Routing, Switching, and Network Infrastructure

Intro to Networks Module 5.0

Brendan Shea, PhD

Rochester Community and Technical College
Intro to Networking

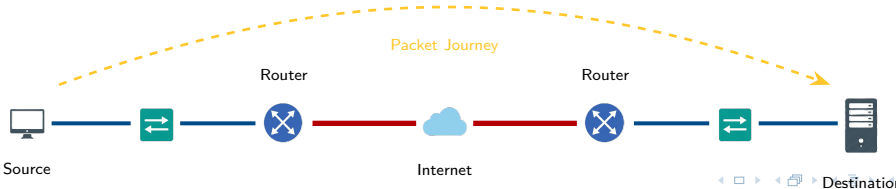

Outline I

Review: From IP Addressing to Routing

- Module 4 covered **IP addresses** and **subnetting** (Layer 3 addressing).
- Today's question: How does data actually *get* to its destination?
- IP addresses tell us **where** to go, but **routing** tells us **how** to get there.
- Routers make decisions about the best path for each packet.

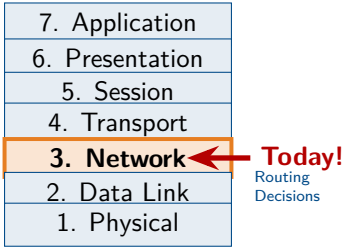

The Journey Ahead

Module 4: Addressing (naming the destination)
Module 5: Routing (finding the path there)

Module Overview and Learning Outcomes

Topics Covered

1. Routing Technologies and Tables
2. Dynamic Routing Protocols
3. Network Address Translation (NAT)
4. Firewalls and Security
5. Enterprise Network Topologies
6. Virtual LANs (VLANs)

Learning Outcomes

After this module, you will be able to:

- Explain how routers forward packets using routing tables
- Compare static vs. dynamic routing protocols
- Describe NAT and PAT operation
- Identify firewall types and proper placement
- Design networks using three-tier hierarchy
- Configure and troubleshoot VLANs

The Router: Gateway Between Networks

- **Routers** connect different networks at Layer 3.
- Each router interface has its own IP address on a **different subnet**.
- Routers make **forwarding decisions** based on the destination IP address.
- Every packet passing through decrements its **TTL** by 1.

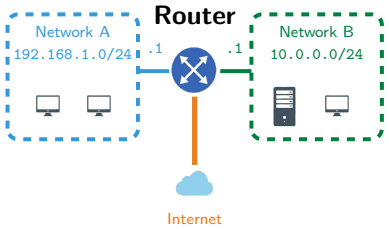

Key Difference from Switches

Switches: MAC addresses (Layer 2)
Routers: IP addresses (Layer 3)

Routing Tables and Path Selection

- Every router maintains a **routing table** listing known networks.
- Key components of each entry:
  - **Destination network** and prefix/mask
  - **Next hop** (next router's IP) OR exit interface
  - **Metric** (cost to reach destination)
- Router uses **longest prefix match** to find best route.

Longest Prefix Match

If a packet for `192.168.1.50` matches both `192.168.0.0/16` and `192.168.1.0/24`, the **/24 wins** because it's more specific.

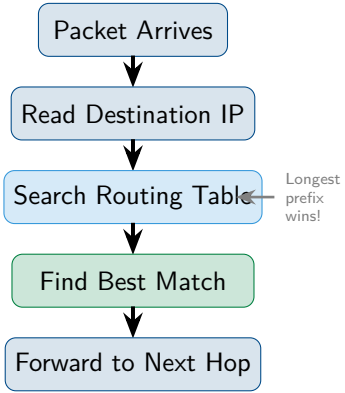

Static and Default Routes

Static Routes

- **Manually configured** by administrator
- Don't change unless admin modifies
- Good for: small networks, specific paths
- Downside: won't adapt to failures

```
ip route 10.0.0.0 255.0.0.0 192.168.1.1
```

Default Route

- The "**route of last resort**"
- Used when **no specific match** found
- Points to gateway for "everything else"
- Written as `0.0.0.0/0`

```
ip route 0.0.0.0 0.0.0.0 203.0.113.1
```

When to Use Static

Small networks, stub networks (one exit), or when you need security control over paths.

Analogy

Static: "Take Highway 10 to Grandma's."
Default: "For anywhere else, head to the main highway."

Routing Table Example

Sample Routing Table

Network	Mask	Next Hop	Metric	Type
192.168.1.0	/24	Gig0/0	0	Direct
10.0.0.0	/24	Gig0/1	0	Direct
172.16.0.0	/16	10.0.0.2	1	Static
0.0.0.0	/0	192.168.1.254	1	Default

Directly Connected | Static Route | Default Route

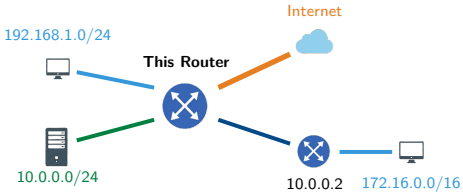

Reading the Table

To reach `172.16.0.0/16`, send packets to next hop `10.0.0.2`. For unknown destinations, use the **default route** via `192.168.1.254`.

Packet Forwarding Process

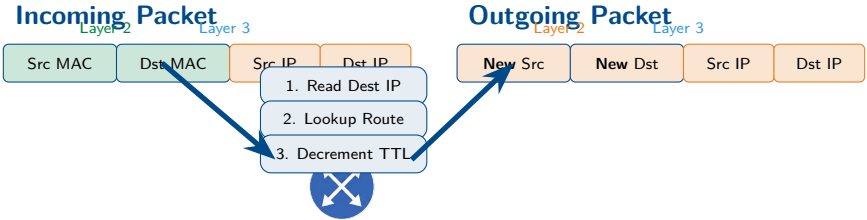

Layer 2 headers change at each hop (new MACs)
Layer 3 headers stay the same end-to-end (same IPs)

Fragmentation and MTU

- **MTU** (Maximum Transmission Unit) is the largest packet size a link can carry.
- Standard Ethernet MTU: **1500 bytes**.
- If packet $>$ MTU, router must **fragment** it into smaller pieces.
- Fragments are reassembled at the **destination**.

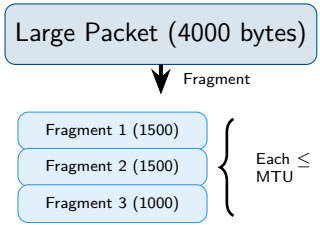

Problems with Fragmentation

- Adds processing overhead
- If one fragment lost, entire packet must be retransmitted
- Slower overall performance

Path MTU Discovery

Modern systems discover the smallest MTU along the entire path to avoid fragmentation.

Case Study: Westley's Journey to Florin

▶ Case Study: Westley's Journey to Florin

Westley is sending a message from the Pirate Ship network (`172.16.0.0/24`) to Princess Buttercup at Castle Florin (`10.0.0.0/24`). The ship's router has this routing table:

Network	Next Hop	Type
172.16.0.0/24	Directly Connected	Direct
10.0.0.0/24	192.168.50.1	Static
0.0.0.0/0	203.0.113.1	Default

Westley's PC: `172.16.0.25` Buttercup's PC: `10.0.0.100`

Review Questions

1. Which routing table entry will the router use for Westley's packet?
2. What is the next hop IP address?
3. What would happen if someone deleted the `10.0.0.0/24` entry?

Case Study Solution: Westley's Journey to Florin

✓ Solution: Westley's Journey to Florin

1. Router uses **10.0.0.0/24**—longest prefix match for `10.0.0.100`.
2. Next hop is **192.168.50.1** (path to Castle Florin).
3. If deleted, router uses **default route** (`0.0.0.0/0`) via `203.0.113.1`—a longer path through the kingdom!

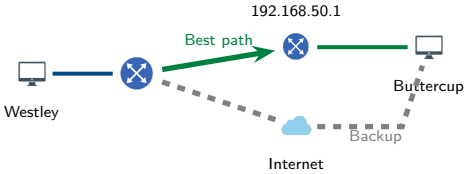

Key Lesson

"As you wish!" Specific routes always win over the default. Without proper entries, packets may take the long way—or never arrive!

Basic Router Configuration

Essential Settings

- **Hostname** – Identify the router
- **Interface IPs** – Address each port
- **Enable interfaces** – `no shutdown`
- **Routes** – Static and/or default
- **Passwords** – Secure access

Two Configuration Methods

CLI: Command-line interface (Cisco IOS)
GUI: Web-based management interface

Sample CLI Commands

```
Router> enable
Router# configure terminal
Router(config)# hostname Florin-GW
Florin-GW(config)# interface g0/0
Florin-GW(config-if)# ip address
  192.168.1.1 255.255.255.0
Florin-GW(config-if)# no shutdown
Florin-GW(config-if)# exit
Florin-GW(config)# ip route 0.0.0.0
  0.0.0.0 203.0.113.1
```

Tip

Always save config: `copy run start`

Routing Table Tools

Cisco IOS Commands

`show ip route`
Display full routing table

`show ip route` *`network`*
Show specific route entry

`show ip interface brief`
Interface status summary

`show running-config`
View current configuration

Windows / Linux Commands

`route print` (Windows)
`netstat -r` (Windows/Linux)
Display host routing table

`ip route` (Linux)
Modern Linux routing display

`netstat -rn`
Numeric output (no DNS lookup)

Quick Check

Use `ip route get` *IP* on Linux to see which route would be used for a specific destination.

tracert and traceroute

- Shows the **path** packets take to a destination.
- Uses **TTL manipulation**:
  - Send packet with TTL=1
  - First router sends back error
  - Increment TTL, repeat
- Each router returns **ICMP Time Exceeded**.
- `tracert` (Windows)
  `traceroute` (Linux/Mac)

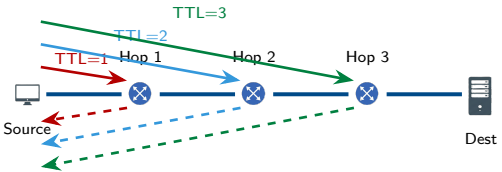

Sample Output

```
1   2ms    1ms    1ms    192.168.1.1
2   10ms   9ms    11ms   10.0.0.1
3   15ms   14ms   15ms   203.0.113.50
```

Use Cases

Find where packets get stuck, identify slow links, verify routing path.

Dynamic Routing Protocols Overview

Why Dynamic Routing?

- Static routes don't adapt to failures
- Large networks have hundreds of routes
- Dynamic protocols **share route info automatically**
- Network changes propagate without admin intervention

Two Categories

IGP (Interior Gateway Protocol)
Within an organization: RIP, OSPF, EIGRP

EGP (Exterior Gateway Protocol)
Between organizations: BGP

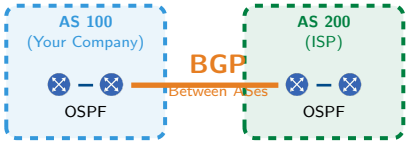

Autonomous System (AS)

A network under single administrative control.
Each AS has a unique number (ASN).

RIP and EIGRP

RIP (Routing Information Protocol)

- Simple, easy to configure
- Uses **hop count** as metric
- Maximum 15 hops (16 = unreachable)
- Broadcasts updates every 30 seconds
- Slow convergence

Best for: Small, simple networks

EIGRP (Enhanced Interior Gateway RP)

- Cisco-developed (now partially open)
- Uses **bandwidth + delay** for metric
- Fast convergence
- Only sends updates when changes occur
- Supports unequal-cost load balancing

Best for: Cisco enterprise networks

RIP Limitation

Hop count ignores link speed! A 10 Gbps path with 3 hops loses to a 1 Mbps path with 2 hops.

EIGRP Advantage

Keeps backup routes ready—failover is nearly instant!

OSPF (Open Shortest Path First)

- **Industry standard** (open protocol)
- Uses **cost** based on bandwidth
- Creates complete network map (**link-state**)
- Fast convergence
- Supports **hierarchical design** with areas

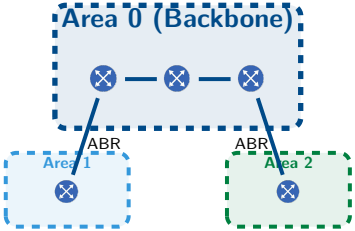

OSPF Cost Formula

$$\text{Cost} = \frac{10^8}{\text{bandwidth (bps)}}$$

100 Mbps	Cost = 1
10 Mbps	Cost = 10
1 Gbps	Cost = 1

ABR

Area Border Router connects areas to the backbone (Area 0).

BGP (Border Gateway Protocol)

- The **routing protocol of the Internet**
- Connects different organizations (ASes)
- **Path vector** protocol
- Uses **policies** more than metrics
- Very slow, deliberate convergence

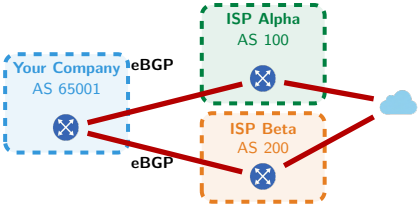

When to Use BGP

- Connecting to multiple ISPs
- Running your own AS
- Need fine control over routing policy

Multihoming

Connecting to multiple ISPs provides redundancy and load balancing.

Warning

BGP misconfigurations can break parts of the Internet! Handle with care.

Route Selection and Administrative Distance

Routing Protocol Metrics

Protocol	Metric
RIP	Hop count
EIGRP	Bandwidth + delay
OSPF	Cost (bandwidth)
BGP	Path attributes

What If Multiple Protocols?

When the same route is learned from different protocols, **Administrative Distance** decides which to trust.

Administrative Distance (AD)

Lower AD = more trusted

Route Source	AD
Directly Connected	0
Static Route	1
EIGRP (internal)	90
OSPF	110
RIP	120
External EIGRP	170
Unknown	255

Example

If OSPF (AD 110) and RIP (AD 120) both know a route, OSPF wins!

Case Study: Inigo's Path to Count Rugen

▶ Case Study: Inigo's Path to Count Rugen

Inigo Montoya has configured routers to help him find Count Rugen (the six-fingered man). His router has learned two paths to Rugen's network (`10.20.30.0/24`):

Protocol	Next Hop	Metric	AD
OSPF	192.168.1.1	Cost: 20	110
RIP	192.168.2.1	Hops: 3	120

Both paths are currently working. Inigo needs to reach Rugen as quickly as possible!

Review Questions

1. Which route will the router prefer and why?
2. What is Administrative Distance and how does it apply here?
3. If the OSPF path fails, what happens?

Case Study Solution: Inigo's Path to Count Rugen

✓ Solution: Inigo's Path to Count Rugen

1. Router prefers **OSPF route**—AD of 110 beats RIP's AD of 120.
2. **Administrative Distance** is a "trustworthiness" rating. Lower AD = more trusted. When multiple protocols know a route, lowest AD wins.
3. If OSPF fails, router automatically uses **RIP route** as backup!

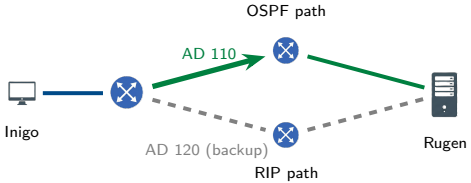

Key Lesson

"Hello. My name is Inigo Montoya..." Administrative Distance determines which protocol to trust. The backup route is ready if the primary fails!

Edge Routers and Network Boundaries

- **Edge router** connects internal network to outside (ISP).
- Also called: border router, gateway router.
- Key responsibilities:
  - Route between internal and external
  - Apply **NAT** (translate IPs)
  - First line of security (ACLs)
  - Connect to ISP via WAN link

Demarcation Point

The boundary where your network ends and the ISP's begins. Usually at the edge router.

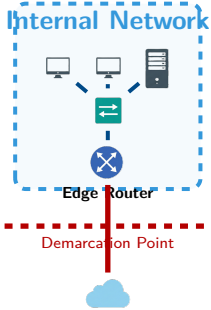

Network Address Translation (NAT)

- **NAT** translates private IPs to public IPs.
- Allows many internal devices to share limited public addresses.
- Essential because IPv4 addresses are exhausted!
- Performed by edge router or firewall.

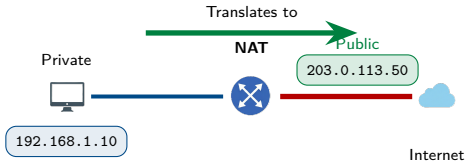

NAT Terminology

- **Inside Local**: Private IP (192.168.1.10)
- **Inside Global**: Public IP seen by internet
- **Outside**: External destination

Why NAT Matters

Your home network might have 20 devices, but your ISP only gives you **one** public IP. NAT makes this work!

NAT Types

Static NAT

- One private IP $\leftrightarrow$ One public IP
- Permanent, fixed mapping
- Use: Servers needing consistent public address

`192.168.1.10` $\leftrightarrow$ `203.0.113.10`

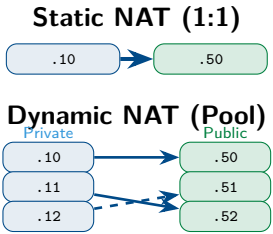

Dynamic NAT

- Pool of private $\rightarrow$ Pool of public
- First-come, first-served
- Mapping changes over time
- Use: When you have several public IPs

Limitation

Both Static and Dynamic NAT require one public IP per active connection. Not scalable!

PAT: Port Address Translation

- **PAT** = Many private IPs share ONE public IP
- Uses **port numbers** to track connections
- Also called: NAT Overload, NAPT
- **Most common** type in homes and businesses

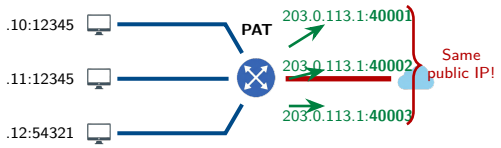

How PAT Works

1. Host sends packet (src port 12345)
2. Router changes src IP to public
3. Router assigns unique port (e.g., 40001)
4. Tracks mapping in NAT table
5. Return traffic uses port to find host

Capacity

With 65,000 available ports, one public IP can support thousands of simultaneous connections!

Firewall Uses and Types

What Firewalls Do

- Filter traffic based on rules
- Block unauthorized access
- Allow legitimate traffic through
- Log security events
- Enforce security policies

Default Behavior

Most firewalls: **Deny all, permit by exception**. Only explicitly allowed traffic passes.

Firewall Types

- **Packet Filtering**
  Checks IP/port only (stateless)
- **Stateful Inspection**
  Tracks connection state
- **Application Layer**
  Inspects content (Layer 7)
- **Next-Gen (NGFW)**
  All above + IPS, deep inspection

Trend

Modern networks use NGFW for comprehensive protection.

Firewall Placement and DMZ

Placement Options

- **Perimeter**: Between LAN and internet
- **Internal**: Between network segments
- **Host-based**: On individual devices

DMZ (Demilitarized Zone)

- Screened subnet for public servers
- Web, email, DNS servers go here
- Isolated from internal LAN
- If compromised, LAN still protected

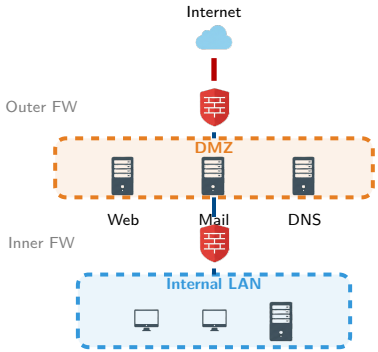

Defense in Depth

Multiple firewall layers provide better protection than a single firewall.

Introduction to VLANs

- **VLAN** = Virtual Local Area Network
- Logically separates one physical switch into multiple **broadcast domains**.
- Devices on different VLANs cannot communicate directly—need a router.
- Configured per switch port.

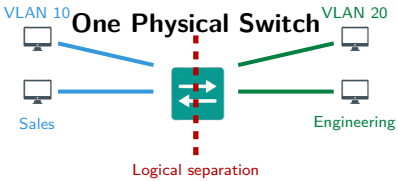

Benefits of VLANs

- **Security**: Isolate sensitive traffic
- **Performance**: Reduce broadcast traffic
- **Flexibility**: Group by function, not location
- **Cost**: One switch, many networks

Key Point

VLANs create Layer 2 isolation. Cross-VLAN traffic requires Layer 3 routing!

VLAN Trunking and 802.1Q

The Problem

How do VLANs span multiple switches? We need a way to carry VLAN info between switches.

Trunk Ports

- Carry traffic for **multiple VLANs**
- Connect switches to switches
- Connect switches to routers
- Use **802.1Q tagging**

Access vs Trunk

Access port: One VLAN, untagged
Trunk port: Multiple VLANs, tagged

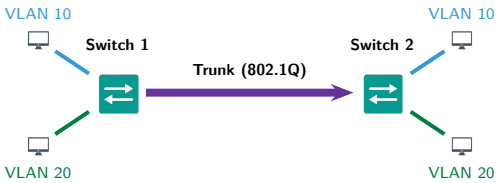

802.1Q Tag

Adds 4-byte tag to Ethernet frame:

- **VLAN ID**: 12 bits (1–4094)
- **Priority**: 3 bits (QoS)

Voice VLANs

- VoIP phones need **priority treatment** for quality calls.
- **Voice VLAN** separates phone traffic from data traffic.
- One port supports both:
  - Data VLAN (for PC)
  - Voice VLAN (for phone)
- Phone traffic gets higher QoS priority.

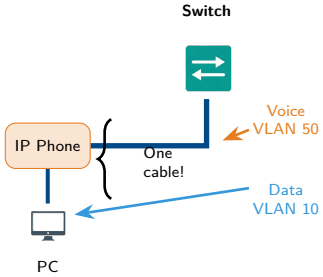

Configuration Example

```
switchport mode access
switchport access vlan 10
switchport voice vlan 50
```

Why Separate?

Voice is sensitive to delay and jitter. Separating it ensures calls stay clear even when data traffic is heavy.

Native VLAN and Security

What is Native VLAN?

- Traffic on trunk that has **no 802.1Q tag**
- Default: VLAN 1
- Used for backward compatibility
- Both ends of trunk must match!

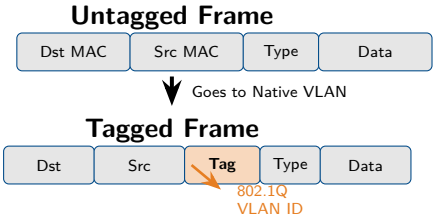

Security Risk

Attackers can exploit native VLAN mismatches to hop between VLANs. This is called **VLAN hopping**. Best practices:

- Change native VLAN from default
- Use unused VLAN for native
- Ensure both trunk ends match

Configuration

```
switchport trunk native vlan 999
```

VLAN 1 Concerns

VLAN 1 carries management traffic by default. Keep it separate from user data!

Case Study: Miracle Max's Potion Shop

▶ Case Study: Miracle Max's Potion Shop

Miracle Max runs a potion shop with three departments, each on its own VLAN:

Department	VLAN	Subnet
Potion Brewing	VLAN 10	192.168.10.0/24
Customer Sales	VLAN 20	192.168.20.0/24
Billing Office	VLAN 30	192.168.30.0/24

The Sales team (`192.168.20.50`) needs to check inventory on the Brewing server (`192.168.10.10`), but their packets aren't getting through!

Review Questions

1. Why can't Sales communicate directly with Brewing?
2. What device is needed to enable this communication?
3. What is this process called?

Case Study Solution: Miracle Max's Potion Shop

✓ Solution: Miracle Max's Potion Shop

1. VLANs are separate broadcast domains—**Layer 2 isolation** prevents direct communication between VLANs.
2. A **router** (or Layer 3 switch) is needed to route between VLANs.
3. This is called **inter-VLAN routing**.

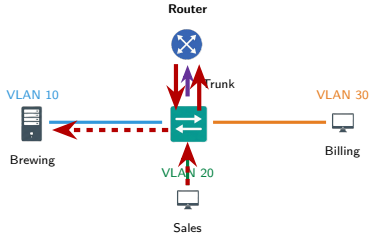

Key Lesson

"Have fun storming the castle!" VLANs provide security through isolation, but inter-VLAN routing lets authorized traffic flow when needed.

Module 5.0 Summary I

Key Concepts:

- **Routers** connect networks at Layer 3. **Routing tables** store known destinations; **longest prefix match** selects best route.
- **Static routes**: Manually configured, predictable. **Default route** (0.0.0.0/0) is gateway of last resort.
- Dynamic routing: **RIP** (hop count, max 15), **EIGRP** (Cisco, fast convergence), **OSPF** (open standard, link-state), **BGP** (internet backbone).

Module 5.0 Summary II

- **Administrative Distance (AD)** determines protocol trustworthiness: Direct (0) > Static (1) > EIGRP (90) > OSPF (110) > RIP (120).
- **NAT** translates private IPs to public. Types: **Static NAT** (1:1), **Dynamic NAT** (pool), **PAT** (many-to-one using ports).
- **Firewalls** filter traffic: Packet filter, Stateful, Application-level, **NGFW** (Next-Generation). **DMZ** = screened subnet for servers.

Module 5.0 Summary III

VLANs & Network Design:

- **VLANs** logically segment networks on one switch, creating separate broadcast domains for security, performance, and flexibility.
- **802.1Q** tags frames with VLAN ID (1-4094). **Trunk ports** carry multiple VLANs. **Native VLAN**: untagged traffic.
- **Voice VLANs** separate voice from data traffic, apply QoS priority (one port, two VLANs).
- **Inter-VLAN routing** requires Layer 3 device (router or Layer 3 switch with **SVI** - Switched Virtual Interface).
- **Three-tier hierarchy**: **Core** (backbone), **Distribution** (policy/routing), **Access** (end-device connection).
- **Spine-leaf architecture**: Modern data center design with predictable latency and no oversubscription.