Troubleshooting an Enterprise Network

Introducing Routing and Switching in the Enterprise – Chapter 9

Objectives

After completion of this chapter, you should be able to:

- Explain the importance of uptime and the types of issues that cause failure.
- Isolate and correct switching problems.
- Isolate and correct routing issues.
- Isolate and correct WAN configurations.
- Isolate and correct ACL issues.

Uptime and downtime

- Network **uptime** is the time that the **network is available** and functioning as expected.
- Network **downtime** is any time that the network is **not performing as required**.
- Network outages also prevent customers from placing orders or obtaining the information they require.
- **Downtime results in lost productivity, customer frustration, and often the loss of customers to competitors!**

Uptime and downtime

- For organizations, **any downtime is extremely costly**.
- Many factors cause network downtime. These include:
  - Weather and natural **disasters**
  - Security **breaches**
  - Man-made **disasters**
  - Power surges
  - **Virus** attacks
  - Equipment failure
  - Misconfiguration of devices
  - **Lack of resources**

Uptime and downtime

- To ensure the proper and efficient flow of traffic, a good design includes **redundancy of all critical components and data paths** to <u>eliminate single points of failure</u>
- The **three-layer hierarchical network** design model separates the functionality of the various networking devices and links. This separation ensures <u>efficient network performance</u>.
- In addition, the use of **enterprise class equipment** provides a high degree of <u>reliability</u>.
- Even with proper network design, **some downtime is inevitable** (failure is **<u>always</u>** an option!!).
- To guarantee service levels, an enterprise should have service level agreements (SLAs) with key suppliers.

3-layer network design

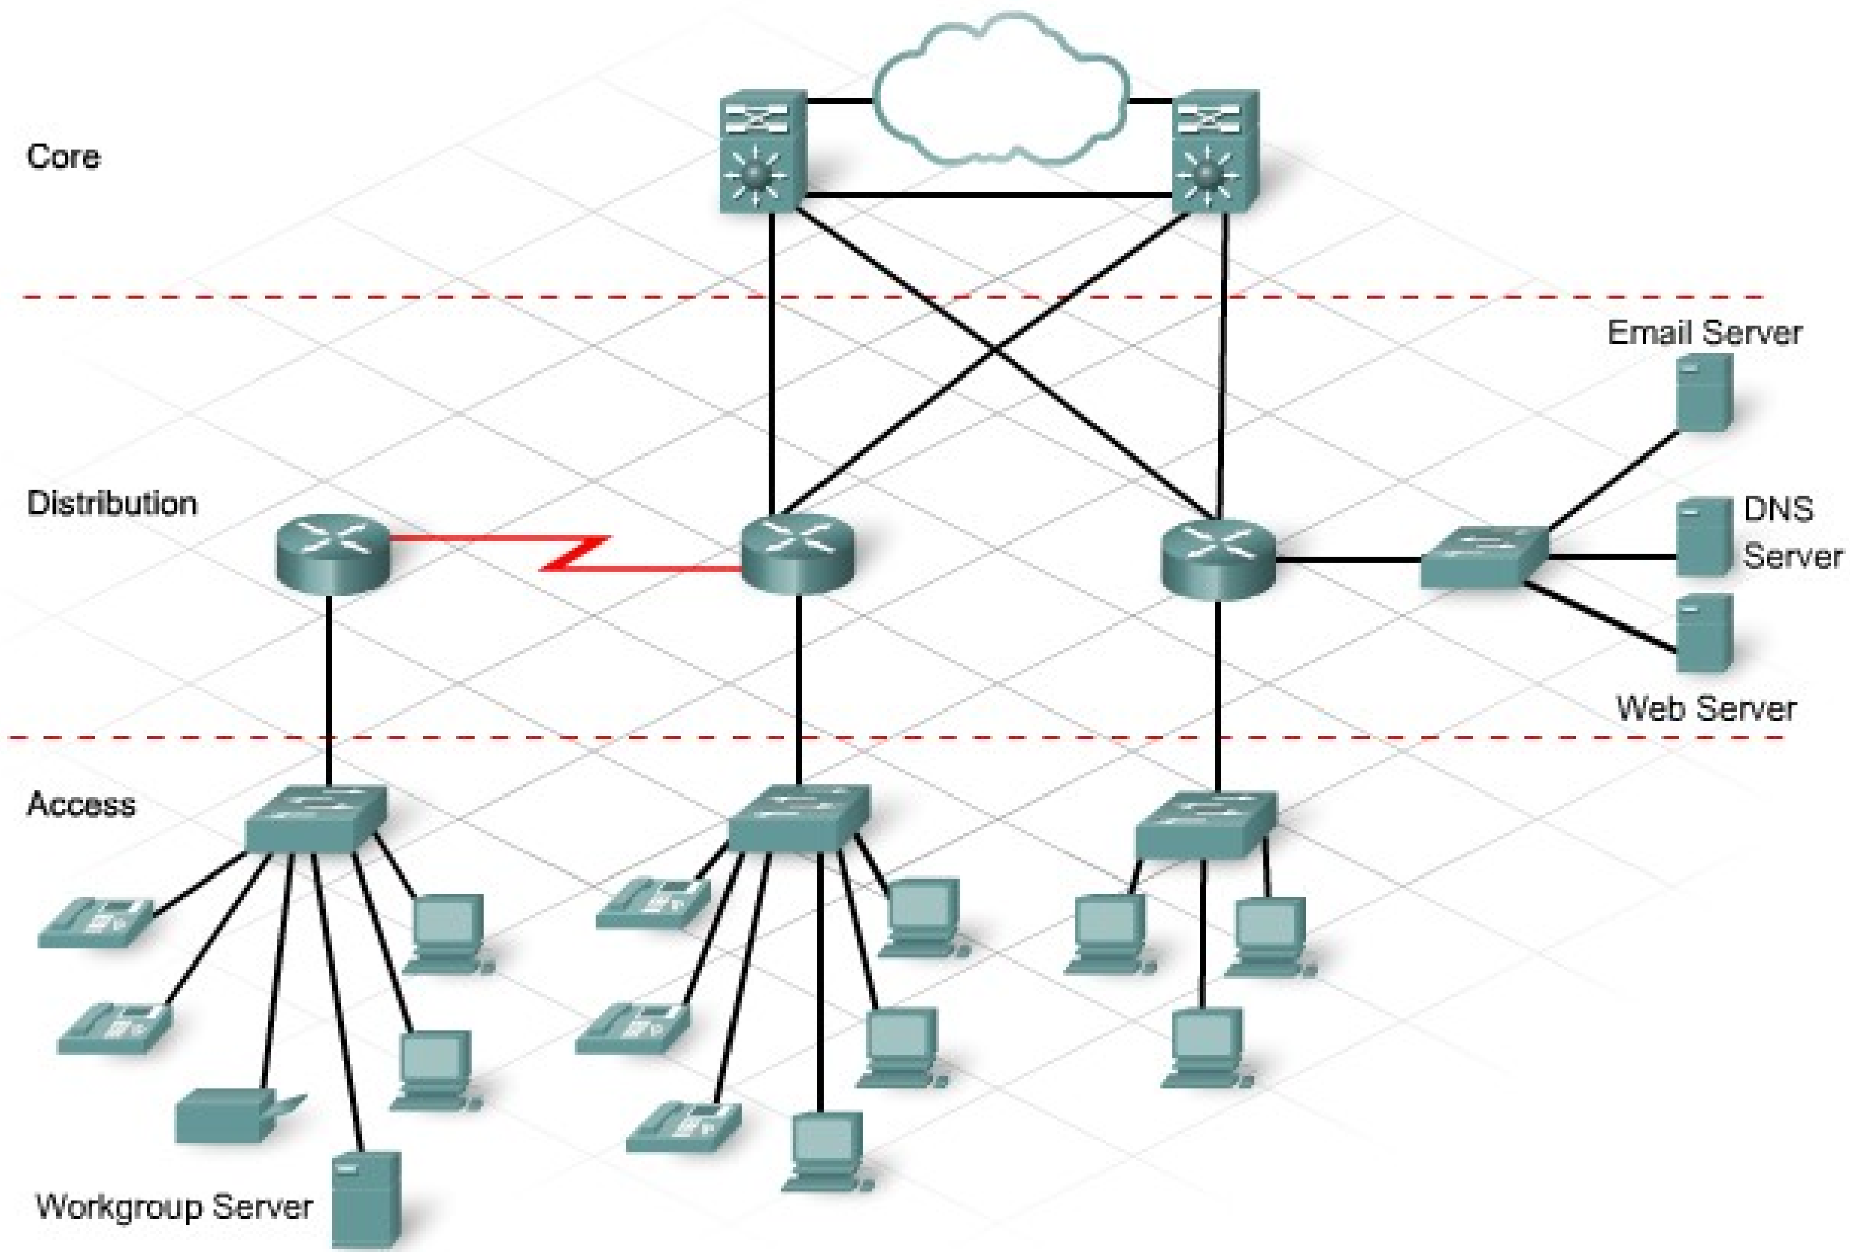

Network Monitoring

- One way of ensuring uptime is to **monitor** current network functionality and perform **proactive maintenance**.
- The purpose of network monitoring is to watch network performance in comparison to a predetermined **baseline**. Any observed deviations from this baseline indicate potential problems with the network and require investigation. **As soon as** the **network administrator** determines the cause of **degraded performance**, corrective actions can be taken to **prevent** a serious network outage.
- Several groups of tools are available for monitoring network performance levels and gathering data. These tools include:
  - Network utilities
  - Packet sniffing tools
  - SNMP monitoring tools
- A network administrator performs **proactive maintenance** on a **regular basis** to verify and service equipment. By doing this, the administrator can detect weaknesses **prior to a critical error** that could bring down the network. Like regular servicing on a car, **proactive maintenance extends the life of a network device**.

Network documentation

- Network monitoring tools, techniques, and programs rely on the availability of a complete set of accurate and current **network documentation**. This documentation includes:
  - **Physical** and **logical topology** diagrams
  - **Configuration files** of **all** network devices
  - A **baseline** performance level
- It is best practice to determine **baseline** network performance levels when the network is **first installed** and then again **after any major changes** or upgrades occur.

Network Monitoring Tools

- Simple network utilities, like **ping** and **tracert**, provide information on the performance of the network or network link. Performing these commands at **multiple times** shows the difference in time required for a packet to travel between two locations.
- **Packet sniffing tools** monitor the types of traffic on various parts of the network. These tools indicate if there is an <u>excessive amount of a particular traffic type</u>. They <u>examine the contents of the packets</u>, which provides a quick way of locating the source of this traffic.
- These tools may also be able to remedy the situation **before** network congestion becomes critical. For example, traffic sniffing can detect whether a type of traffic or a particular transaction occurring on the network is unexpected. This detection might stop a potential denial of service attack **before** it impacts network performance.

SNMP

- **Simple Network Management Protocol** (SNMP) allows monitoring of individual devices on the network.
- SNMP-compliant devices **use agents** to monitor a number of predefined parameters for specific conditions.
- These agents collect information and store it in a database known as the **management information base** (MIB).
- **SNMP polls devices at regular intervals** to collect information about managed parameters.
- SNMP also traps certain events that **exceed a predefined threshold or condition**.

Failure domain

- When designing an enterprise network limit the size of a **failure domain**.
- The failure domain is **the area of the network that is impacted by the failure or misconfiguration of a network device**.
- The actual size of the domain depends on the device and the type of failure or misconfiguration.
- When troubleshooting a network, determine the scope of the issue and **isolate the issue to a specific failure domain**.

Troubleshooting

- Many different structured and unstructured **problem-solving techniques** are available to the network technician. These include:
  - Top-down
  - Bottom-up
  - Divide-and-conquer
  - Trial-and-error
  - Substitution
- Most **experienced network technicians** rely on the knowledge gained from past experience and start the troubleshooting process using a **trial-and-error** approach. Correcting the problem in this manner **saves a great deal of time**.
- Unfortunately, less experienced technicians cannot rely solely on previous experience. Additionally, many times the trial-and-error approach does not provide a solution. Both of these cases require a more structured approach to troubleshooting.

Troubleshooting

- When a situation requires a more structured approach, most network personnel **use a layered process based on the OSI or TCP/IP models.**
- The technician uses previous experience to determine if the issue is associated with the **lower layers** of the OSI model or the **upper layers**.
- The **layer dictates** whether a top-down or bottom-up approach is appropriate.
- When approaching a problem situation, follow the **generic problem-solving model**, regardless of the type of troubleshooting technique used.
  - Define the problem
  - Gather facts
  - Deduce possibilities and alternatives
  - Design plan of action
  - Implement solution
  - Analyze results

Troubleshooting switches

- Faults with the switch hardware or configuration prevent connection between local and remote devices.
- <u>The most common problems with switches occur at the Physical Layer</u>.
- If a switch is installed in an **unprotected environment**, it can suffer damage such as dislodged or damaged data or power cables. Ensure that switches are <u>placed in a physically secure area</u>.
- If an end device cannot connect to the network and **the link LED is not illuminated**, the link or the switch port is defective or shutdown, perform the following steps:
  - Ensure that the **power LED** is illuminated.
  - Ensure that the **correct type of cable** connects the end device to the switch.
  - **Reseat the cables** at both the workstation and the switch end.
  - **Check the configuration** to ensure that the port is in a no shutdown state.

Troubleshooting switches

- If a switch port fails or malfunctions, the easiest way to test it is to **move the physical connection to another port** and see if this corrects the problem.
- <u>Ensure that switch port security has not disabled the port.</u> Confirm this using the following commands:
  - **`show running-config`**
  - **`show port-security interface interface_id`**
- If the switch security settings are disabling the port, review the security policy to **see if altering the security is acceptable**.

Troubleshooting switches: MAC addresses

- Switches function at Layer 2 and keep a record of the MAC address of all connected devices. If the **MAC address in this table is not correct**, the switch forwards information to the wrong port and communication does not occur.
- To display the MAC address of the device connected to each switch port, use:
  - **`show mac-address-table`**
- To clear the dynamic entries in the table, issue the command:
  - **`clear mac-address-table dynamic`**
- The switch then repopulates the MAC address table with updated information.

Troubleshooting switches: Switching Loops

- **STP** prevents bridging loops and broadcast storms by **shutting down redundant paths** in a switched network. If STP bases its decisions on inaccurate information, loops may occur.
- <u>Indicators that a loop is present</u> in a network include:
  - **Loss of connectivity** to, from, and through affected network regions
  - **High CPU utilization** on routers connected to affected segments
  - **High link utilization** up to 100%
  - **High switch backplane utilization** as compared to the baseline utilization
  - **Syslog messages indicating packet looping**, constant address relearning, or MAC address flapping messages
  - **Increasing number of output drops on many interfaces**

Troubleshooting switches: Switching Loops

- A loop develops when the switch does not receive BPDUs or is unable to process them. This problem could be due to:
  - Misconfigurations
  - Defective transceivers
  - Hardware and cabling issues
  - Overloaded processors
- Overloaded processors disrupt STP and prevent the switch from processing the BPDUs.
- A port that is **flapping** causes multiple transitions to occur.
- **Multiple transactions** can overload the processors. This should be a <u>rare occurrence</u> in a properly configured network.

Troubleshooting switches: suboptimal switching

- Left to default values, **STP does not always identify** the best root bridge or root ports.
- **Changing the priority value** on a switch can force the selection of the root bridge.
- The root bridge should normally be at the **center** of the network to provide for optimum switching.
- When troubleshooting STP, use the following commands:
  - To provide information about the STP configuration:
    - `show spanning-tree`
  - To provide information about the STP state of an individual port:
    - `show spanning-tree interface interface_id`

Troubleshooting switches: VLANs

- If the non-functioning ports are in the **same VLAN**, the hosts must have IP addresses on the **same network** or subnet in order to communicate.
- If the non-functioning ports are in **different VLANs**, communication is only possible with the aid of a **Layer 3 device**, such as a router.
- If information is required on a specific VLAN, use the following command **show vlan id vlan_number** to display the ports assigned to each VLAN.
- If **inter-VLAN routing** is required, verify the following configurations:
  - <u>One port from each VLAN connects into a router</u> interface or subinterface.
  - Both the switch port and the router interface are <u>configured with trunking</u>.
  - Both the switch and router interface are <u>configured with the same encapsulation</u>.
- **Newer switches default to 802.1Q**, but some Cisco switches support both 802.1Q and <u>Cisco proprietary Inter-Switch Link (ISL) format</u>.
- **IEEE 802.1Q should be used whenever possible**, because it is the defacto standard and <u>802.1Q and ISL are not compatible</u>.

Troubleshooting switches: VLANs

- When troubleshooting inter-VLAN issues**, ensure that there is no IP address on the physical interface of the router**. The interface must be active.
- To verify the interface configuration, use:
  - `show ip interface brief`
- The network associated with each VLAN should be **visible in the routing table**. If not, recheck all physical connections and trunk configuration on both ends of the link.
- If it is not directly connected to the VLAN subnets, check the configuration of the routing protocol to verify that there is a route to each of the VLANs. Use the command:
  - `show ip route`

Troubleshooting switches: VLANs

- **Access or Trunk Port**
  - <u>Each switch port is either an access port or a trunk port.</u>
  - On some switch models, **other switch port modes are available** and the switch automatically configures the port to the appropriate status.
  - It is sometimes advisable to **lock the port into either <u>access</u> or <u>trunk</u> status** to avoid potential problems with this detection process.
- **Native and Management VLANs**
  - <u>The native VLAN and management VLAN are VLAN1 by default.</u>
  - **Untagged frames** sent across a trunk are assigned to the native VLAN of the trunk line.
  - <u>If the native VLAN assignment is changed on a device</u>, each end of the 802.1Q trunk should be configured with **the same native VLAN number**.
    - If one end of the trunk is configured for native VLAN10 and the other end is configured for native VLAN14, a frame sent from VLAN10 on one side is received on VLAN14 on the other. VLAN10 "leaks" into VLAN14. This can create unexpected connectivity issues and increase latency.
- For smoother, quicker transitions, <u>verify that the native VLAN assignment is the same on all devices throughout the network</u>.

Troubleshooting switches: VTP

- To display the VLAN Trunk Protocol (VTP) version in use on a device, the VTP domain name, the VTP mode, and the VTP revision number, issue the command:
  - **`show vtp status`**
- To modify the VTP version number, use:
  - **`vtp version <1 | 2>`**
- It is also a problem if a rogue switch joins the domain and modifies VLAN information. To prevent this situation, it is important to **configure a password** on the VTP domain with the global configuration command:
  - **`vtp password`** ***password***
- When configured, the **authentication password must be the same on all devices** in the VTP domain.
- If updates are not propagating to a new switch in the VTP domain, suspect the password. To verify the password, use the command:
  - **`show vtp password`**

Troubleshooting EIGRP

- A number of IOS show commands and debug commands are the same for troubleshooting EIGRP routing issues as they are for RIP. Commands specific to troubleshooting EIGRP include:
- **show ip eigrp neighbors**
  - Displays neighbor IP addresses and the interface on which they were learned.
- **show ip eigrp topology**
  - Displays the topology table of known networks with successor routes, status codes, feasible distance, and interface.
- **show ip eigrp traffic**
  - Displays EIGRP traffic statistics for the AS configured, including hello packets sent/received, updates, and so on.
- **debug eigrp packets**
  - Displays real-time EIGRP packet exchanges between neighbors.
- **debug ip eigrp**
  - Displays real-time EIGRP events, such as link status changes and routing table updates.

Troubleshooting OSPF

- In addition to the standard show and debug commands, the following commands assist troubleshooting OSPF issues:
  - **`show ip ospf`**
  - **`show ip ospf neighbor`**
  - **`show ip ospf interface`**
  - **`debug ip ospf events`**
  - **`debug ip ospf packet`**

Troubleshooting Route Redistribution

- With each routing protocol, configure a default quad 0 static route **on the edge router**:
  - **`ip route 0.0.0.0 0.0.0.0 interface`**
- Next, configure the edge router to send or propagate its default route to the other routers.
  - With **RIP** and **OSPF**, enter router configuration mode and use the command **default-information originate**.
  - **EIGRP** redistributes default routes directly; the **redistribute static** command can also be used.
- Failure to properly implement default route redistribution prevents users that are connected to internal routers from accessing external networks.

Troubleshooting WAN

- To display the type of cable and the detection and status of DTE, DCE, and clocking, use the following command:
  - **`show controllers <serial_port>`**
- To see the encapsulation in use on a serial line, use the command:
  - **`show interfaces <serial_port>`**
- The IP address configured on an interface and the status of the port and line protocol is viewable with the command:
  - **`show ip interface brief`**
- If the **interface is up but the line protocol is down**, check that the proper cable is connected and is firmly attached to the port. If this step still does not correct the problem, replace the cable.
- If the status of an interface is **administratively down**, the most probable cause is that the **no shutdown** command was not entered on the interface. Interfaces are shutdown by default.

Troubleshooting WAN

- When **troubleshooting PPP connectivity,** verify that:
  - LCP phase is complete
  - Authentication has passed, if configured
  - NCP phase is complete
- To show the status of the LCP and NCP phase, use:
  - **`show interface`**
- To display PPP packets transmitted during the startup phase where PPP options are negotiated, use:
  - **`debug ppp negotiation`**
- To display real-time PPP packet flow, use:
  - **`debug ppp packet`**

Troubleshooting WAN

- If using PAP authentication on a current version of the IOS, activate it with the command:
  - **ppp pap sent-username *user* password *pass***
- Debugging the authentication process provides a quick method of determining what is wrong.
- To display packets involved in the authentication process as they are exchanged between end devices, use the command:
  - **debug ppp authentication**

Troubleshooting ACLs

- When networks or hosts become unreachable and ACLs are in use, it is critical to determine if the ACL is the problem. Ask the following **questions to help to isolate the problem**:
  - **Is an ACL applied** to the problem router or interface?
  - **Has it been applied recently**?
  - Did the **issue exist before** the ACL was applied?
  - Is the ACL performing as expected?
  - Is the problem with all hosts connected to the interface or **only specific hosts**?
  - Is the problem with all protocols being forwarded or **only specific protocols**?
  - Are the networks appearing **in the routing table as expected**?
- One way to determine the answer to several of these questions is to **enable logging**.
- Logging shows the effect that ACLs are having on various packets. By default, the number of matches display with the **`show access-list`** command.
- To display details about packets permitted or denied, add the **`log`** keyword to the end of ACL statements.

Troubleshooting ACLs

- To <u>display all ACLs</u> configured on the router, whether applied to an interface or not, use the following command:
  - **`show access-lists`**
- To <u>clear the number of matches</u> for each ACL statement, use:
  - **`clear access-list counters`**
- To <u>display the source and destination IP address for each packet received or sent by any interface on the router</u>, use:
  - **`debug ip packet`**

Troubleshooting ACLs

- In some cases, the ACL may permit or deny the intended traffic but <u>can also have unintended effects</u> on other traffic.
- If it appears that the ACL is the problem, there are several issues to check.
  - If the ACL statements are **<u>not in the most efficient order</u>** to permit the highest volume traffic early in the ACL, <u>check the logging results to determine a more efficient order</u>.
- The <u>implicit deny may be having unintended effects</u> on other traffic.
  - If so, use an explicit **`deny ip any any log`** command so that packets that do not match any of the previous ACL statements can be monitored.

Troubleshooting ACLs

- In addition to determining whether the ACL is correctly configured, it is also important to apply the ACL to the right router or interface, and in the appropriate direction. A correctly configured ACL that is incorrectly applied is one of the most common errors when creating ACLs.
- **Standard ACLs** filter only on the source IP address; therefore, place them **as close to the destination as possible**.
  - Placing a Standard ACL close to the source may unintentionally block traffic to networks that should be allowed.
  - Placing the ACL close to the destination unfortunately allows traffic to flow across one or more network segments prior to being denied. This is a waste of valuable bandwidth.
- Using an Extended ACL resolves both of these issues.

End of lesson