative routes are available to reach the destination network
 - Networks that is no longer available should not be seen in the routing table
- Depends on the routing protocol

Routing Decisions

- The main goal of the routing decision is to maintain a path that is valid and free from routing loops to the destination network, regardless of whether it is single hop or a multiple hop path
- The decision is made based on the metric value in the routing table
 - Using the sum of the metrics associated with the default routing protocol value and the intermediate connections

Router Traffic Forwarding

Forwards traffic to the out-bound interface

- Routing table entry ensures that the network topology is learnt
- Routing table entries that contain the information of the routes learned from other routers

Classful and Classless Distinction

Types of Routing Protocols

- Routing protocols are essentially applications inside the router designed to ensure correct and timely exchange of information within the network
- The IP routing protocol has several distinctions which can be divided into different groups
 - The first is the group is how the routing protocol handles the subnet mask sent during the routing update
 - The early routing protocols don't support this but the newer one's like RIP2, OSPF, ISIS, BGP4 support it. The distinctions are called “**classful** and **classless**”

Classful routing protocol (obsolete)

- Periodic updates are done by the routing protocol
 - does not carry out subnet or routing masks because the assumption is always based on network bit boundaries
 - does not support VLSM which makes it inefficient for addressing the network
- This has been **obsolete** for a long time but for knowledge purposes there are two protocols designed for it.
 - RIPv1
 - IGRP
- It has created constraints to IP network design due to its limitations

Classless Routing Protocol

- Classless routing protocol was designed to overcome the constraints from classful routing
 - OSPF
 - IS-IS
 - EIGRP (Cisco)
 - BGP
 - lead also to the development of RIP2

Characteristics of Classless Routing

- Support for different subnet mask values
 - Routers can be configured to have different subnet masks (VLSM)
- Supports route summarisation
 - Manual configuration
- Supports Classless Interdomain Routing (CIDR) architecture

Distance Vector Routing Protocol

- Another distinction based on the behavior of learning the path to the destination network
 - Distance vector and
 - Link-State protocols
- Distance vector routing protocol
 - Early technology of IP routing protocols (RIPv2)
 - Concept design was for small networks

Distance Vector Operation

- Maintains its own table by sending its own modified table for updates
- Sends updates to directly connected neighbors
 - is done in a periodic manner. This is commonly known as the (routing by rumor)
 - the timer needs to reach the expiration stage before the next update will be sent

NOTE: Updating of the tables affects the entire routing table, except those networks learned from interfaces where update is sent

APNIC



Link-State Routing Protocol

- Link-state routing protocol was designed to overcome the limitation of Distance Vector routing protocol. The main goal is to achieve the following:
 - Maintain a loop-free and accurate table
 - Utilises multicast address and make updated based on incremental
 - Fast convergence of the network
 - Reduce the network overhead during updates
 - Selection of best path based on link status
- Routing protocols that support link-state are:
 - OSPF
 - IS-IS

Link-state Operation

- Link state operates through its main concern focusing on the link connected to the router (not the routes)
- Changes in the link state is propagated to its neighboring routers
 - to maintain the same image of the network topology among each neighbors
- Uses the router's CPU resources instead of bandwidth like Distance Vector
- When there is a state change the *incremental update* is sent to neighboring routers immediately,
 - remains silent if there's no change in the link state.

Link-state Operation

- Maintains the topology map of the network for easy local network table updates via incremental updates
 - OSPF = Link State Advertisement (LSA)
 - Dijkstra algorithm used to compute the new topology map of the network
- Metric used in Link-state is stated as “*cost*”
 - Equipment vendors default value setting can be overridden in manual configuration

Link-State Comparison Chart

Characteristic	OSPF	IS-IS
Hierarchical topology	X	X
Retains knowledge of all possible routes	X	X
Routes summarisation - manual	X	X
Event-triggered announcement	X	X
Load balancing – equal path	X	X
VLSM support	X	X
Routing algorithm	Dijkstra	IS-IS
Metric	Cost	Cost
Hop Count limit	Unlimited	1024
Scalability	Large	Vry-Lrg

Interior and Exterior Routing Protocols

- Other group distinctions with routing protocols are simplified as the protocols used for internal or external networks
- Interior
 - Routing protocol used to maintain routes within the organisation
 - Routing protocols that support it are:
 - *RIPv2, OSPF, IS-IS, EIGRP (Cisco)*
- Exterior
 - Routing protocol used to maintain routes connecting to different organisations
 - Exchange routing information between organisations
 - Using Exterior Gateway Protocol (EGP)
 - Border Gateway Protocol version 4 (BGP-4)

Addressing Issues and Solutions

IP Addressing Issues

- Exhaustion of IPv4 addresses
 - Wasted address space in traditional subnetting
 - Limited availability of /8 subnet addresses
- Internet routing table growth
 - Size of the routing table due to higher number prefix announcement
- Tremendous growth of the Internet

IP Addressing Solutions

- Subnet masking and summarisation
 - Variable-length subnet mask definition
 - Hierarchical addressing
 - Classless InterDomain Routing (CIDR)
 - Routes summarisation (RFC 1518)
- Private address usage (RFC 1918)
 - Network address translation (NAT)
- Development of IPv6 addresses


Subnetting Overview

- Allows the creation of additional sub-networks by simply moving the network boundary to the *right*
- When the contiguous 1s are added, it indicates by how many bits the network portion will be extended
- The sub-network is calculated by the 2^n where “ n ” is the number of extended bits.

Addressing Hierarchy

- Support for easy troubleshooting, upgrades and manageability of networks
- Performance optimisation
 - Scalable and more stable
 - Less network resources overhead (CPU, memory, buffers, bandwidth)
- Faster routing convergence

Addressing Hierarchy Example




Addressing Hierarchical (cont.)

Network Number
192.168.0.0/16

Core
192.168.32.0/19

Distribution/Core
192.168.32.0/21

Access/Distribution
192.168.48.0/21




Variable Length Subnet Mask

- Allows the ability to have more than one subnet mask within a network
- Allows re-subnetting
 - create sub-subnet network addresses
- Increase the route capabilities
 - Addressing hierarchy
 - Summarisation

Calculating VLSM Example

- Subnet 192.168.0.0/24 into smaller subnet
 - Subnet mask with /27 and /30 (point-to-point)



Calculating VLSM Example (cont.)

- Subnet 192.168.0.0/24 into smaller subnets
 - Subnet mask with /30 (point-to-point)

Description	Decimal	Binary
Network Address	192.168.0.0/30	x.x.x.00000000
1 st valid IP	192.168.0.1/30	x.x.x.00000001
2 nd valid IP	192.168.0.2/30	x.x.x.00000010
Broadcast address	192.168.0.3/30	x.x.x.00000011

Calculating VLSM Example (cont.)

- Subnet 192.168.0.0/24 into smaller subnets
 - Subnet mask with /27

Description	Decimal	Binary
Network Address	192.168.0.32/27	x.x.x.00000000
Valid IP range 192.168.0.33 - 192.168.0.62		x.x.x.00000001
		x.x.x.00000010
Broadcast address	192.168.0.63/30	x.x.x.00011111

Calculating VLSM Example (cont.)

- Subnet 192.168.0.0/24 into smaller subnets
 - Subnet mask with /27

Description	Decimal	VSLM	Host	Host range
1 st subnet	192.168.0.0/27	x.x.x.000	00000	0-31
2 nd subnet	192.168.0.32/27	x.x.x.001		31-63
3 rd subnet	192.168.0.64/27	x.x.x.010		64-95
4 th subnet	192.168.0.96/27	x.x.x.011		96-127

$n = 5$ (n is the remaining subnet bits)

$2^n - 2 = 30$ host per subnet


Summarisation of Routes

Route Summarisation

- Allows the presentation of a series of networks in a single summary address.
- Advantages of summarisation
 - Faster convergence
 - Reducing the size of the routing table
 - Simplification
 - Hiding Network Changes
 - Isolate topology changes

Summarisation Example

- Router C summarises its networks (2 x/24) before announcing to its neighbors (routers B and D)
- Router A combined the networks received from B, C, D and announce it as single /16 routing to Internet



Route summarisation

- Subnet 192.168.0.0/24 and 192.168.1.0/24 combining then to become a bigger block of address “/23”


Network	Subnet Mask	Binary
192.168.0.0	255.255.255.0	x.x.00000000.x
192.168.1.0	255.255.255.0	x.x.00000001.x
Summary	192.168.0.0/23	x.x.00000010.x
192.168.0.0	255.255.252.0	x.x.00000010.x

Discontiguous Networks

- A network not using routing protocols that support VLSM creates problems
 - Router will not know where to send the traffic
 - Creates routing loops or duplication
- Summarisation is not advisable for networks that are discontiguous
 - Turn off summarisation
 - Alternative solution but understand the scaling limitation
 - Find ways to re-address the network
 - Can create disastrous situation


CIDR Solution Advantage

- CIDR offers the advantage of reducing the routing table size of the network by summarising the ISP announcement into a single /21 advertisement





A day in a life of a router

- find path
- forward packet, forward packet, forward packet, forward packet...
- find alternate path
- forward packet, forward packet, forward packet, forward packet...
- repeat until powered off



Routing versus Forwarding

- Routing = building maps and giving directions
- Forwarding = moving packets between interfaces according to the “directions”



IP Routing – finding the path

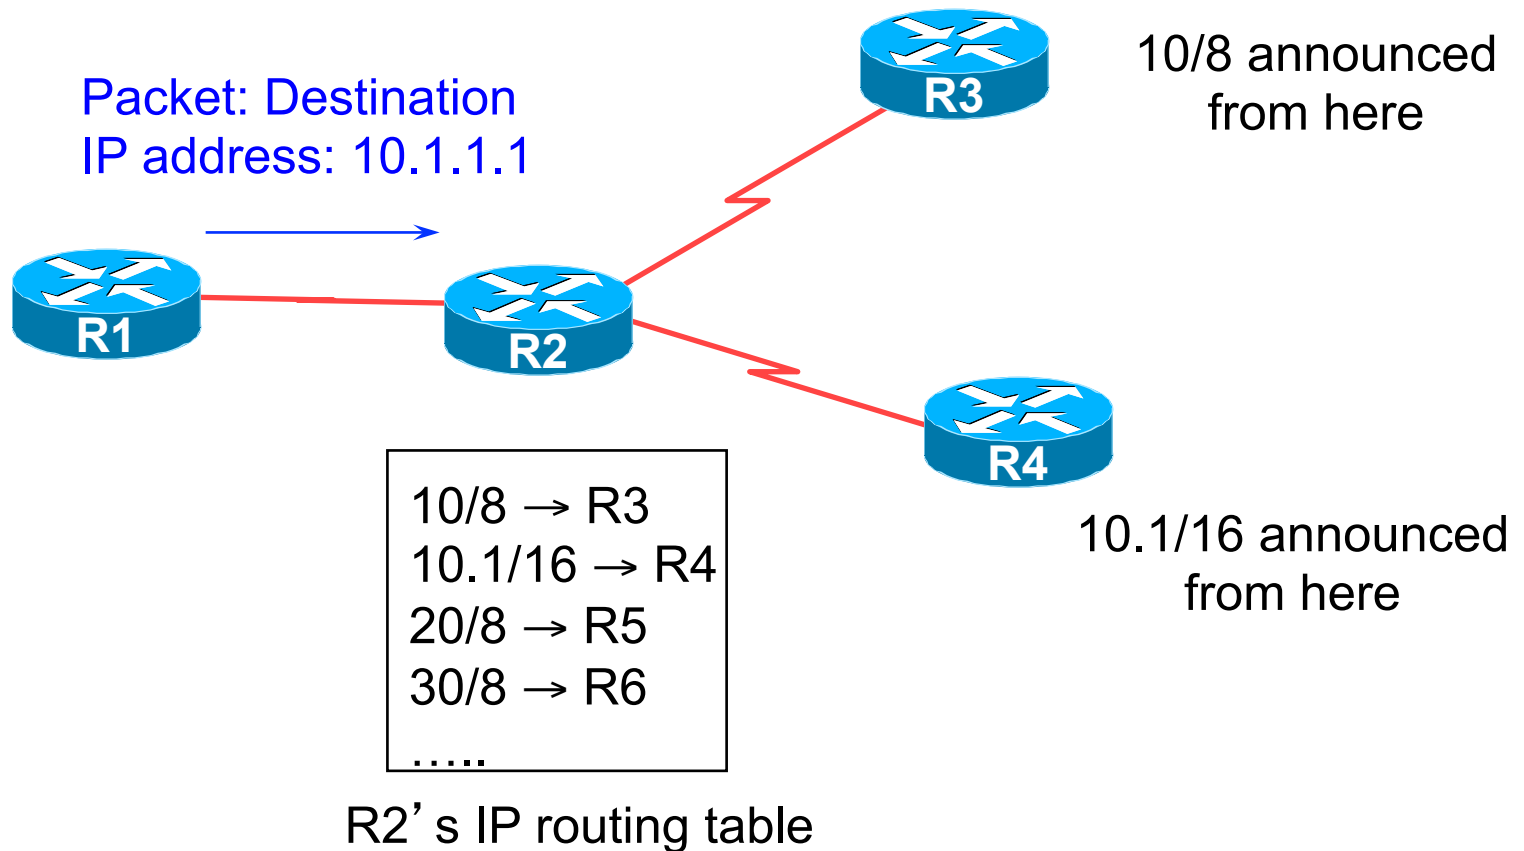
- Path derived from information received from a routing protocol
- Several alternative paths may exist
 - best path stored in forwarding table
- Decisions are updated periodically or as topology changes (event driven)
- Decisions are based on:
 - topology, policies and metrics (hop count, filtering, delay, bandwidth, etc.)

IP route lookup

- Based on destination IP address
- “longest match” routing
 - More specific prefix preferred over less specific prefix
 - **Example:** packet with destination of 10.1.1.1/32 is sent to the router announcing 10.1/16 rather than the router announcing 10/8.

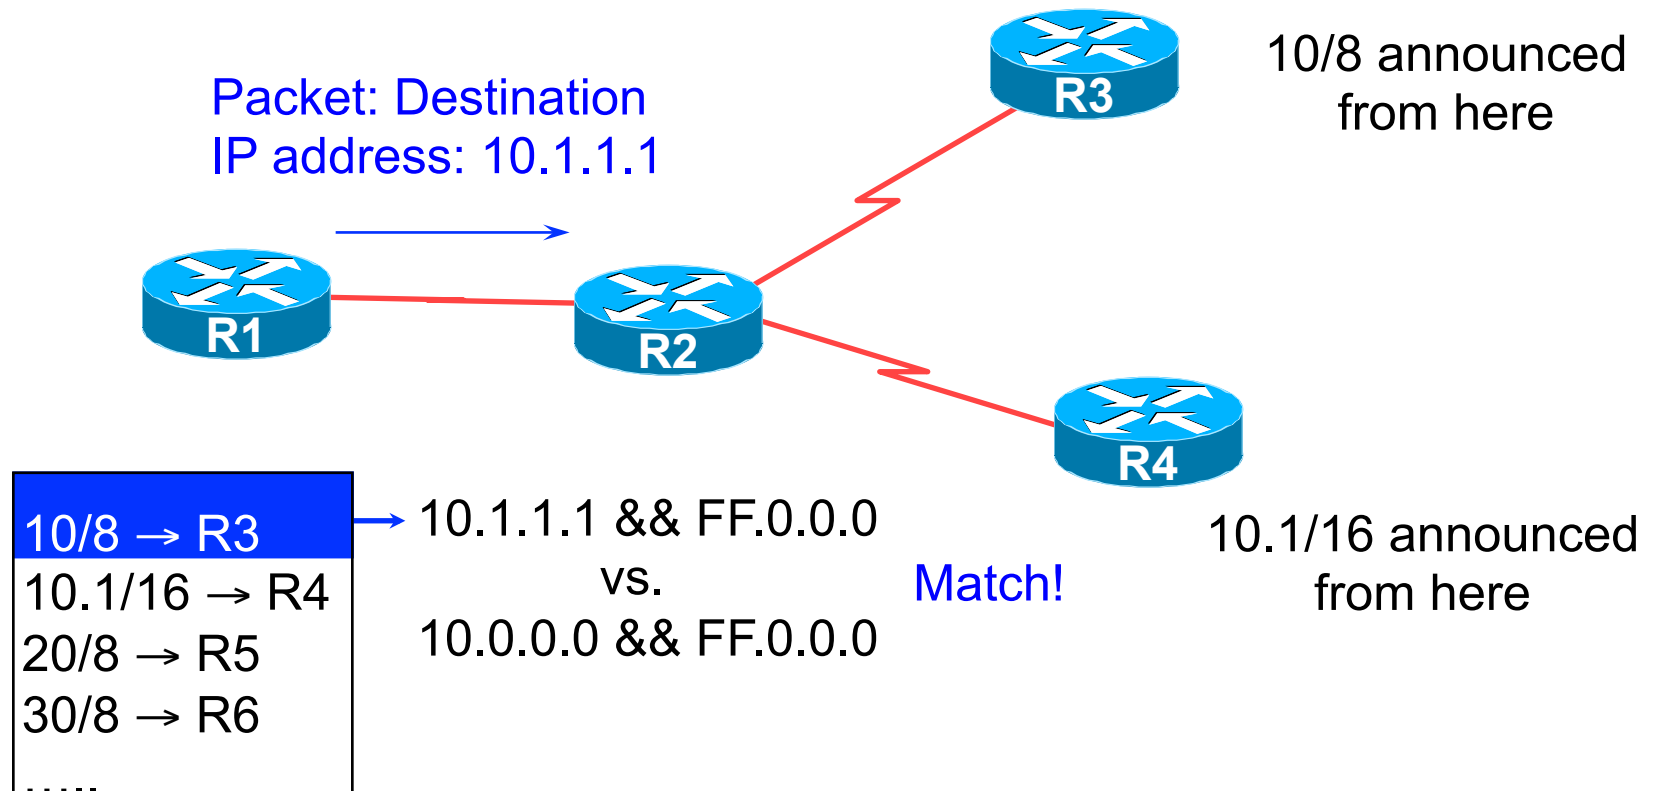
IP route lookup

- Based on destination IP address



IP route lookup: Longest match routing


- Based on destination IP address



R2's IP routing table

IP route lookup: Longest match routing


- Based on destination IP address



R2' s IP routing table

IP route lookup: Longest match routing

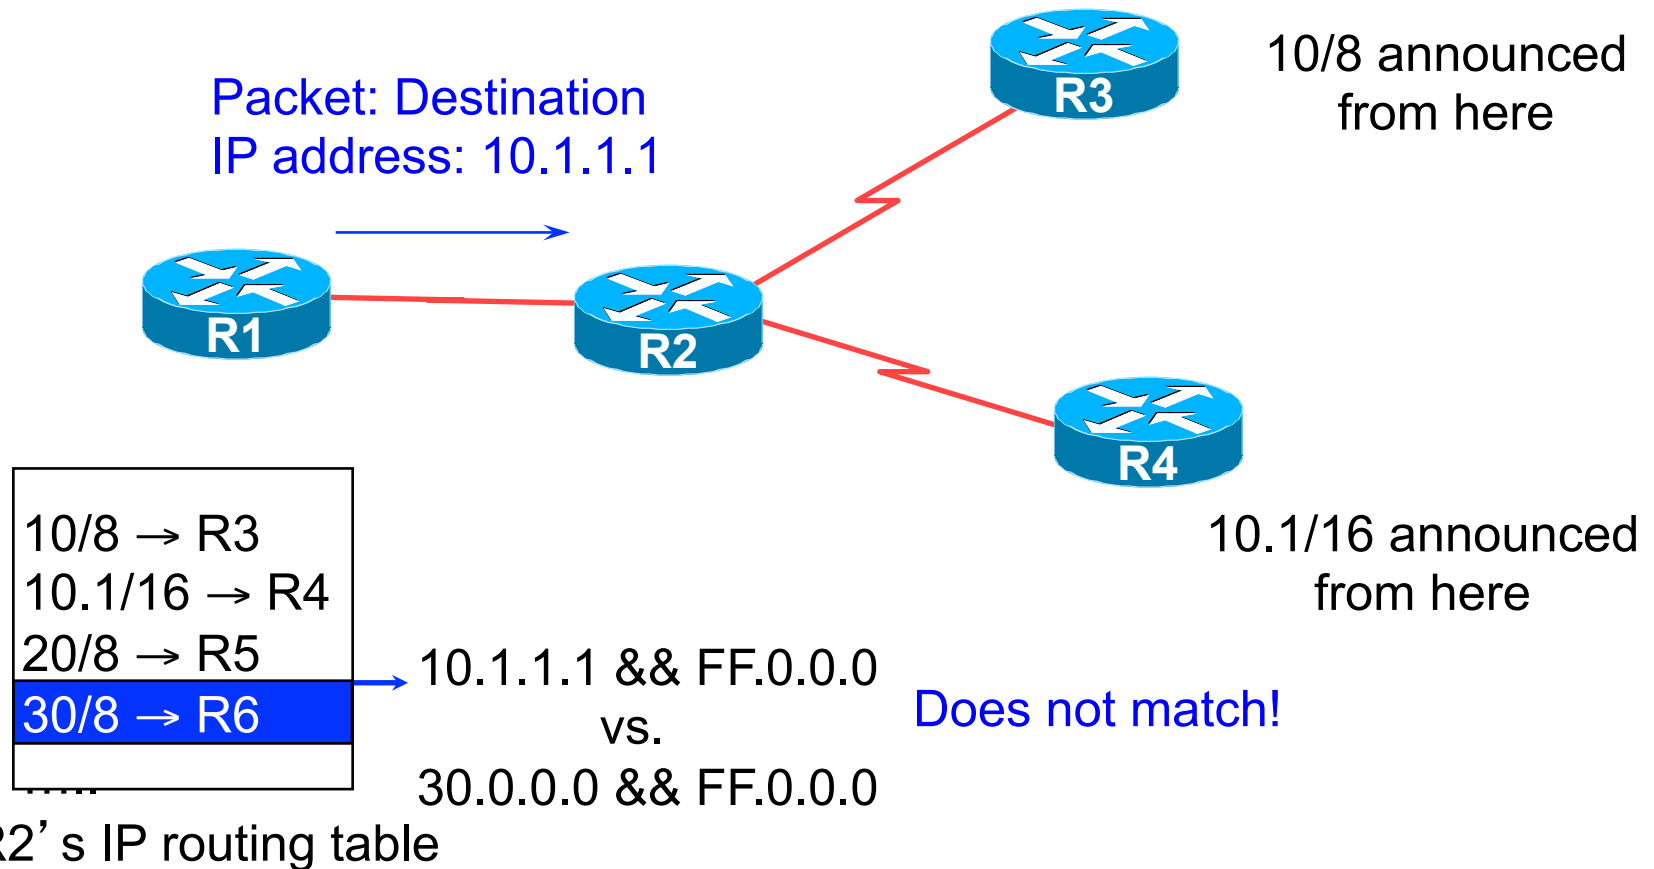
- Based on destination IP address



R2' s IP routing table

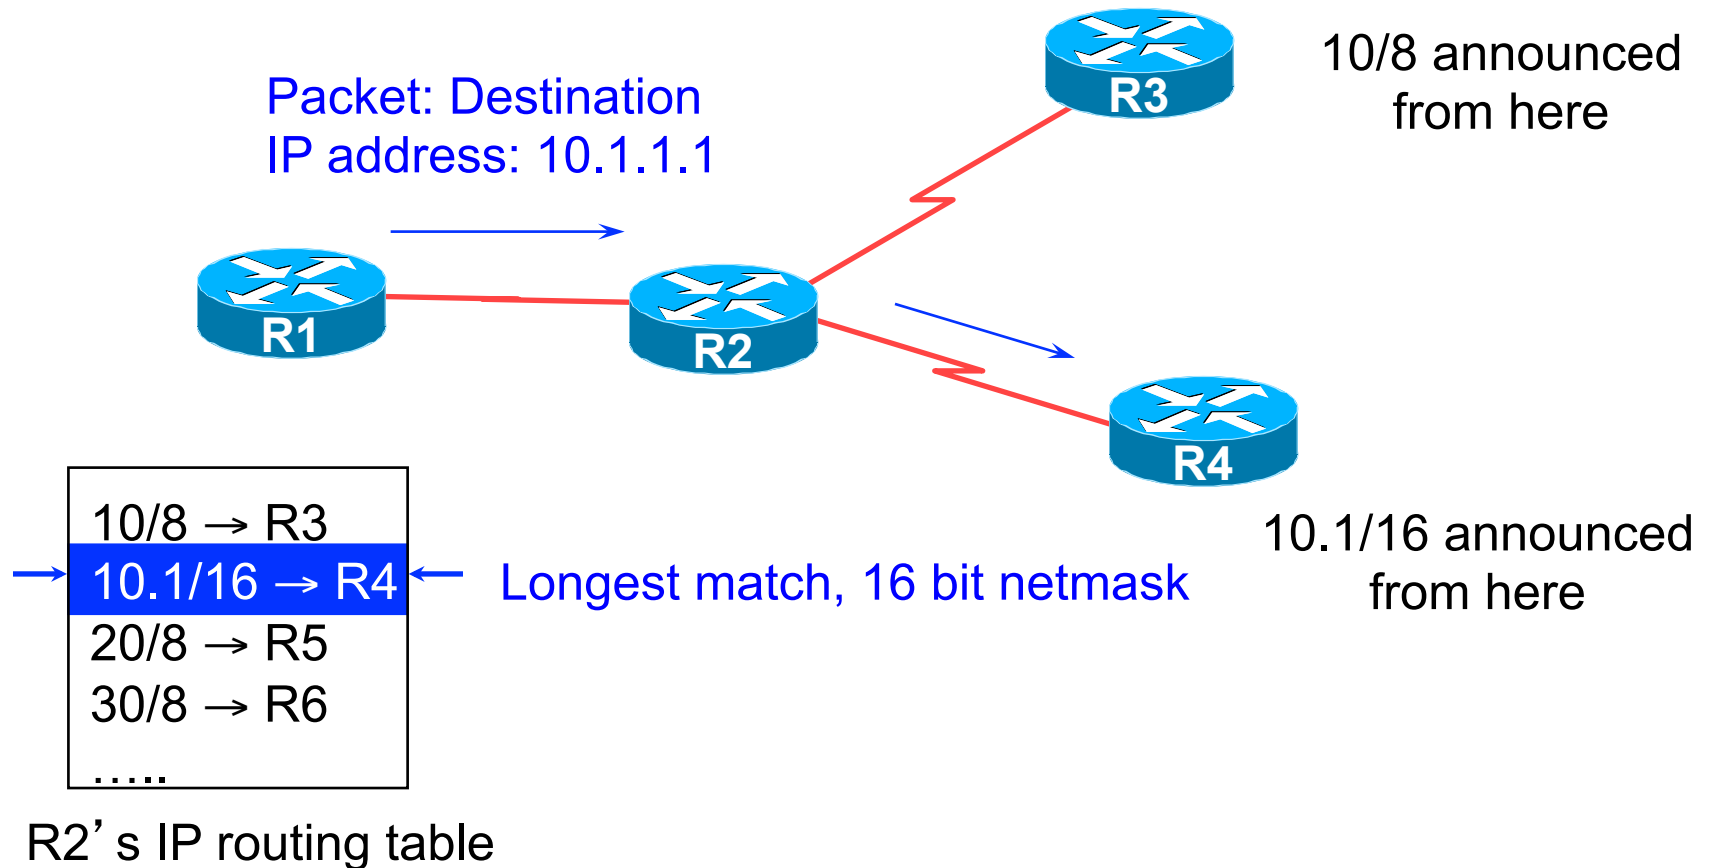
IP route lookup: Longest match routing

- Based on destination IP address



IP route lookup: Longest match routing

- Based on destination IP address



FYI: Cisco IOS Default Administrative Distances

Route Source	Default Distance
Connected Interface	0
Static Route	1
Enhanced IGRP Summary Route	5
External BGP	20
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
External Enhanced IGRP	170
Internal BGP	200
Unknown	255

Questions?