# Campus Networking Workshop
# CIS 399

Core Network Design

UNIVERSITY OF OREGON

# Routing Architectures

- Where do we route?
  - At the point where we want to limit our layer-2 broadcast domain
  - At your IP subnet boundary
  - We can create more complex topologies using routers and at the same time keep things sane

# Routing Architectures

- If we start with the right topology it will make our network more stable
- Use a hierarchical approach that makes good use of your traffic patterns and IP address allocations
- Be aware that topology and logical design are not the same

# Routing Architectures

- What is the right topology?
- Continue to think of three layers
  - Access
  - Distribution
  - Core
- Thinking of layers helps reduce convergence time because of the scope of information to process
- These layers should not be confused with your L2 architecture

# Routing Architectures

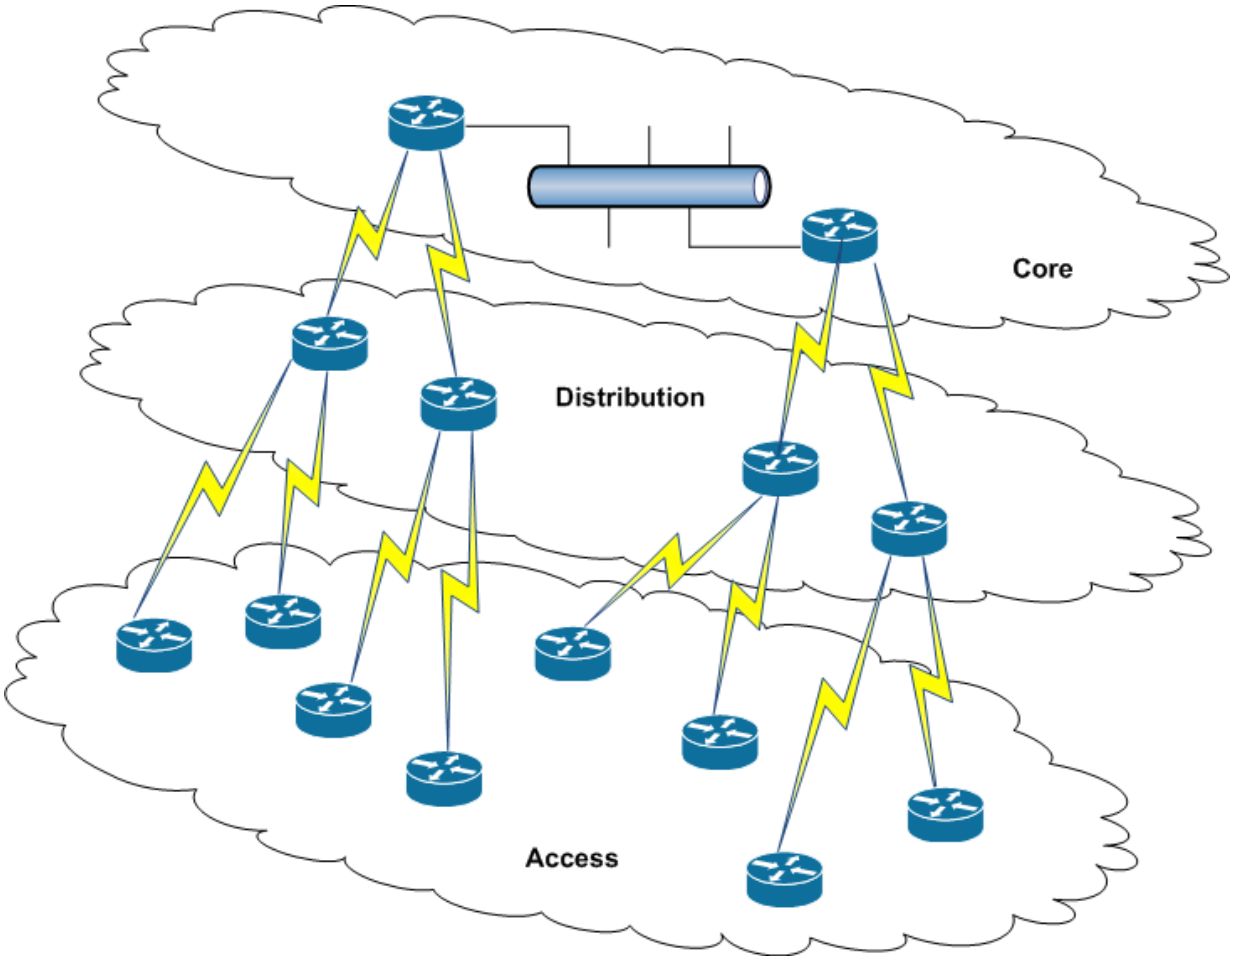

# Routing Architectures

- Access Layer
  - Minimum routing information
  - Feeds traffic into the network
    - Link sizing
  - Provides network access control
    - No spoofing
    - No broadcast sources
    - No directed broadcasts
  - Provides other edge services
    - Tagging for QoS
    - Tunnel termination
    - Traffic metering and accounting
    - Policy-based routing

# Routing Architectures

- Distribution Layer
  - Goals
    - Isolates topology changes
    - Controls the routing table size
    - Aggregates traffic
  - Strategies
    - Route summarization
    - Minimize the number of connections to the core

# Routing Architectures

- Core Layer
  - Goal
    - Forwarding packets fast
  - Strategies
    - Clear of network policies
    - Every device has full reachability to every destination
      - Facilitates core redundancy
      - Reduces suboptimal routing
      - Prevents routing loops

# Routing Architectures

- Depending in how large your campus is you could use the typical hierarchical model or a subset
  - Two collapse core models
    1. Single router acts as the network core
       - All other routers in the distribution layer
    2. Single router acts as the network core
       - No distribution layer
       - All access layer routers connected to the core

# Routing Architectures

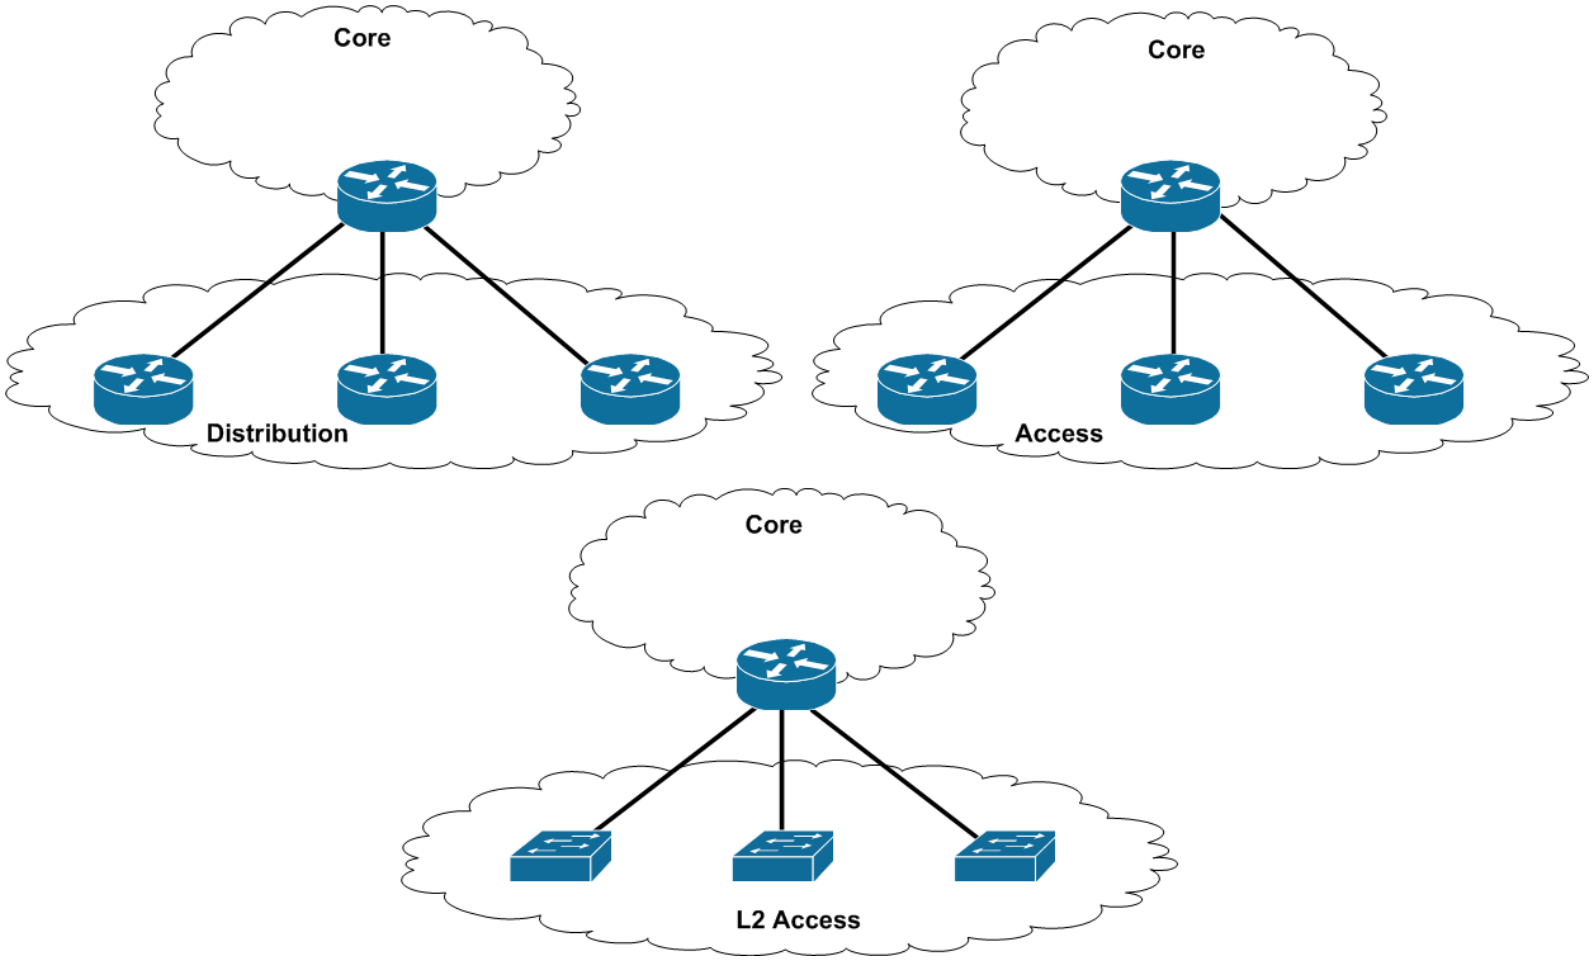

# Routing Architectures

- What to do about your address space
  - Assign it as you need it ….. WRONG!
    - Poor summarization has an impact on your network's stability
    - Very difficult to correct poor allocations
  - Spend some time thinking about how you will assign address space
    - Routing stability is affected by the number of routes propagated through your network

# Routing Architectures

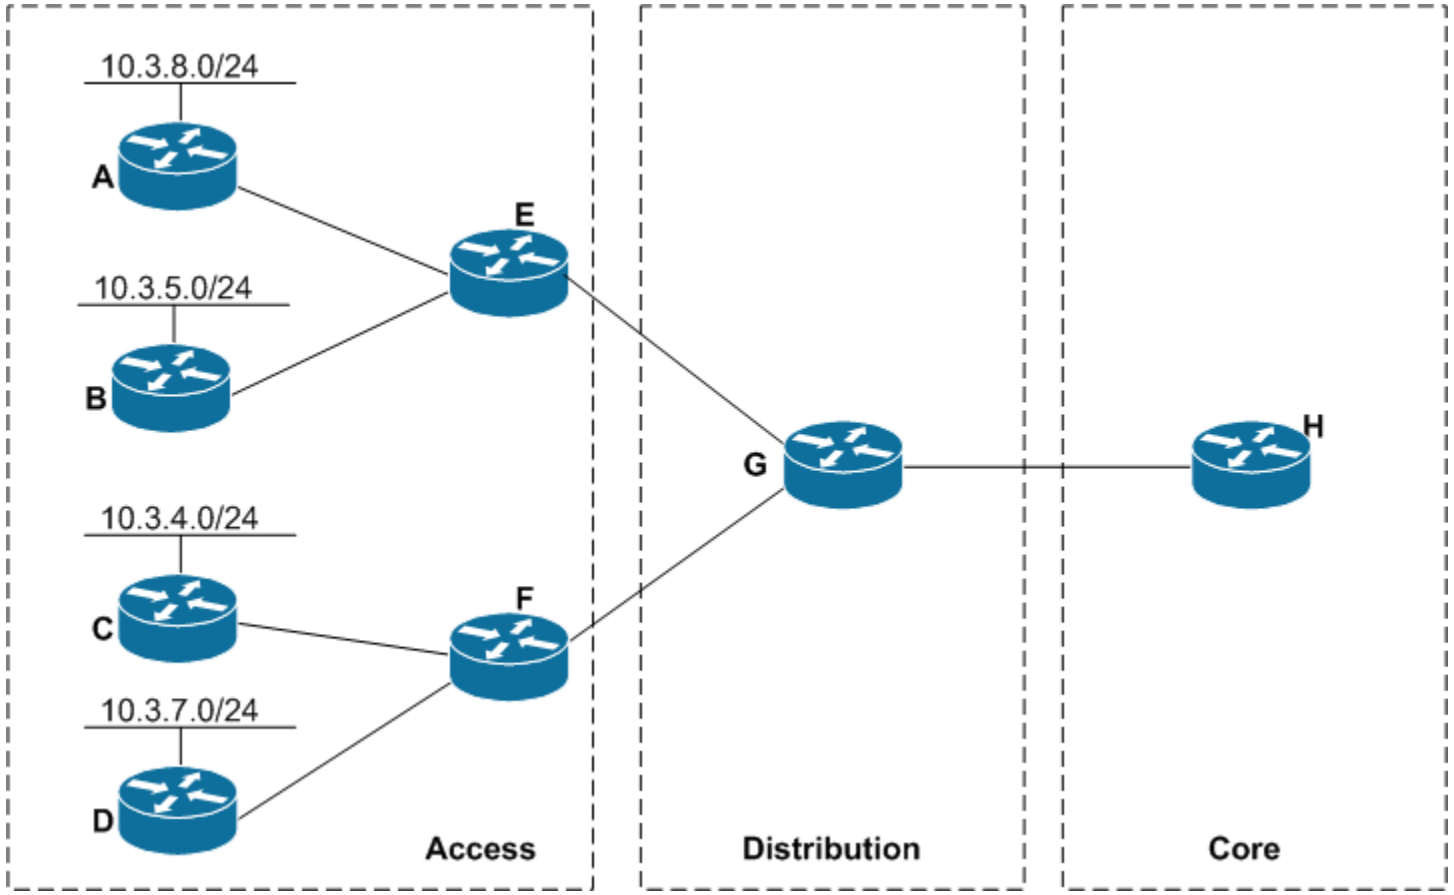

# Routing Architectures

- What happens if the link to router D fails?
- How is the distribution layer affected?
- How is the core layer affected?
- What changes can I make to my address allocation and address summarization to minimize the impact of a link failure on convergence time and network stability?

# Routing Architectures

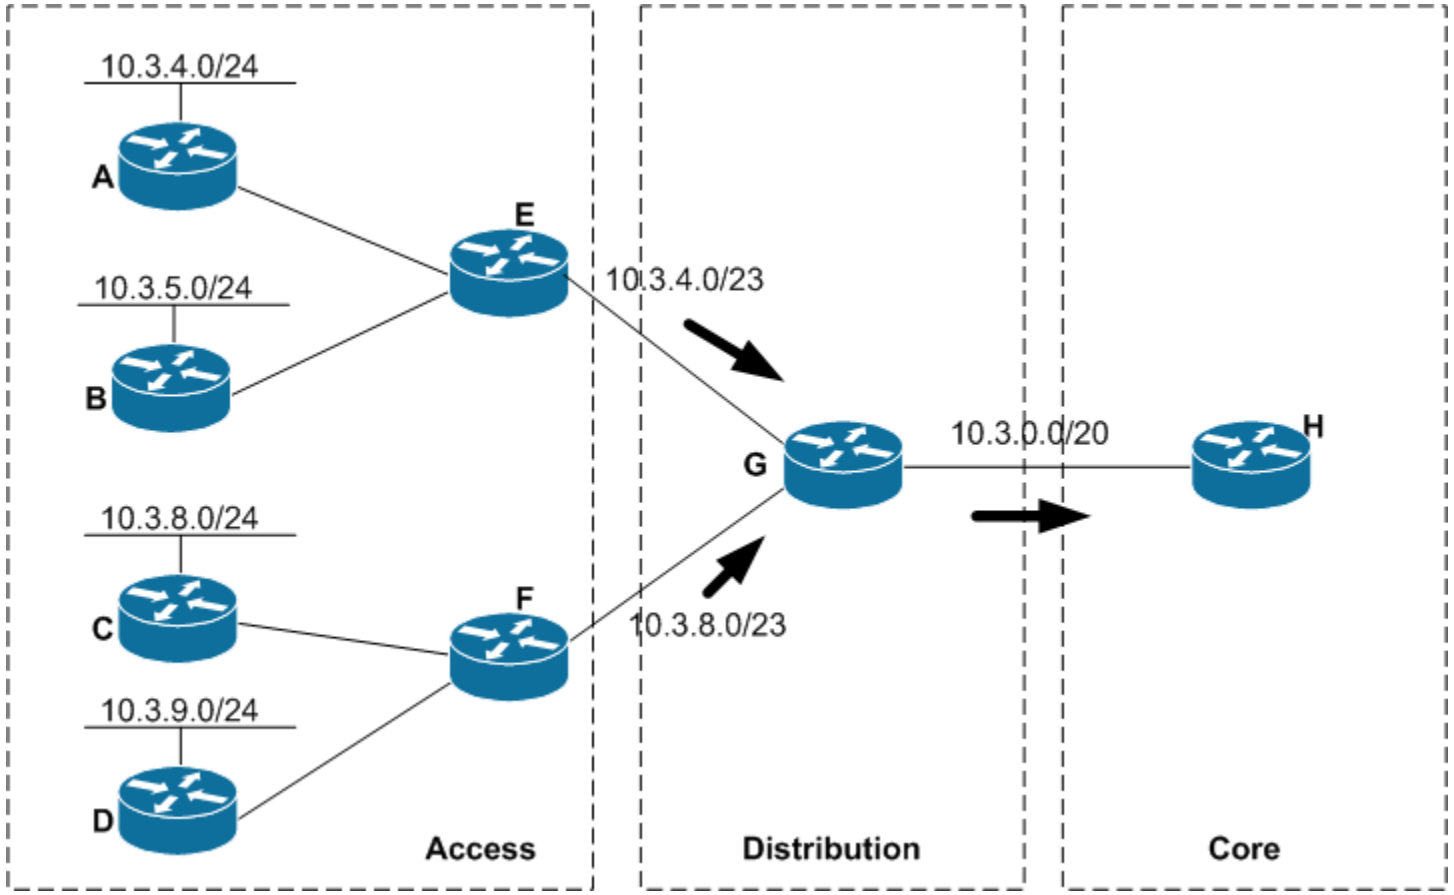

# Routing Architectures

- Where should you summarize?
  - Only provide full topology where it is needed
    - Core routers don't need to know about every single network
    - Access routers don't need to know how to get to every other network
      - They should only carry enough information to reach one (or a couple of) distribution router(s)
  - Summarize at the hierarchy edges
    - Distribution layer to core
    - Distribution layer to access

# Routing Architectures

- Strategies for Successful Addressing
  - First come, first serve
    - Start with a large pool and hand them out as needed
  - Politically
    - Divide the space so each group with in the organization have a pool of addresses available
  - Geographically
    - Divide the space so that every location has a pool of addresses available
  - Topologically
    - Assign addresses based on the point of attachment to the network (maybe same as geographically)

# Routing Architectures

- Addressing & Summarization
  - "Easy for you to say. I already have my network running and it looks nothing like what you show"
    - You are not alone
    - The principles still apply
    - Take it slowly. Define a goal and start working towards it. It can take years.
    - Maybe we can do the right thing with IPv6

# High Availability

- How can we achieve high availability?
  - Introduce hardware resiliency and backup paths into your network
  - Depending on the layer, you will use techniques differently
  - The idea is to protect your network against a single device failure affecting all of your network
  - Direct relationship between reliability, complexity and costs
  - The trick is to balance all variables and come up ahead

# High Availability

- You need to evaluate your needs
  - Minimal need
    - Network just needs to be up for a portion of the day
    - Downtime is easily schedule after working hours
    - Business is not impacted if the network is down
    - Users' productivity is not impacted by a network failure

# High Availability

– Medium need
  - Network needs to be available for most of the day
  - Only centralized servers need to be up 24 hours/ day
  - Downtime needs to be scheduled on weekends
  - If critical parts of the network fail, the business operation is impacted
  - A network failure affects user productivity

# High Availability

– High need
  - Network needs to be up 24x7
  - Downtime needs to be scheduled well in advance and completed within schedule
  - A network failure causes major loss of business
  - User productivity drastically impacted by a network failure

# High Availability

- Methods
  - Component Redundancy
    - Duplicate or backup parts
      - Power supplies, fans, processors, etc.
    - Have spares handy
  - Server Redundancy
    - Protect your data with backups
    - Use of hot standby servers
    - Or better yet use load balancers to distribute access
  - Network Link & Data Path Redundancy
    - Provide physical redundant connections between devices
    - Allow for hot backup paths (STP) and parallelism (routing)

# High Availability

- Core layer
  - Build a dual router core and provide dual paths to it from your distribution layer
    - These could be either L2 or L3 paths
  - Make sure that you have redundant power supplies in your devices
    - This also assumes two different sources of power
    - Think of UPS protected circuits
    - Maybe even a power inverter solution for emergencies
  - Think about the possibility of dual routing/forwarding engines
    - Weigh this against the use of two devices
    - Or just throw that in there as yet another layer of reliability

# High Availability

- Core layer
  - You want to also balance
    - Reduction of the hop count
    - Reduction of the available paths
    - Increase of the number of failures to withstand
  - Easy to do in a single location but complexity and costs directly proportional to the number and distance between the locations

# High Availability

- Distribution Layer
  - Provide dual connections to the core
    - Or provide a redundant link to other distribution layer devices
      - Doubles the core's routing table size
      - Possible use of the redundant path for traffic transiting the core
      - Preferring the redundant link to the core path
      - Routing information leaks
  - Allow for dual-homing of Access layer devices

# High Availability

- Distribution Layer:
  - Make sure that you have redundant power supplies in your devices
    - This also assumes two different sources of power
    - Think of UPS protected circuits
    - Maybe even a power inverter solution for emergencies
  - Think about the possibility of dual routing/forwarding engines
    - Weigh this against the use of two devices
    - Or just throw that in there as yet another layer of reliability
    - Increases the cost of the distribution layer

# High Availability

- Access Layer
  - Same challenges and solutions as the distribution layer
    - Dual home to the same distribution layer branch
      - Make sure to restrict destinations advertised to prevent transit traffic through the access layer router
    - Alternate path to another access layer device
      - Don’t use the redundant link for normal traffic
      - Make sure to restrict destinations advertised to prevent transit traffic through the access layer router

# High Availability

- Access Layer
  - Dual home to different distribution layer branches
    - Don’t use the redundant link for normal traffic
    - Make sure to restrict destinations advertised to prevent transit traffic through the access layer router

# High Availability

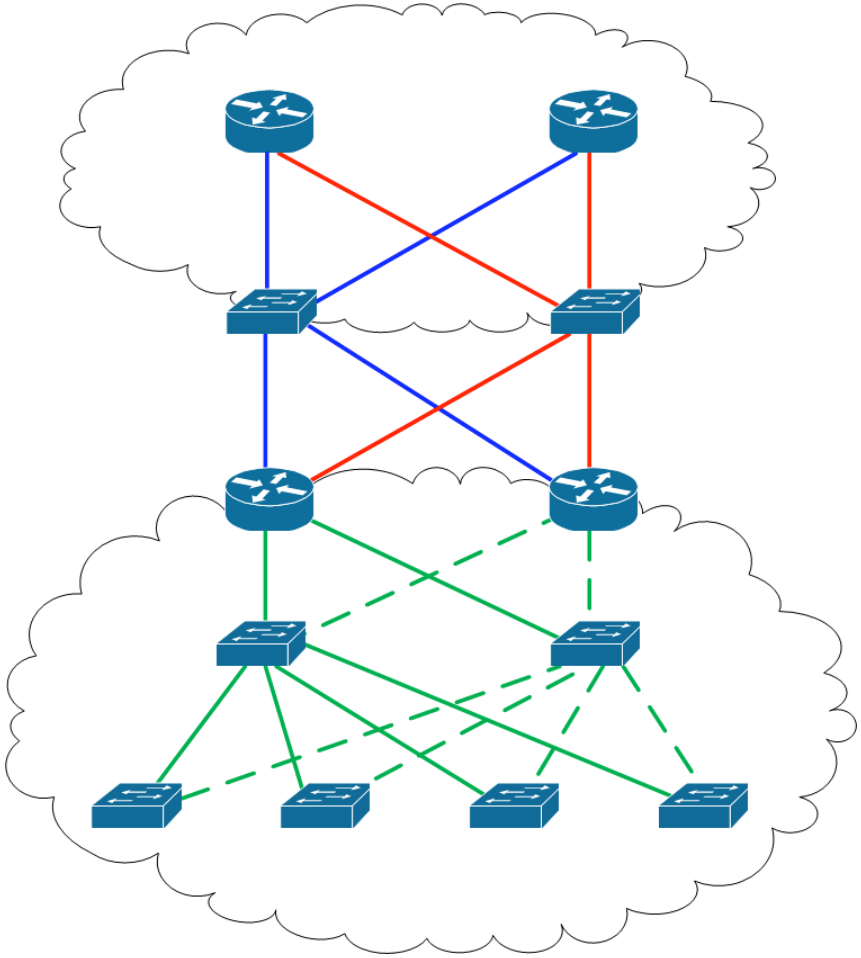

# High Availability

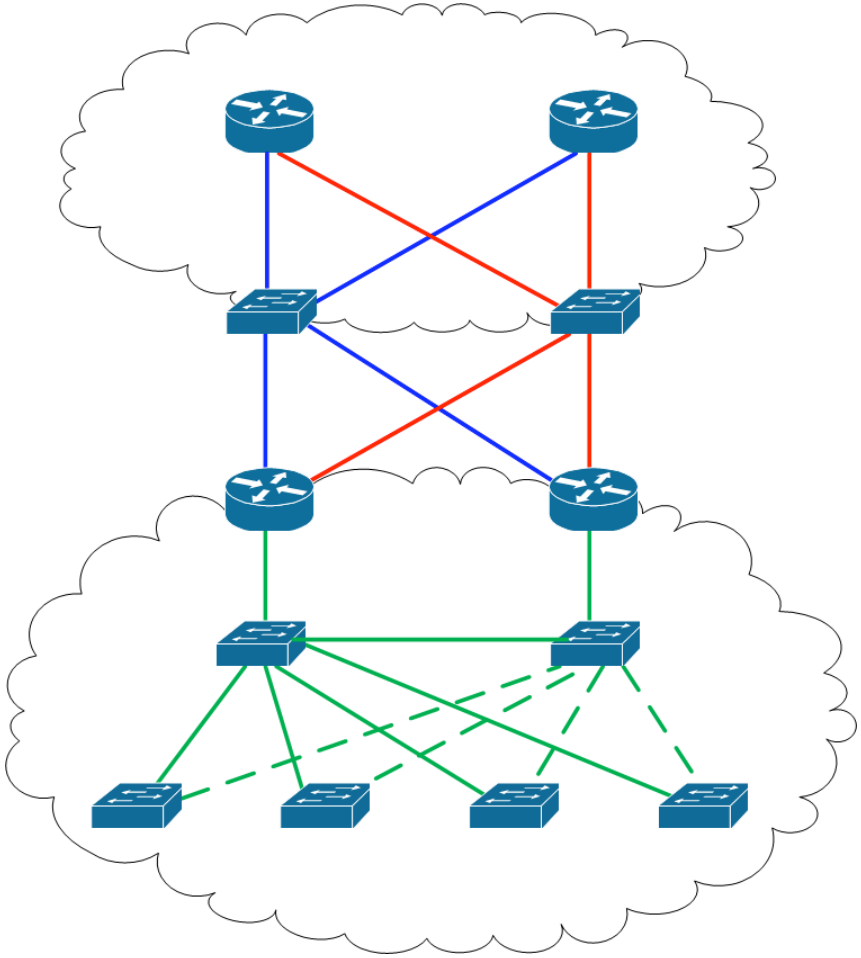

# High Availability

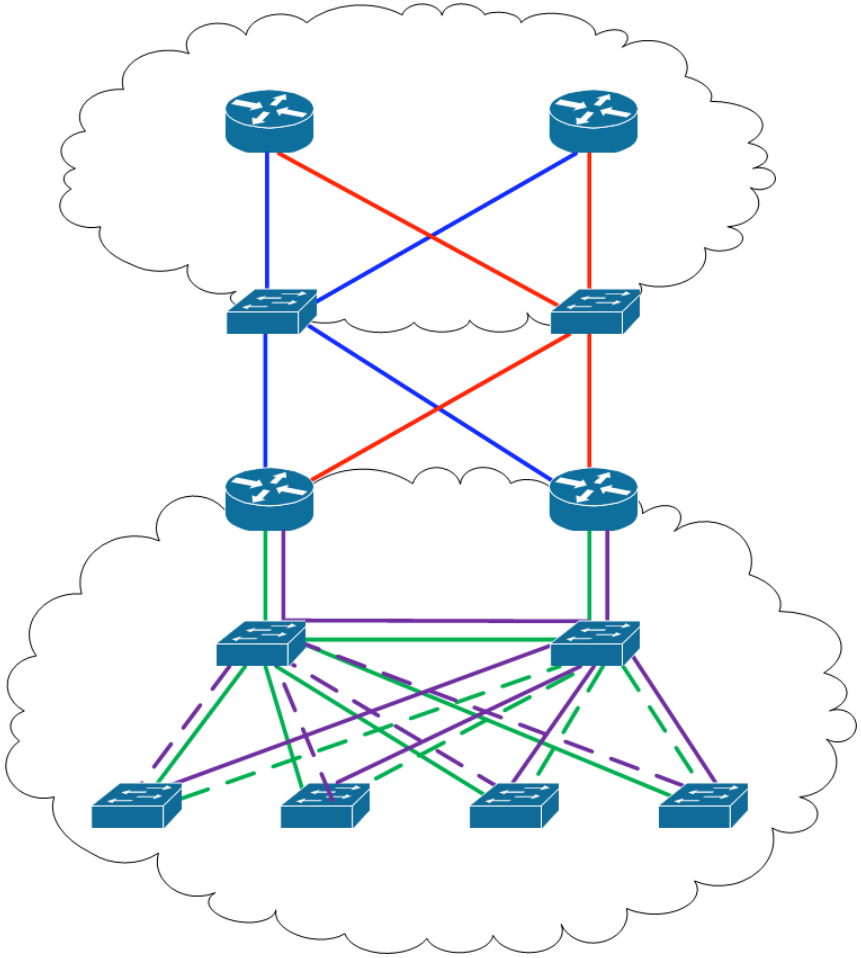

# High Availability

- So I built all this redundancy and high availability in my network, how can my end users take advantage of it?
- You are already providing more that one router for a segment
- You want to provide your users with a way to move their traffic from one default gateway to another

# High Availability

- If one of the routers fails the other one will continue to provide services to the segment
- Be aware that redundancy is not the same as load balancing

# High Availability

- How can we accomplish that?
  - Have the routers do proxy-ARP … Yikes!
  - Run a routing protocol between your workstations and the routers … Yikes!
  - Split your workstations into two groups
    - One uses one router as its default gateway
    - The other group uses the other router
  - Use ICMP Router Discovery Protocol (IRDP)
  - There is got to bit a better and simpler way to do this

# High Availability

- Current solutions:
  - Hot Standby Redundancy Protocol – HSRP (Cisco Proprietary, RFC2281)
  - Virtual Router Redundancy Protocol – VRRP (RFC3768)
  - Gateway Load Balancing Protocol – GLBP (Cisco Proprietary)

# High Availability

- The concept is very similar
  - Workstations get configured with a single default gateway
  - The routers in the segment will negotiate who will provide services to the workstations and keep track of the state of the other routers
  - In the event of a primary/active router failure, one of the standby routers will take over the task of forwarding traffic for the workstations and become the primary/active
  - Traffic to the workstations will go to the primary/active router
  - Incoming traffic into the segment will follow the routing decisions made by routers in the network

# High Availability

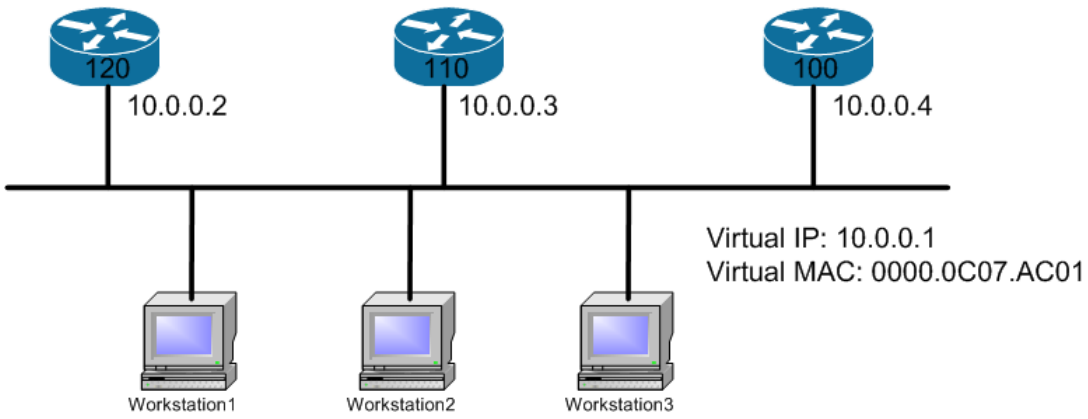

# High Availability

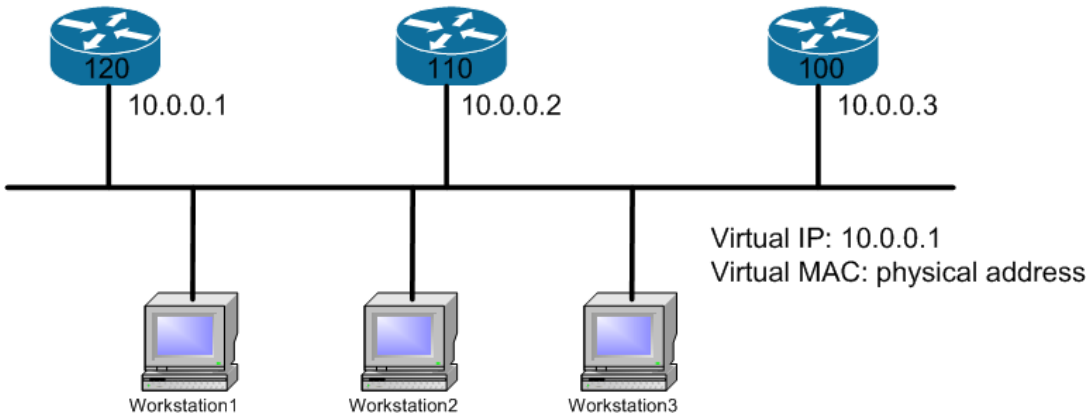

# High Availability

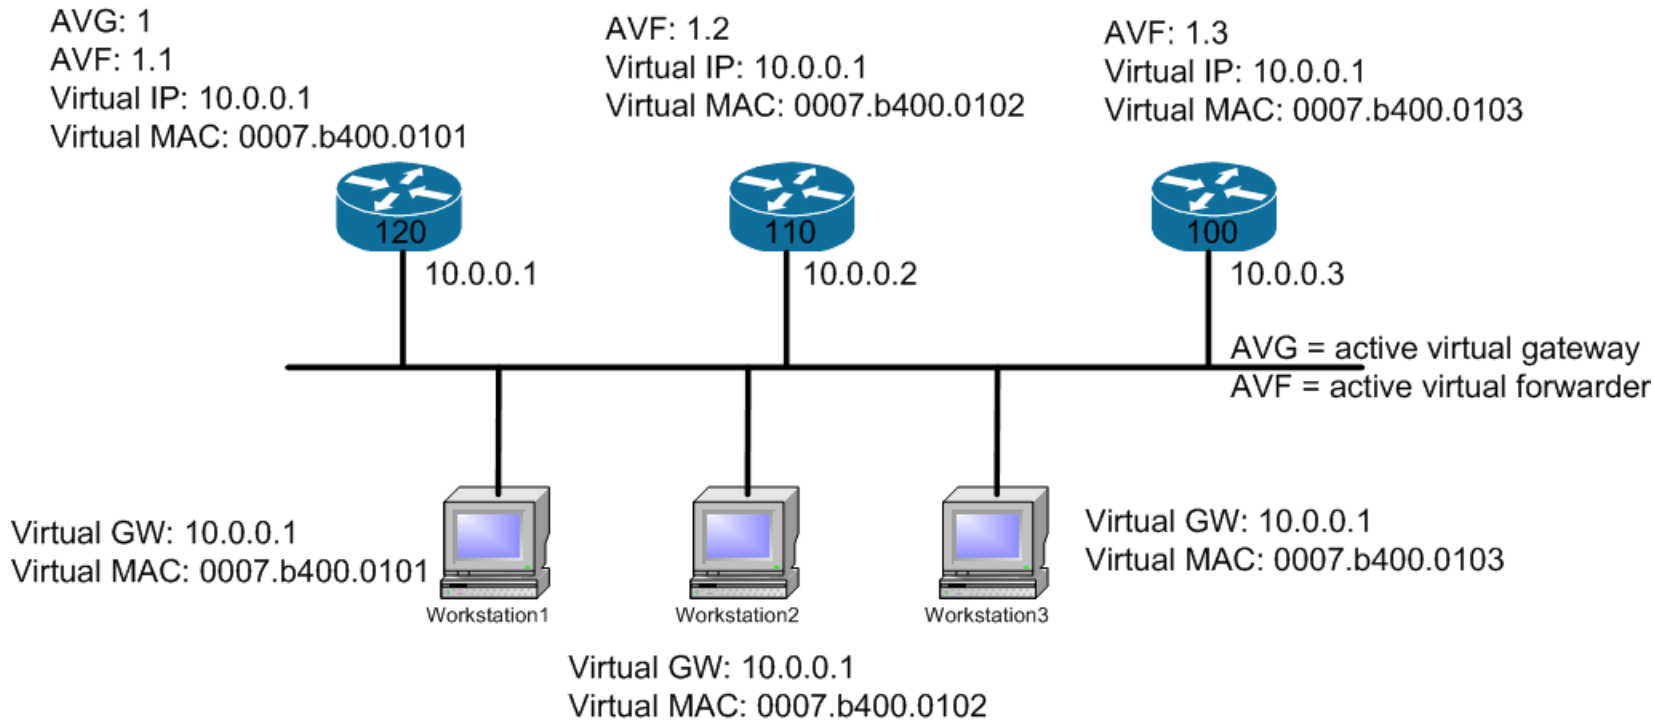

# High Availability

- Which one should I use?
  - They all allow for a common default gateway and MAC address
  - VRRP is standardized
    - HSRP/GLBP are Cisco proprietary
  - GLBP provides load balancing
    - HSRP/VRRP do not (without introducing complexity)
  - GLBP/HSRP can track an uplink interface
    - VRRP does not

# High Availability

- VRRP can reuse the default gateway IP
  - HSRP does not
- HSRP/GLBP support IPv6
  - VRRP does not yet
- VRRP uses protocol 112 & 224.0.0.18
  - HSRP uses UDP/1985 & 224.0.0.2
  - GLBP uses UDP/3222 & 224.0.0.102

# Routing Protocols

- So, now I know what my network is going to look like … or is that true?
- We need to figure out how packets will be forwarded.
- That is a function of the router and the routing protocols that we will implement
- There are many options
  - RIPv2/EIGRP/OSPF/IS-IS/BGP

# Routing Protocols

- Routing protocols can be classified in
  - Interior Gateway Protocols (IGP)
    - RIP, EIGRP, OSPF, IS-IS
    - We will talk about OSPF later on
  - Exterior Gateway Protocols (EGP)
    - BGP
    - We will talk about BGP later on

# Routing versus Forwarding

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

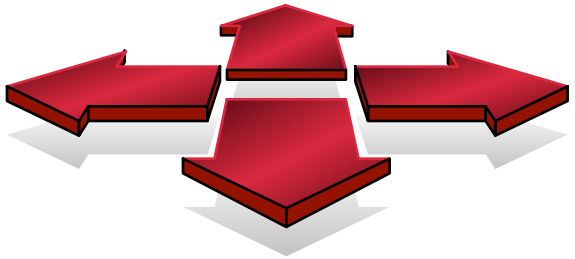

# IP Routing – finding the path

- Path derived from information received from a routing protocol
- Several alternative paths may exist
  - best next hop 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 packet
- “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 packet

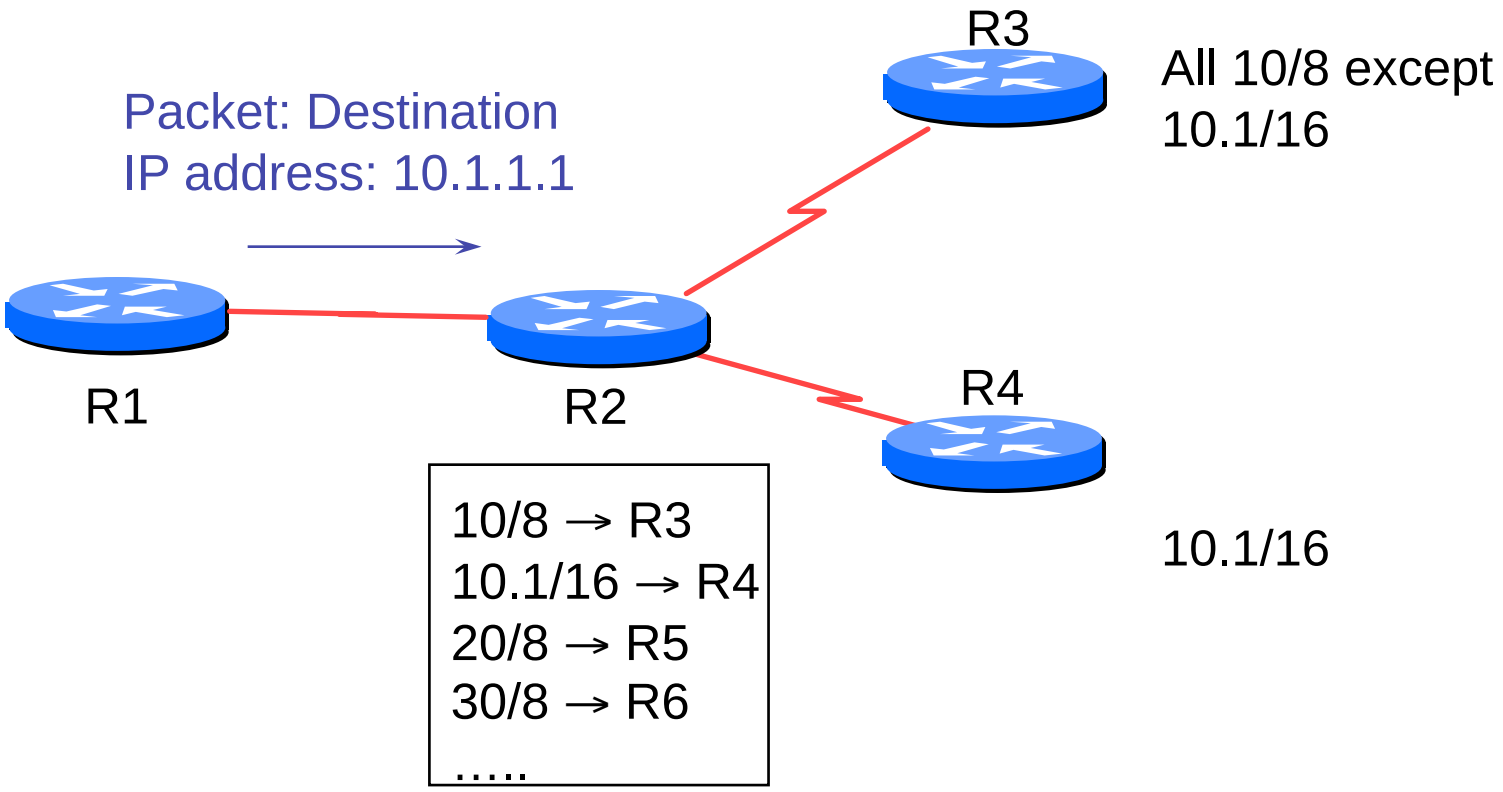

# IP route lookup: Longest match routing

- Based on destination IP packet

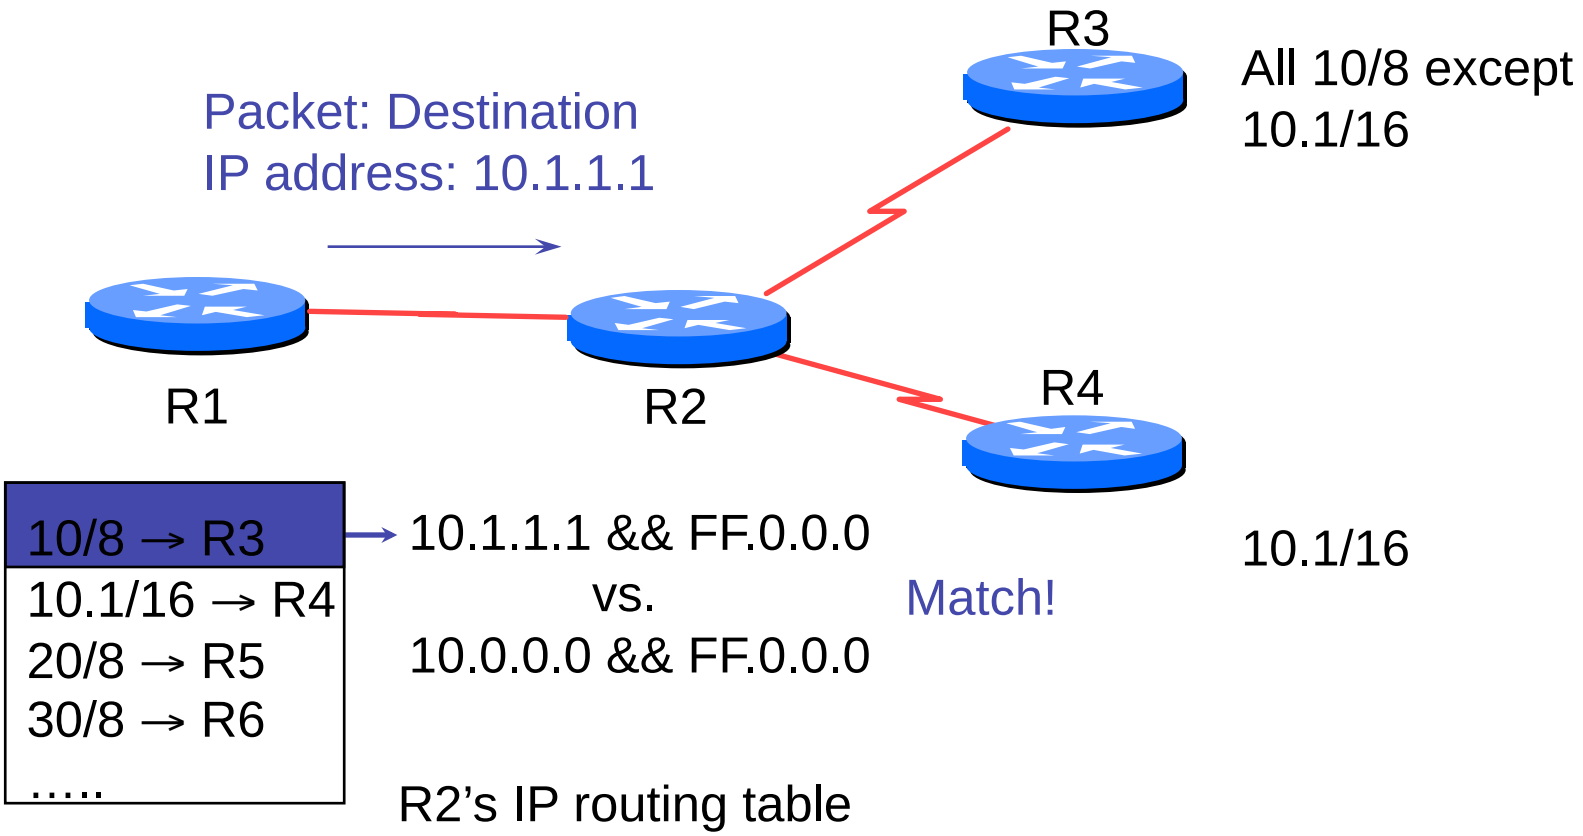

# IP route lookup: Longest match routing

- Based on destination IP packet

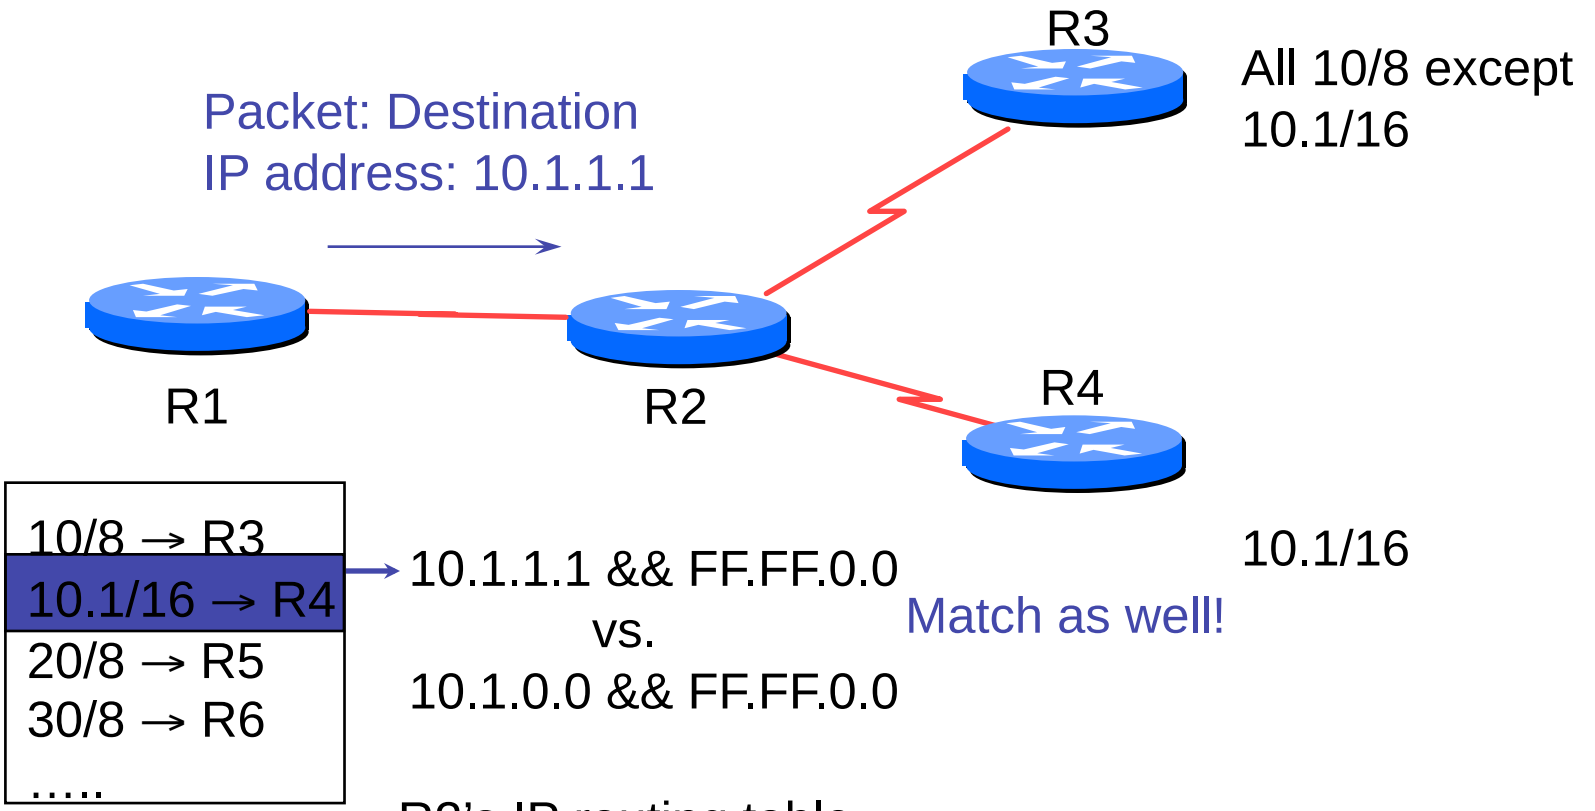

# IP route lookup: Longest match routing

- Based on destination IP packet

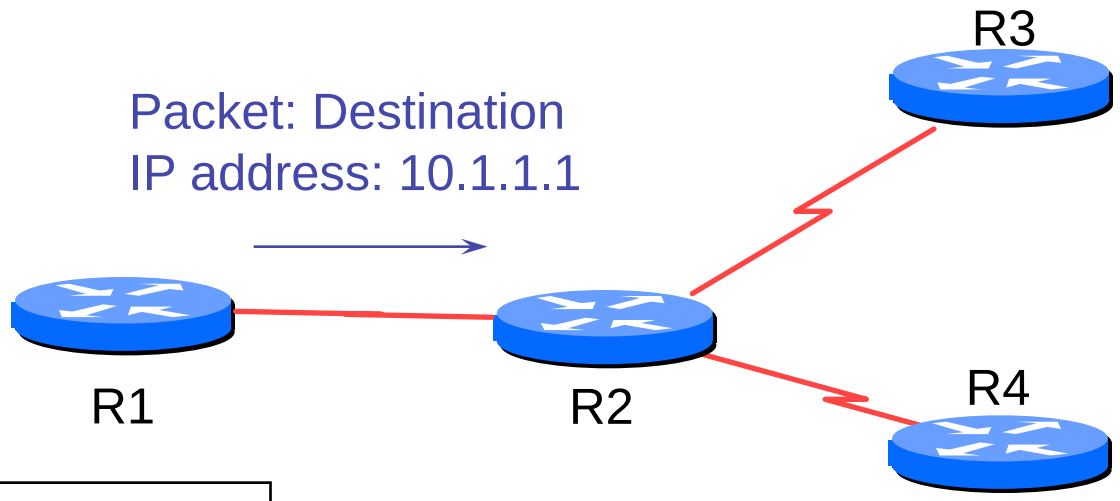


10/8 → R3
10.1/16 → R4
20/8 → R5
30/8 → R6
…..

10.1.1.1 && FF.0.0.0
vs.
20.0.0.0 && FF.0.0.0

Does not match!

R2's IP routing table

# IP route lookup: Longest match routing

- Based on destination IP packet

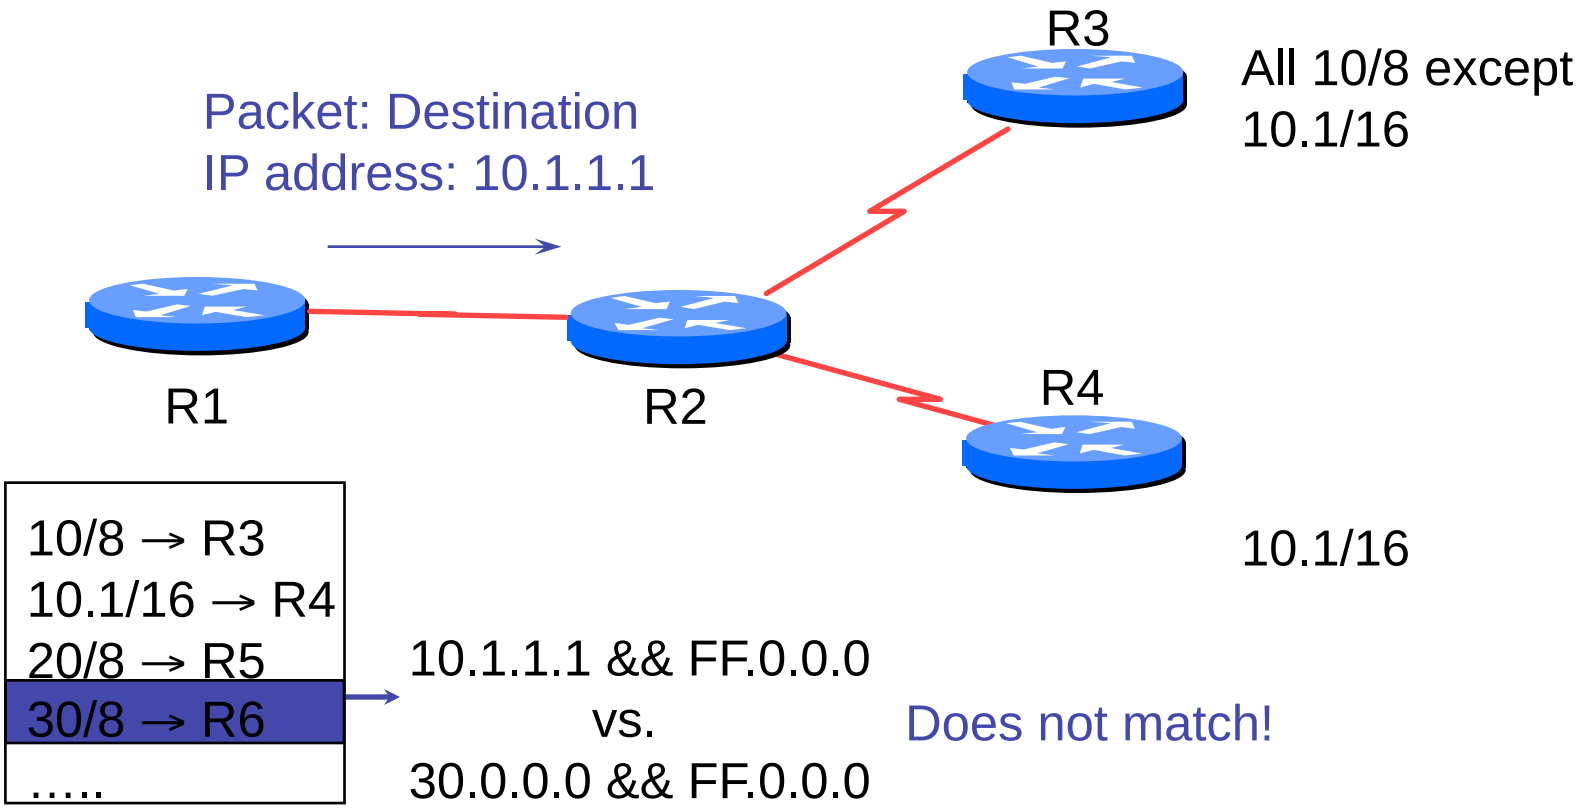

# IP route lookup: Longest match routing

- Based on destination IP packet

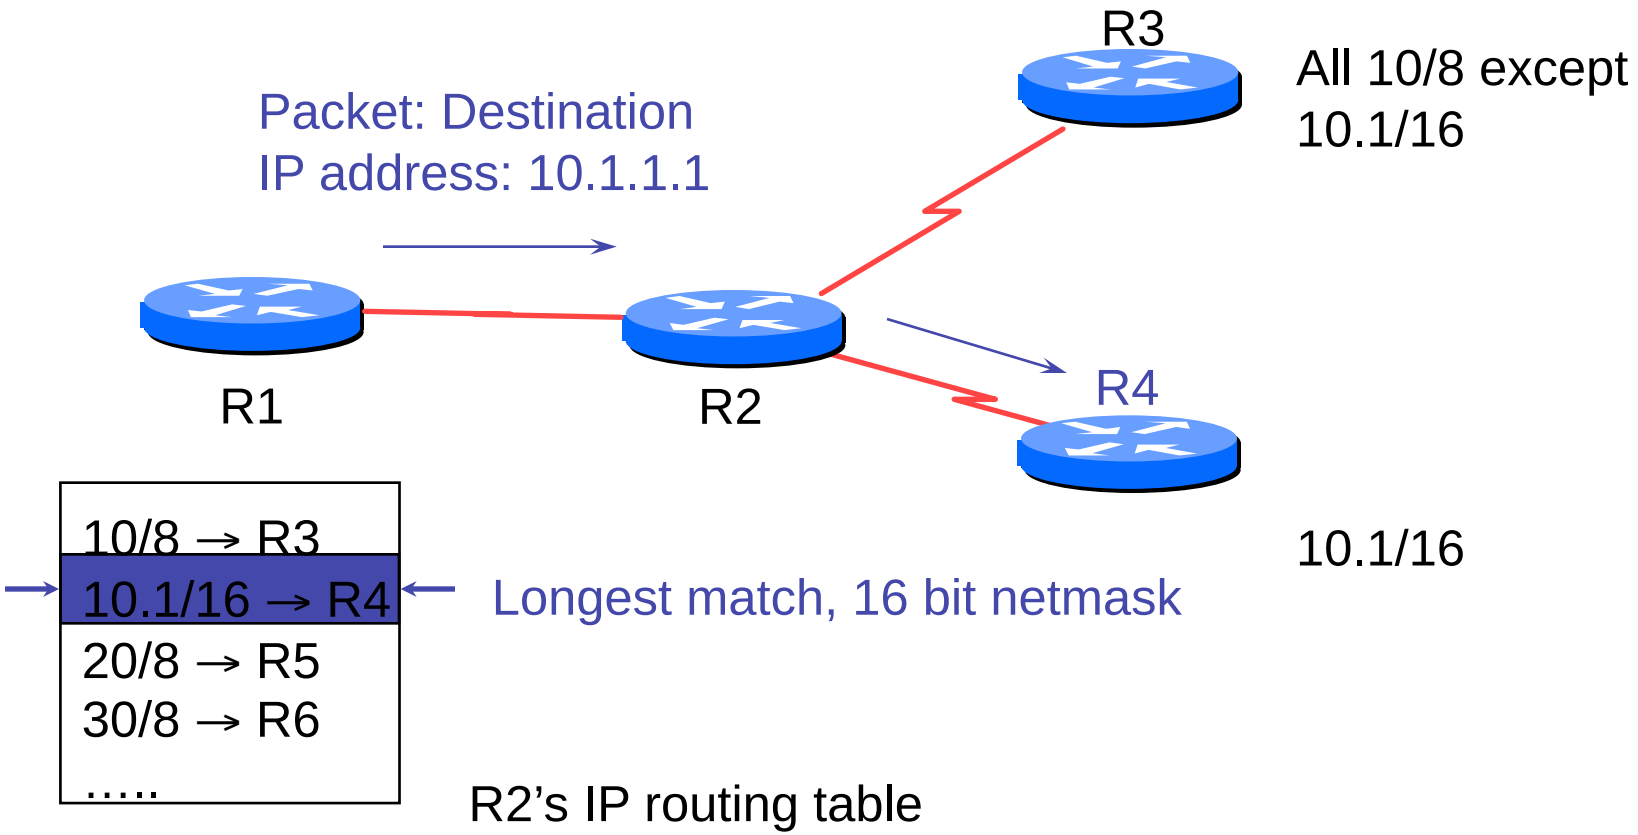

# IP Forwarding

- Router makes decision on which interface a packet is sent to
- Forwarding table populated by routing process
- Forwarding decisions:
  - destination address
  - class of service (fair queuing, precedence, others)
  - local requirements (packet filtering)
- Can be aided by special hardware

# Routing Tables Feed the Forwarding Table