

CompTIA®



The Official CompTIA

Network+

Study Guide

Exam N10-009



Official CompTIA Content Series for CompTIA Performance Certifications

**The Official
CompTIA
Network+
Instructor Guide
(Exam N10-009)**

Acknowledgments



James Pengelly, Author

Juan Guerrero, Author

Katie Hoenicke, Senior Vice President, Product Development

Becky Mann, Vice President, Content Development

Craig Jenkins, Senior Director, Technical Content Development

Michael Olsen, Director, Content Development

Lizbeth Johnson, Senior Manager, Content Development

Notices

Disclaimer

While The Computing Technology Industry Association ("CompTIA") takes care to ensure the accuracy and quality of these materials, we cannot guarantee their accuracy, and all materials are provided without any warranty whatsoever, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. The use of screenshots, photographs of another entity's products, or another entity's product name or service in this book is for editorial purposes only. No such use should be construed to imply sponsorship or endorsement of the book by nor any affiliation of such entity with CompTIA. This courseware may contain links to sites on the Internet that are owned and operated by third parties (the "External Sites"). CompTIA is not responsible for the availability of, or the content located on or through, any External Site. Please contact CompTIA if you have any concerns regarding such links or External Sites.

Trademark Notice

CompTIA®, Network+®, and the CompTIA logo are registered trademarks of CompTIA, Inc., in the United States and other countries. All other product and service names used may be common law or registered trademarks of their respective proprietors.

Copyright Notice

Copyright © 2024 CompTIA, Inc. All rights reserved. Screenshots used for illustrative purposes are the property of the software proprietor. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of CompTIA, 3500 Lacey Road, Suite 100, Downers Grove, IL 60515-5439.

This book conveys no rights in the software or other products about which it was written; all use or licensing of such software or other products is the responsibility of the user according to terms and conditions of the owner. If you believe that this book, related materials, or any other CompTIA materials are being reproduced or transmitted without permission, please call 1-866-835-8020 or visit help.comptia.org.

Table of Contents

Module 1: Explaining Network Topologies	1
Lesson 1.1: Networking Overview	2
Lesson 1.2: OSI Model Concepts	7
Lesson 1.3: SOHO Networks	16
Lesson 1.4: Troubleshooting Methodology	24
 Module 2: Supporting Cabling and Physical Installations.....	 33
Lesson 2.1: Ethernet	34
Lesson 2.2: Copper Cables and Connectors	39
Lesson 2.3: Wiring Implementation.....	45
Lesson 2.4: Fiber Optic Cables and Connectors.....	51
Lesson 2.5: Physical Installation Factors	58
Lesson 2.6: Cable Troubleshooting	64
 Module 3: Configuring Interfaces and Switches.....	 73
Lesson 3.1: Network Interfaces	74
Lesson 3.2: Ethernet Switches.....	81
Lesson 3.3: Switch Port Configuration	88
Lesson 3.4: Switch Troubleshooting	95
 Module 4: Configuring Network Addressing.....	 105
Lesson 4.1: Internet Protocol Basics.....	106
Lesson 4.2: IP Version 4 Addressing.....	113
Lesson 4.3: IP Version 4 Subnetting.....	123
Lesson 4.4: IP Troubleshooting Tools	133
Lesson 4.5: IP Version 6	139
Lesson 4.6: IP Troubleshooting.....	147

Module 5: Configuring Routing and Advanced Switching	153
Lesson 5.1: Routing Technologies	154
Lesson 5.2: Dynamic Routing Technologies	164
Lesson 5.3: Network Address Translation	171
Lesson 5.4: Firewalls	176
Lesson 5.5: Enterprise Network Topologies	179
Lesson 5.6: Virtual LANs	183
Lesson 5.7: Routing and VLAN Troubleshooting	191
 Module 6: Implementing Network Services	 197
Lesson 6.1: Transport and Application Layer Protocols	198
Lesson 6.2: Dynamic Host Configuration Protocol	207
Lesson 6.3: APIPA and SLAAC	211
Lesson 6.4: DHCP Relay and Troubleshooting	215
Lesson 6.5: Domain Name System	218
Lesson 6.6: DNS Troubleshooting	229
 Module 7: Explaining Application Services	 235
Lesson 7.1: Application Security and Time Synchronization	236
Lesson 7.2: Web, File/Print, and Database Services	240
Lesson 7.3: Email and Voice Services	248
Lesson 7.4: Disaster Recovery and High Availability	254
 Module 8: Supporting Network Management	 265
Lesson 8.1: Organizational Policies and Documentation	266
Lesson 8.2: Host Discovery and Monitoring	276
Lesson 8.3: Simple Network Management Protocol	283
Lesson 8.4: Event Management	286
Lesson 8.5: Packet Capture and Analysis	293
Lesson 8.6: Traffic Monitoring	297

Module 9: Explaining Network Security Concepts	307
Lesson 9.1: Security Concepts	308
Lesson 9.2: Network Threats and Attacks	315
Lesson 9.3: Spoofing Attacks	320
Lesson 9.4: Rogue System Attacks	323
Lesson 9.5: Social Engineering	326
 Module 10: Applying Network Security Features	 331
Lesson 10.1: Authentication	332
Lesson 10.2: Authorization and Account Management	343
Lesson 10.3: Network Hardening	347
Lesson 10.4: Switch Security	350
Lesson 10.5: Network Security Rules	355
 Module 11: Supporting Network Security Design	 363
Lesson 11.1: Zone-Based Security	364
Lesson 11.2: Internet of Things	369
Lesson 11.3: Physical Security	373
 Module 12: Configuring Wireless Networks	 379
Lesson 12.1: Wireless Concepts and Standards	380
Lesson 12.2: Enterprise Wireless Network Design	388
Lesson 12.3: Wireless Security	396
Lesson 12.4: Wireless Troubleshooting	403
 Module 13: Comparing Remote Access Methods	 413
Lesson 13.1: WAN and Internet Connectivity	414
Lesson 13.2: Virtual Private Networks	417
Lesson 13.3: Remote Management	425

Module 14: Summarizing Cloud Concepts..... 433

Lesson 14.1: Datacenter and Storage Networks 434

Lesson 14.2: Cloud Concepts..... 439

Lesson 14.3: Cloud Networking..... 443

Lesson 14.4: Modern Network Environments 449

Appendix A: Mapping Course Content to CompTIA Certification A-1

Glossary.....G-1

Index..... I-1

About This Course

CompTIA, a not-for-profit trade association, is here to help you get the tech career you deserve with industry-leading certifications, courses, and expert knowledge. Today's job market demands individuals have demonstrable skills, and the information and activities in this course can help you build your network administration skill set so that you can confidently perform your duties in any entry-level network support technician role.

With CompTIA Network+, you can unlock a diverse range of career paths, from network administration and support to cybersecurity, creating opportunities for advancement and specialization in the rapidly evolving IT industry. CompTIA Network+ is a global IT certification validating candidates have the core skills necessary to establish, maintain, troubleshoot, and secure networks in any environment, regardless of technology and platform.

This course can prepare you for the CompTIA Network+ (Exam N10-009) certification examination and a job role in network administration. It utilizes a learning progression model to help you learn and build skills related to the course objectives and job task requirements. This learning methodology uses a series of steps to contextualize what you're learning, elaborate on areas where additional instruction is needed, and provide relevance through practice and personalized feedback. You'll then apply what you learned and demonstrate the skills you've gained through a series of lab activities and quizzes.

On course completion, you will be able to do the following:

- Deploy and troubleshoot Ethernet networks.
- Support IPv4 and IPv6 networks.
- Configure and troubleshoot routers.
- Support network services and applications.
- Ensure network security and availability.
- Deploy and troubleshoot wireless networks.
- Support WAN links and remote access methods.
- Support organizational procedures and site security controls.
- Summarize cloud and datacenter architecture.

Using the Course

Course Design

This course is designed to optimize knowledge acquisition and skills development related to the learning objectives and associated job task requirements through a learning progression model. The learning progression model follows a series of steps to contextualize, elaborate, and provide relevance through practice and personalized feedback, contextualized application, and demonstrable evidence of skills gained.

Different activities throughout the course will help you practice and develop your skills as well as gauge your understanding of the various topics covered. The course is broken into modules and lessons. After each lesson, there will be a review as well as lab activities to try the skills.

Prerequisites

To ensure your success in this course, you should have basic IT skills comprising nine to 12 months' experience. CompTIA A+ certification, or the equivalent knowledge, is strongly recommended.



The prerequisites for this course might differ significantly from the prerequisites for the CompTIA certification exams. For the most up-to-date information about the exam prerequisites, complete the form on this page: www.comptia.org/training/resources/exam-objectives.

How to Use the Study Notes

The following sections will help you understand how the course structure and components are designed to support mastery of the competencies and tasks associated with the target job roles and will help you to prepare to take the certification exam.

As You Learn

At the top level, this course is divided into **modules**, each representing an area of competency within the target job roles. Each lesson is composed of a number of topics. A **lesson** contains subjects that are related to a discrete job task, mapped to objectives and content examples in the CompTIA exam objectives document. Rather than follow the exam domains and objectives sequence, modules and lessons are arranged in order of increasing proficiency. Each lesson is intended to be studied within a short period (typically 30 minutes at most). Each lesson is concluded by one or more activities, designed to help you to apply your understanding of the study notes to practical scenarios and tasks.

In addition to the study content in the lessons, there is a glossary of the terms and concepts used throughout the course. There is also an index to assist in locating particular terminology, concepts, technologies, and tasks within the lesson and topic content.



In many electronic versions of the book, you can click links on keywords in the topic content to move to the associated glossary definition, and on page references in the index to move to that term in the content. To return to the previous location in the document after clicking a link, use the appropriate functionality in your eBook viewing software.

Watch throughout the material for the following visual cues.

Student Icon	Student Icon Descriptive Text
	A Note provides additional information, guidance, or hints about a topic or task.
	A Caution note makes you aware of places where you need to be particularly careful with your actions, settings, or decisions so that you can be sure to get the desired results of an activity or task.

As You Review

Any method of instruction is only as effective as the time and effort you, the student, are willing to invest in it. In addition, some of the information that you learn in class may not be important to you immediately, but it may become important later. For this reason, we encourage you to spend some time reviewing the content of the course after your time in the classroom.

Following the lesson content, you will find a table mapping the modules and lessons to the exam domains, objectives, and content examples. You can use this as a checklist as you prepare to take the exam and to review any content that you are uncertain about.

As a Reference

The organization and layout of this book make it an easy-to-use resource for future reference. Guidelines can be used during class and as after-class references when you're back on the job and need to refresh your understanding. Taking advantage of the glossary, index, and table of contents, you can use this book as a first source of definitions, background information, and summaries.

Module 1

Explaining Network Topologies

Module Introduction

Computer networks are complex systems that incorporate multiple functions, standards, and proprietary technologies. The Open Systems Interconnection (OSI) model is used to try to simplify some of this complexity. It divides network technologies among seven functional layers. This makes it easier to separate and focus on individual concepts and technologies while retaining an understanding of relationships between functions of technologies placed in other layers.

This module uses the OSI model to give you an overview of the technologies that you will be studying in the rest of the course. You will compare the functions of these layers in the OSI model and apply those concepts to the installation and configuration of a small office/home office network.

You will also learn how to apply a methodology to structure troubleshooting activity, so that you can diagnose and resolve problems efficiently.

Module Objectives

In this lesson, you will do the following:

- Explain network types and characteristics.
- Compare and contrast OSI model layers.
- Configure SOHO networks.
- Explain CompTIA's troubleshooting methodology.

Lesson 1.1

Networking Overview



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 1.6 Compare and contrast network topologies, architectures, and types.

Network types and topologies determine the scale and flow of data through a network. Identifying the different topologies is essential to designing or troubleshooting a network. No matter what your specific role in network implementation and management, you will need to understand the characteristics of the network topology you are working with and identify how the topology affects network design, performance, and troubleshooting.

As you study this lesson, answer the following questions:

- Why is it useful to categorize network types, appliances, applications, functions, and topologies?
- What are the advantages of a client-server network compared to a peer-to-peer network?
- How does a logical topology differ from a physical topology?
- What do diagrams of point to point, star, and mesh topologies look like?

Networking Concepts

A network is two or more computer systems that are linked by a transmission medium and share one or more protocols that enable them to exchange data. You can think of any network in terms of nodes and links. The nodes are devices that send, receive, and forward data, and the links are the communications pathways between them.

There are two general kinds of nodes: intermediate nodes and end systems. Intermediate nodes perform a forwarding function, while end system nodes are those that send and receive data traffic. End systems are often also referred to as hosts.

Client-Server Versus Peer-to-Peer Networks

End system nodes can be classified as either clients or servers:

- A server makes network applications and resources available to other hosts.
- A client consumes the services provided by servers.

A **client-server** network is one where some nodes, such as PCs, laptops, and smartphones, act mostly as clients. The servers are more powerful computers. Application services and resources are centrally provisioned, managed, and secured.

A **peer-to-peer** network is one where each host acts as both client and server. This is a decentralized model where provision, management, and security of services and data are distributed around the network. A small peer-to-peer network can also be referred to as a workgroup.

Business and enterprise networks are typically client-server, while residential networks are more often peer-to-peer. However, note that in a client-server network, often, hosts will function as both clients and servers at the same time. For example, a computer hosting a web application acts as a server to browser clients but is itself a client of database services running on other server computers. It is the centrally administered nature of the network that really defines it as client-server.

Appliances, Applications, and Functions

You can also think of a network as having appliances, applications, and functions:

- **Appliances**—Networks make use of many types of specialized platforms. Unlike general-purpose Windows or macOS computers and servers, an appliance is a computer with an operating system and software designed to perform a particular network role. Examples of these roles include the switches, routers, and wireless access points that forward data, the firewalls and intrusion detection systems that enforce security rules, and the load balancers and proxies that improve network performance.

An appliance can be deployed as **physical hardware**, meaning that the appliance OS/software runs on its own CPU, memory, storage, and network interfaces. It is also possible to deploy **virtual appliances**. This means that the appliance OS is deployed as a virtual machine running on a hypervisor computing platform. The same hypervisor computer could run multiple virtual appliances.

- **Applications**—The nodes and links of networking infrastructure are deployed to run services. Services are shared applications that allow the network to do useful work, such as sharing files or allowing employees to send email.
- **Functions**—Networks can be configured with additional properties to perform different functions. For example, the security properties of a virtual private network allow devices to join a local network from across the Internet. As another example, quality of service functionality allows optimization of a network to suit a particularly time-sensitive application, such as voice or video.

Network Types

A network type refers primarily to its size and scope. The size of a network can be measured as the number of nodes, while the scope refers to the area over which nodes sharing the same network address are distributed.

Local Area Networks

A **local area network (LAN)** is confined to a single geographical location. In a LAN, all nodes and segments are directly connected with cables or short-range wireless technologies. Most of the network infrastructure in a LAN would be directly owned and managed by a single organization. Some examples of LANs include the following:

- **Home/residential network**—With an Internet router and a few computers, plus mobile devices, gaming consoles, and printers.

- **Small office/home office (SOHO) network**—A business-oriented network possibly using a centralized server in addition to client devices and printers, but often still using a single Internet router/switch/access point to provide connectivity.
- **Small and medium-sized enterprise (SME) Network**—A network supporting dozens of users. Such networks would use structured cabling and multiple switches and routers to provide connectivity.
- **Enterprise LAN**—A larger network with hundreds or thousands of servers and clients. Such networks would require multiple enterprise-class switch and router appliances to maintain performance levels.
- **Datacenter**—A network that hosts only servers and storage, not end user client devices.

Wide Area Networks

A **wide area network (WAN)** is a network of networks, connected by long-distance links. A typical enterprise WAN would connect a main office site with multiple branch office sites, possibly in different countries. A WAN could link two or more large LANs or could be used for remote workers connecting to an enterprise network via a public network such as the Internet. WAN links are also used to connect datacenters to one another. WANs are likely to use leased network devices and links, operated and managed by a service provider.

Network Topology

Where the type defines the network scope, the **topology** describes the physical or logical structure of the network in terms of nodes and links.

A network's physical topology describes the placement of nodes and how they are connected by the transmission media. For example, in one network, nodes might be directly connected via a single cable; in another network, each node might connect to a switching appliance via separate cables. These two networks have different physical topologies.

The logical topology describes the flow of data through the network. For example, given the different physical network topologies described previously, if in each case the nodes can send messages to one another, the logical topology is the same. The different physical implementations—directly connected via a cable versus connected to the same switch—achieve the same logical layout.

In the simplest type of topology, a single link is established between two nodes. This is called a **point to point** link. Because only two devices share the connection, they are guaranteed a level of **bandwidth**.

1

In a point-to-point (or duplex) network, only two nodes are connected to the network media.



2

When the network media is half-duplex, a node cannot transmit and receive at the same time.



3

When the network media is full-duplex, nodes can transmit and receive simultaneously.



Physical point to point topologies using different media types for half-duplex and duplex communications.


A point to point link can be a physical or logical topology. For example, on a WAN, two router appliances might be physically linked via multiple intermediate networks and physical devices but still share a logical point to point link, where each can address only the other router. With either a physical or logical topology, it is the 1:1 relationship that defines a point to point link.

Star Topology

In a **star topology**, each endpoint node is connected to a central forwarding appliance, such as a switch or router. The central node mediates communications between the endpoints. The star topology is the most widely used physical topology. For example, a typical SOHO network is based around a single Internet router appliance that clients can connect to with a cable or wirelessly. The star topology is easy to reconfigure and easy to troubleshoot because all data goes through a central point, which can be used to monitor and manage the network. Faults are automatically isolated to the media, node (network card), or the switch, router, or wireless access point at the center of the star.

1

In a star topology, each node is connected to a concentrator over dedicated network media.



2


When a node transmits, the signal is sent over the network media to the concentrator.

3

The concentrator forwards signals to other nodes. It may do this by repeating the signal to all nodes, working as logical bus topology (implemented as an Ethernet hub).

4

More commonly, the concentrator tracks node addresses and selectively switches communication paths to establish point-to-point links.



5


The concentrator is a single point of failure. Also, each node must be within the media's maximum transmission distance of the concentrator or must use a repeater to overcome the distance limitation.

Star topologies using different types of concentrator.

You may also encounter the **hub-and-spoke** topology, which has the same physical layout as a star topology but is primarily used in a different context. While the star topology is often seen in local area networks (LANs), the hub-and-spoke topology is more commonly applied to wide area networks (WANs) with remote sites.

Mesh Topology

A **mesh topology** is commonly used in WANs, especially public networks such as the Internet. A full mesh network requires that each device has a point to point link with every other device on the network. This approach is normally impractical, however. The number of links required by a full mesh is expressed as $n(n-1)/2$, where n is the number of nodes. For example, a network of just four nodes would require six links, while a network of 40 nodes would need 780 links! Consequently, a hybrid approach is often used, with only the most important devices interconnected in the mesh, perhaps with extra links for fault tolerance and redundancy. This type of topology is referred to as a partial mesh.



Fully connected and partial mesh topology examples.

Mesh networks provide excellent redundancy, because other routes, via intermediary devices, are available between locations if a link failure occurs.

Lesson 1.2

OSI Model Concepts



Exam Objectives Covered

1.1 Explain concepts related to the Open Systems Interconnection (OSI) reference model.


Networks are built on common standards and models that describe how devices and protocols interconnect. In this lesson, you will identify how the implementation and support of these systems refer to an important common reference model: the Open Systems Interconnection (OSI) model. The OSI model breaks the data communication process into discrete layers. Being able to identify the OSI layers and compare the functions of devices and protocols working at each layer will help you to implement and troubleshoot networks.

As you study this lesson, answer the following questions:

- Why are protocols important for networking?
- What is the OSI model? Why is it important for understanding networking?
- At which OSI model layer do common network appliances, applications, and functions work?
- How does data encapsulation facilitate data transmission? How does encapsulation relate to frames, datagrams, segments, and ports?

Open Systems Interconnection Model

The International Organization for Standardization (ISO) developed the **Open Systems Interconnection (OSI) reference model** ([iso.org/standard/20269.html](https://www.iso.org/standard/20269.html)) to promote understanding of how components in a network system work. It does this by separating the functions of hardware and software components into seven discrete layers. Each layer performs a different group of tasks required for network communication.



The OSI model.


Although not all network systems implement layers using this precise structure, they all implement each task in some way. The OSI model is not a standard or a specification; it serves as a functional guideline for designing network protocols, software, and appliances and for troubleshooting networks.

Data Encapsulation and Decapsulation

A **network protocol** is a set of rules for exchanging data in a structured format. A network protocol has two principal functions:

- **Addressing**—Describing where data messages should go. At each OSI model layer, there are different mechanisms for identifying nodes and rules for how they can send and receive messages.
- **Encapsulation**—Describing how data messages should be packaged for transmission. Encapsulation is like an envelope for a letter, with the distinction that each layer requires its own envelope. At each layer, the protocol adds fields in a header to whatever payload data it receives from an application or other protocol.

A network will involve the use of many different protocols operating at different layers of the OSI model. At each layer, for two nodes to communicate they must be running the same protocol. The protocol running at each layer communicates with its peer layer on the other node. This communication *between* nodes at the same layer is described as a same layer interaction. To transmit or receive a communication, *on* each node, each layer provides services for the layer above and uses the services of the layer below. This is referred to as adjacent layer interaction.



Encapsulation and decapsulation. (Images © 123RF.com.)

When a message is sent from one node to another, it travels down the stack of layers on the sending node, reaches the receiving node using the transmission media, and then passes up the stack on that node. At each level (except the Physical layer), the sending node adds a header to the data payload, forming a “chunk” of data called a **protocol data unit (PDU)**. This is the process of encapsulation.

For example, on the sending node, data is generated by an application, such as the HyperText Transfer Protocol (HTTP), which will include its own application header. At the Transport layer, a Transmission Control Protocol (TCP) header is added to this application data. At the Network layer, the TCP segment is wrapped in an Internet Protocol (IP) header. The IP packet is encapsulated in an Ethernet frame at the Data Link layer, then the stream of bits making up the frame is transmitted over the network at the Physical layer as a modulated electrical signal.

The receiving node performs the reverse process, referred to as decapsulation. It receives the stream of bits arriving at the Physical layer and decodes an Ethernet frame. It extracts the IP packet from this frame and resolves the information in the IP header, then does the same for the TCP and Application headers, eventually extracting the HTTP application data for processing by a software program, such as a web browser or web server.



You might notice that this example seems to omit some OSI layers. This is because “real-world” protocols do not conform exactly to the OSI model.

Layer 1 - Physical

The **Physical layer (PHY)** of the OSI model is defined as layer 1. The Physical layer is responsible for the transmission and receipt of the signals that represent bits of data. Transmission media can be classified as cabled or wireless:

- **Cabled**—A physical signal conductor is provided between two nodes. Examples include copper or fiber optic cable types. Cabled media can also be described as bounded media.
- **Wireless**—Uses free space between nodes, such as microwave radio. Wireless media can also be described as unbounded media.

The Physical layer specifies the following:

- **Physical topology**—The layout of nodes and links as established by the transmission media. An area of a larger network is called a segment. A network is typically divided into segments to cope with the physical restrictions of the network media used, to improve performance, or to improve security. At the Physical layer, a segment is where all the nodes share access to the same media.
- **Physical interface**—Mechanical specifications for the network medium. For cabled media, this means the construction of the cable, the interface/connector form factor, and the number and functions of the pins in a connector. For wireless media, it means radio transceiver and antenna specifications.
- **Signaling**—The process of transmitting and receiving encoded data over the network medium. A modulation scheme describes how electrical, light, or radio signals represent bits. Timing and synchronization schemes ensure senders and receivers can identify groups of signals as a chunk or frame of data.

Devices that operate at the Physical layer include the following:


- **Transceiver**—The part of a network interface that sends and receives signals over the network media.
- **Repeater**—A device that amplifies an electronic signal to extend the maximum allowable distance for a media type.
- **Hub**—A multiport repeater, deployed as the central point of connection for nodes.
- **Media converter**—A device that converts one media signaling type to another.

Layer 2 - Data Link

Layer 2 is referred to as the **Data Link layer**. It is responsible for transferring data between nodes on the same logical segment. At the Data Link layer, a segment is one where all nodes can send traffic to one another using hardware addresses, regardless of whether they share access to the same media. A layer 2 segment might include multiple physical segments. This is referred to as a logical topology.

Local networks do not typically connect hosts directly with point to point or mesh links. To reduce cabling and interface costs, each host is connected to a central node, such as a switch or a wireless access point. The central node provides a forwarding function, receiving the communication from one node and sending it to another. To do this, each node interface must have a Data Link layer address. The addresses of interfaces within the same layer 2 segment are described as local addresses or hardware addresses.

The Data Link layer also performs an encapsulation function. It organizes the stream of bits arriving from the Physical layer into structured units called **frames**. Each frame contains a Network layer packet as its payload. The Data Link layer adds control information to the payload in the form of header fields. These fields include source and destination hardware addresses, plus a basic error check to test if the frame was received intact.




Communications at layer 2 of the OSI model. (Images © 123RF.com.)

Devices that operate at the Data Link layer include the following:

- **Network adapter or network interface card (NIC)**—A NIC joins an end system host to network media (cabling or wireless) and enables it to communicate over the network by assembling and disassembling frames.
- **Bridge**—A bridge is a type of intermediate system that joins physical network segments while minimizing the performance reduction of having more nodes on the same network. A bridge has multiple ports, each of which functions as a network interface.
- **Switch**—An advanced type of bridge with many ports. A switch creates links between large numbers of nodes more efficiently.
- **Wireless access point (AP)**—An AP allows nodes with wireless network cards to communicate and creates a bridge between wireless networks and wired ones.

Layer 3 - Network

Layer 3 is the **Network layer**. This layer is responsible for moving data around a network of networks, known as an internetwork. While the Data Link layer is capable of forwarding data by using hardware addresses within a single segment, the Network layer moves information around an internetwork by using logical network and host IDs. The networks are often heterogeneous; that is, they use a variety of Physical layer media and Data Link protocols. The main appliance working at layer 3 is the **router**.



Communications at layer 3 of the OSI model. (Images © 123RF.com.)

At layer 3, each packet is given a destination network address. Routers are configured with information about how to reach these different logical networks. The packet is forwarded, router by router (or hop by hop), through the internetwork to the target network. Once it has reached the destination network, the hardware address can be used to deliver the packet to the target node.




The general convention is to describe PDUs packaged at the Network layer as packets or datagrams and messages packaged at the Data Link layer as frames. Packet is often used to describe PDUs at any layer, however.

It is usually important for traffic passing between networks to be filtered. A basic **firewall** operates at layer 3 to enforce an **access control list (ACL)**. A network ACL is a list of the addresses and types of traffic that are permitted or blocked.

Layer 4 - Transport

The first three layers of the OSI model are primarily concerned with moving frames and datagrams between nodes and networks. At the **Transport layer**—also known as the end-to-end or host-to-host layer—the content of the packets becomes significant. Any given host on a network will be communicating with many other hosts using many different types of networking data. One of the functions of the Transport layer is to identify each type of network application by assigning it a port number. For example, data requested from an HTTP web application can be identified as port 80, while data sent to an email server can be identified as port 25.

At the Transport layer, on the sending host, data from the upper layers is packaged as a series of layer 4 PDUs, referred to as segments. Each segment is tagged with the application's port number. The segment is then passed to the Network layer for delivery. Many different hosts could be transmitting multiple HTTP and email packets at the same time. These are multiplexed using the port numbers along with the source and destination network addresses onto the same link.



Communications at layer 4 (Transport layer) of the OSI model. (Images © 123RF.com.)

At the Network and Data Link layers, the port number is ignored—it becomes part of the data payload and is invisible to the routers and switches that implement the addressing and forwarding functions of these layers. At the receiving host, each segment is decapsulated, identified by its port number, and passed to the relevant handler at the Application layer. Put another way, the traffic stream is de-multiplexed.

The Transport layer can also implement reliable data delivery mechanisms, should the application require it. Reliable delivery means that any lost or damaged packets are resent.

Devices working at the Transport layer include multilayer switches—usually working as load balancers—and many types of security appliances, such as more advanced firewalls and intrusion detection systems (IDSs).

Upper Layers

The upper layers of the OSI model are less clearly associated with distinct real-world protocols. These layers collect various functions that provide useful interfaces between software applications and the Transport layer.

Layer 5—Session

Most application protocols require the exchange of multiple messages between the client and server. This exchange of such a sequence of messages is called a session or dialog. The **Session layer (layer 5)** represents functions that administer the process of establishing a dialog, managing data transfer, and then ending (or tearing down) the session.

Layer 6—Presentation

The **Presentation layer (layer 6)** transforms data between the format required for the network and the format required for the application. For example, the Presentation layer is used for character set conversion, such as between American Standard Code for Information Interchange (ASCII) and Unicode.



The Presentation layer can also be conceived as supporting data compression and encryption. However, in practical terms, encryption is implemented by devices and protocols running at lower layers of the stack or simply within a homogenous Application layer.


Layer 7—Application

The **Application layer (layer 7)** is at the top of the OSI stack. An Application layer protocol doesn't encapsulate any other protocols or provide services to any protocol. Application layer protocols provide an interface for software programs on network hosts that have established a communications channel through the lower-level protocols to exchange data.

More widely, upper-layer protocols provide most of the services that make a network useful, rather than just functional, including web browsing, email and communications, directory lookup, remote printing, and database services.

OSI Model Summary

The following image summarizes the OSI model, listing the PDUs at each layer, along with the types of devices that work at each layer.



Devices and concepts represented at the relevant OSI model layer.

Lesson 1.3

SOHO Networks



Exam Objectives Covered

1.1 Explain concepts related to the Open Systems Interconnection (OSI) reference model.

The OSI model involves quite a lot of abstraction. As a practical example, it is worth examining how a basic network is implemented. In this topic, you will describe the connection and configuration options for components within a typical small office/home office (SOHO) router. You will also explore some of the organizations responsible for managing the Internet and the various numbering schemes used for network address notation.

As you study this lesson, answer the following questions:

- What separate functions are packed into a typical SOHO router appliance?
- How are Internet services provisioned and managed?
- What types of numbering schemes are used to represent network addresses?

SOHO Routers

Networks of different sizes are classified in different ways. A network in a single location is often described as a local area network (LAN). This definition encompasses many different sizes of networks with widely varying functions and capabilities. It can include both residential networks with a couple of computers and enterprise networks with hundreds of servers and thousands of workstations.


Small office/home office (SOHO) is a category of LAN with a small number of computing hosts that typically rely on a single integrated appliance for local and Internet connectivity.

Networks such as the Internet that are located in different geographic regions but with shared links are called wide area networks (WANs). The intermediate system powering SOHO networks is usually described as a SOHO router because one of its primary functions is to forward traffic between the LAN and the WAN. However, routing is actually just one of its functions. We can use the OSI model to analyze each of these in turn.

Physical Layer Functions

Starting at layer 1, the SOHO router provides the following physical interfaces:

- A number of RJ45 ports (usually four) to implement a local cabled network. These are typically labeled as the LAN ports.
- Radio antennae to transmit and receive wireless signals.
- A type of modem to connect to the **Internet service provider's (ISP's)** network. This is typically labeled as the WAN port. In the example in the diagram, the WAN interface is another RJ45 port designed to connect to a fiber to the premise's Internet service using the same Ethernet technology as the local network. On other SOHO routers, there may be a different type of WAN modem, such as an RJ11 port to connect to a digital subscriber line (DSL) service, or a F-connector coax jack to connect to a cable service.




Physical layer connectivity options on a SOHO router.

Data Link Layer Functions

At layer 2, the SOHO router implements the following functions to make use of its Physical layer adapters:

- **Ethernet switch**—The RJ45 jacks are connected internally by an Ethernet switch.
- **Wireless access point**—The radio antennae implement some version of the Wi-Fi standard. The access point functions as a wireless hub, allowing stations (PCs, tablets, smartphones, and printers) to form a wireless network. The access point is also wired to the Ethernet switch via an internal port. This forms a bridge between the cabled and wireless segments, creating a single logical Data Link network.




Data Link layer local network segment. (Images © 123RF.com.)

At this layer, each host interface is identified by a media access control (MAC) address. For example, each NIC port in the computers and each radio in the mobile devices has a unique MAC address.


Network Layer Functions

At layer 3, the Network layer, the routing part of the SOHO router makes forwarding decisions between the local private network and public Internet. These zones are distinguished by Internet Protocol (IP) addresses. The local network uses a private IP address range, such as 192 . 168 . 1 . 0 / 24. The SOHO router itself is identified by an address in this range, such as 192 . 168 . 1 . 1 or 192 . 168 . 1 . 254.




Network layer private and public segments. (Images © 123RF.com.)

The router runs a Dynamic Host Configuration Protocol (DHCP) server to allocate a unique address to each host that connects to it over either an Ethernet port or via the wireless access point. The addresses assigned to clients use the same first three octets as the router's address: 192 . 168 . 1. The last octet can be any value from 1 to 254, excluding whichever value is used by the router.



Configuring the LAN addresses using DHCP on a wireless router.
(Screenshot courtesy of TP-Link Technologies Co., Ltd.)

The SOHO router's WAN interface is allocated a public IP address, for example 203.0.113.1, by the Internet service provider. When a host on the local network tries to access any valid IP address outside the 192.168.1.0/24 range, the router forwards that packet over its WAN interface and directs any replies back to the host on the LAN.



Configuring the WAN (Internet) interface on a wireless router. These parameters are supplied by the ISP. Many ISP services use DHCP to allocate a dynamic WAN address, but some offer static addressing. (Screenshot courtesy of TP-Link Technologies Co., Ltd.)

Transport and Application Layer and Security Functions

There is no separate OSI model layer for security. Instead, security issues can arise, and solutions are needed at every layer. Network security is essentially a matter of allowing or preventing devices, users, and services (applications) from using the network. The WAN interface is the network perimeter. The SOHO router can apply filtering rules to traffic sent between the public and private zones, implementing a firewall. The firewall can be configured to block traffic based on source or destination IP addresses and also on the type of application.

At layer 4, each application is identified by a port number, such as 80 for Hypertext Transfer Protocol (HTTP) web traffic or 25 for Simple Mail Transfer Protocol (SMTP) email traffic.

The firewall in the router can be configured with rules specifying behavior for each port. For example, computers on the network might use the Server Message Block (SMB) protocol to share files. It would not be appropriate for hosts on the Internet to be able to access these shared files, so the SMB port would be blocked by default on the WAN interface but allowed on the LAN and WLAN interfaces.

Any host can connect to the RJ45 ports on the router and join the network. The wireless network is usually protected by an encryption system that requires each station to be configured with a passphrase-based key to join the network.

Access to the router's management interface and its configuration settings is protected by an administrative account passphrase. As the router is connected to the Internet, it is critical to configure a strong passphrase.

The screenshot shows the TP-Link Archer VR900 web interface. The top navigation bar includes 'Quick Setup', 'Basic', and 'Advanced' tabs, with 'Advanced' selected. On the left, a sidebar lists various settings: NAT Forwarding, USB Settings, Parental Controls, Bandwidth Control, Security, System Tools (highlighted), Time Settings, Diagnostics, Firmware Upgrade, Backup & Restore, Administration, and System Log. The main content area is divided into three sections: 'Account Management' with fields for Old Password, New Password (with a strength indicator: Low, Middle, High), and Confirm New Password (6arP2UzYpH5QSRE); 'Local Management' with fields for Port (80) and IP/MAC Address; and 'Remote Management' with a checkbox for 'Enable' (unchecked), Port (80), and IP/MAC Address. Each section has a 'Save' button.

*Configuring a management interface on a wireless router.
(Screenshot courtesy of TP-Link Technologies Co., Ltd.)*

The Internet

The WAN interface of the router connects the SOHO network to the Internet.

The Public Switched Telephone Network

Most SOHO subscriber Internet access is facilitated via the **public switched telephone network (PSTN)**. The SOHO router is described as customer premises equipment (CPE). More widely, this is any termination and routing equipment placed at the customer site. Some of this equipment may be owned or leased from the telecommunications company (or telco); some may be owned by the customer.

The CPE is connected via its modem and WAN port to the local loop. This is cabling from the customer premises to the local exchange. The point at which the telco's cabling enters the customer premises is referred to as the demarcation point (often shortened to demarc).

Internet Service Providers

The major infrastructure of the Internet consists of high-bandwidth trunks connecting Internet eXchange Points (IXPs). Within an IXP datacenter, ISPs establish links between their networks, using transit and peering arrangements to carry traffic to and from parts of the Internet they do not physically own. There is a tiered hierarchy of ISPs that reflects to what extent they depend on transit arrangements with other ISPs.

Internet Standards

Although no single organization owns the Internet or its technologies, several organizations are responsible for the development of the Internet and on the agreement of common standards and protocols.

- **Internet Assigned Numbers Authority (IANA) (iana.org)**—Manages allocation of IP addresses and maintenance of the top-level domain space. IANA is currently run by Internet Corporation for Assigned Names and Numbers (ICANN). IANA allocates addresses to regional registries that then allocate them to local registries or ISPs.
- **Internet Engineering Task Force (IETF) (ietf.org)**—Focuses on solutions to Internet problems and the adoption of new standards, published as Requests for Comments (RFCs). Some RFCs describe network services or protocols and their implementation, while others summarize policies. An older RFC is never updated. If changes are required, a new RFC is published with a new number. Not all RFCs describe standards. Some are designated informational, while others are experimental. The official repository for RFCs is at rfc-editor.org.



References to RFCs in this course are for your information should you want to read more. You do not need to learn them for the certification exam.



The OSI model has a stricter definition of the Session, Presentation, and Application layers than is typical of actual protocols used on networks. The Internet model (tools.ietf.org/html/rfc1122) uses a simpler four-layer hierarchy, with a Link layer representing OSI layers 1 and 2, layer 3 referred to as the Internet layer, a Transport layer mapping approximately to layers 4 and 5, and an Application layer corresponding to layers 6 and 7.

Binary and Hexadecimal Notation

To interpret network addresses, you must understand the concept of base numbering systems. To start with the familiar; decimal numbering is also referred to as base 10. Base 10 means that each digit can have one of ten possible values (0 through 9). A digit positioned to the left of another has 10 times the value of the digit to the right. For example, the number 255 can be written out as follows:

$$(2 \times 10 \times 10) + (5 \times 10) + 5$$

Binary is base 2, so a digit in any given position can only have one of two values (0 or 1), and each place position is the next power of 2. The binary value 11111111 can be converted to the decimal value 255 by the following sum:

$$(1 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2) + (1 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2) + (1 \times 2 \times 2 \times 2 \times 2 \times 2) + (1 \times 2 \times 2 \times 2 \times 2) + (1 \times 2 \times 2 \times 2) + (1 \times 2 \times 2) + (1 \times 2) + 1$$

As you can see, it takes eight binary digits to represent a decimal value up to 255. An 8-bit value is called a byte or an octet. While computers process everything in binary, the values make for very long strings if they have to be written out or entered into configuration dialogs. An IPv4 address can be expressed as decimal octets. The four decimal numbers in the SOHO router's WAN IP address 203.0.113.1 are

Hexadecimal notation (or hex) is a convenient way of referring to the long sequences of bytes used in some other types of network addresses, such as hardware MAC addresses. Hex is base 16 with the possible values of each digit represented by the numerals 0 through 9 and the characters A, B, C, D, E, and F.

Use the following table to help to convert between decimal, binary, and hexadecimal values.

Decimal	Hexadecimal	Binary
0	0	0000
1	1	0001
2	2	0010
3	3	0011
4	4	0100
5	5	0101
6	6	0110
7	7	0111
8	8	1000
9	9	1001
10	A	1010
11	B	1011
12	C	1100
13	D	1101
14	E	1110
15	F	1111

As you can see from the table, every hex digit lines up neatly with four binary digits (a nibble). Each byte or octet can be expressed as two hex digits. For example, the decimal value 255 is FF in hex. This would sometimes be written as 0xFF for clarity.

Lesson 1.4

Troubleshooting Methodology



Exam Objectives Covered

5.1 Explain the troubleshooting methodology.

Network problems can arise from a variety of sources outside your control. As a network professional, your users, your managers, and your colleagues will all look to you to identify and resolve those problems efficiently. To do that, you will need a strong fundamental understanding of the tools and processes involved in troubleshooting a network. Being able to resolve problems in these areas is a crucial skill for keeping your network running smoothly.

Troubleshooting requires a best practice approach to both problem-solving and customer/client communication. A troubleshooting model provides you with proven processes on which to base your techniques and approach.

As you study this lesson, answer the following questions:

- What are the advantages of applying a structured methodology to troubleshooting?
- What is the sequence of steps in CompTIA's troubleshooting methodology?
- What activities are performed at each step?

Network Troubleshooting Methodology

When you encounter a network problem, you must try to get it resolved as quickly as you reasonably can. However, you must also take enough time to determine what has caused the problem so that you can avoid a recurrence.

You should make sure you familiarize yourself with the order of the steps in the CompTIA Network+ **troubleshooting methodology**. These steps are explained in more detail in the following topics.

1. Identify the problem:
 - Gather information.
 - Question users.
 - Identify symptoms.
 - Determine if anything has changed.
 - Duplicate the problem, if possible.
 - Approach multiple problems individually.

2. Establish a theory of probable cause:
 - Question the obvious.
 - Consider multiple approaches:
 - Top-to-bottom/bottom-to-top OSI model.
 - Divide and conquer.
3. Test the theory to determine cause:
 - If theory is confirmed, determine next steps to resolve problem.
 - If theory is not confirmed, establish new theory or escalate.
4. Establish a plan of action to resolve the problem and identify potential effects.
5. Implement the solution or escalate as necessary.
6. Verify full system functionality and implement preventive measures if applicable.
7. Document findings, actions, outcomes, and lessons learned throughout the process.

Identify the Problem

The first step in the troubleshooting process is to identify the problem. There are several techniques and approaches that can assist with this step.

Troubleshooting is usually managed by a ticket system. A problem is reported to a help desk, and that report, with the user's initial description of the problem, becomes the basis for the ticket. Tickets could also be generated automatically by monitoring and alerting systems.

Gather Information

To start troubleshooting a ticket, identify the location and scope of the problem. Scope means the area or number of systems affected. This is helpful in two ways. First, a problem that is small in scope might not be a priority in relation to other incidents. Secondly, establishing the scope can help to identify the source of the problem.

If the description in the ticket is unclear or incomplete, to assist with identifying the precise nature of a problem, consider what indirect sources of information there may be:

- Check the system documentation, such as installation or maintenance logs, for useful information.
- Check recent job logs or consult any other technicians who might have worked on the system recently or might be working on some related issue.
- Use vendor support sites (knowledge bases) and



Information gathering is the first step in troubleshooting. (Image by rawpixel © 123RF.com.)

Question Users

You will often need to contact users to establish more facts about the problem. The basis of getting troubleshooting information from users is asking good questions. Questions are commonly divided into two types:

- Open questions invite someone to explain in their own words. Examples are: “What is the problem?” or “What happens when you try to switch the computer on?” Open questions are good to start with, as they help to avoid making your own assumptions about what is wrong, and they encourage the user to give you all the information they can.
- Closed questions invite a Yes/No answer or a fixed response. Examples include: “Can you see any text on the screen?” or “What does the error message say?” Closed questions can be used to drill down into the nature of the problem and guide a user toward giving you information that is useful.

Identify Problem Symptoms

If you cannot identify the problem from user responses, there are additional techniques that you can use to diagnose the issue.

Identify Symptoms and Duplicate the Problem

Symptoms are facts and clues in the affected system that can be correlated with known causes and issues. To identify symptoms, complete the following tests:

- Make a physical inspection; look for something out of the ordinary.
- Check system logs or diagnostic software for information.
- Duplicate the problem on the user’s system or a test system. You will need to try to follow the same steps as the user. Issues that are transitory or difficult to reproduce are often the hardest to troubleshoot.

Determine If Anything Has Changed

There are two key questions to ask when trying to identify the cause of a problem:

- Did it ever work? Hopefully, your users will answer the question truthfully, because the correct answer is important—two different approaches are required. If the system worked before 9:00 a.m., you must ask what happened at 9:00 a.m. If the system never worked, then you are not looking for something that stopped working, but for something which was never working in the first place.
- What has changed since it was last working? The change that caused the problem may not be obvious. Maybe the window cleaners were in the building, and one of them tripped over a cable and now the user can't log in. Maybe someone has moved the user's workstation from one end of their desk to another and plugged the cable into a different port. Check for documented changes using the system inventory, but if this does not reveal anything, look for undocumented changes in the local area of the incident.

Approach Multiple Problems Individually

When you start to investigate symptoms, you might discover symptoms of more than one problem. Perhaps a user has reported that a machine has lost Internet connectivity, and you discover that it has also not been receiving maintenance updates. The issues could be related, or one might be incidental to the other.

If the problems do not seem to be related, treat each issue as a separate case. If they seem to be related, check for outstanding support or maintenance tickets that might indicate existing problems.

It may also be the case that a user reports two different problems at the same time, often preceded by "While you're on the line . . ." sort of statements. Treat each problem as a separate case. In most cases, you should advise the user to initiate a separate support ticket.

Establish a Theory of Probable Cause

If you obtain accurate answers to your initial questions, you will have determined the location, scope, and severity of the problem, and whether to look for a recent change or an oversight in configuration.


You diagnose a problem by identifying the symptoms. From knowing what causes such symptoms, you can test each possible cause until you find the right one. Sometimes symptoms derive from more than one cause; while this type of problem is rarer, it is much harder to troubleshoot.

A network system comprises many components. Fault finding needs to identify which component is causing the issue. For difficult problems, be prepared to consider multiple approaches. If one approach does not identify the problem, use a different one. For example, you could consider two different styles of approaching troubleshooting:

- Question the obvious. Step through what should happen and identify the point at which there is a failure or error. This approach can quickly identify obvious oversights, such as a network cable not being plugged in.
- Methodically prove the functionality of each component in sequence. This approach is more time consuming but may be necessary for a difficult problem.

Top-to-Bottom/Bottom-to-Top OSI Model Approach

Methodical validation of network components can be approached by testing at each layer of the OSI model in sequence. There are many components which go to make up a network.



Troubleshooting top-to-bottom or bottom-to-top using the OSI model.

Some, or several, of these components may be at fault when a problem is reported to you. It is important that you tackle the problem logically and methodically. Unless a problem is trivial, break the troubleshooting process into compartments or categories, using the OSI model as a guide. Start from either the top or bottom and only move up or down when you have discounted a layer as the source of the problem. For example, when troubleshooting a client workstation, you might work as follows:

1. Decide whether the problem is hardware or software related (hardware).
2. Decide which hardware subsystem is affected (NIC or cable).
3. Decide whether the problem is in the NIC adapter or connectors and cabling (cabling).
4. Test your theory (replace the cable with a known good one).

When you have drilled down like this, the problem should become obvious. Of course, you could have made the wrong choice at any point, so you must be prepared to go back and follow a different path.



If you are really unlucky, two (or more) components may be faulty. Another difficulty lies in assessing whether a component itself is faulty or if it is not working because a related component is broken.

Divide and Conquer Approach

In a divide and conquer approach, rather than starting at the top or bottom, you start with the layer most likely to be causing the problem and then work either down or up depending on what your tests reveal. For example, if you start diagnosis at layer 3 and cannot identify a problem, you would then test at layer 4. Conversely, if you discovered a problem at layer 3, you would first test layer 2. If there is no problem at layer 2, you can return to layer 3 and work from there up.

Test the Theory to Determine the Cause

By questioning the obvious or by using one or more methodical diagnostic approaches, hopefully you will have gathered enough data to come to an initial theory about the probable cause. Remember that you might be wrong! Without jumping to conclusions, set out to prove or disprove your suspicions by using your troubleshooting skills and toolkit.

If you cannot prove the cause of the problem, you will either need to develop and test a new theory or decide to escalate the problem. **Escalation** means referring the problem to a senior technician, manager, or third party. You may need to escalate a problem for any of these reasons:

- The problem is beyond your knowledge or ability to troubleshoot.
- The problem falls under a system warranty and would be better dealt with by the supplier.
- The scope of the problem is very large and/or the solution requires some major reconfiguration of the network.
- A customer becomes difficult or abusive or demands help on an unsupported item.

Some of the alternatives for escalation include the following:

- Senior staff, knowledge experts, subject matter experts, technical staff, developers, programmers, and administrators within your company.
- Suppliers and manufacturers.
- Other support contractors/consultants.

When you escalate a problem, you should have established the basic facts, such as the scope of the problem and its likely cause, and be able to communicate these clearly to the person to whom you are referring the incident.

If you can prove the cause of the problem, you can start to determine the next steps to resolve the problem.

Establish a Plan of Action

Assuming you choose not to escalate the issue, the next step in the troubleshooting process is to establish an action plan. An action plan sets out the steps you will take to solve the problem. There are typically three solutions to any problem:

- **Repair**—You need to determine whether the cost of repair/time taken to reconfigure something makes this the best option.
- **Replace**—Often, this is more expensive and may be time consuming if a part is not available. There may also be an opportunity to upgrade the device or software.



A basic technique when you are troubleshooting a cable, connector, or device is to have a known good duplicate on hand (that is, another copy of the same cable or device that you know works) and to test by substitution.

- **Accept**—Not all problems are critical. If neither repair nor replace is cost-effective, it may be best either to find a workaround or to document the issue and move on.

When you consider solutions, you must assess the cost and time required. Another consideration is potential effects on the rest of the system. A typical example is applying a software patch, which might fix a given problem but cause other programs not to work. Up-to-date configuration management documentation and standard operating procedures should help you to understand how different systems are interconnected and cause you to seek the proper authorization for your plan.

Implement the Solution

The solution to a problem might just involve resetting a system to its baseline configuration. Perhaps a user installed some unauthorized software, disabled a necessary service, or unplugged a cable. If you are reverting to a known good configuration, you may be able to implement the solution directly. If the solution requires a change to the system or the network environment, you are likely to have to follow a change management plan.

If you do not have authorization to implement a solution, you will need to escalate the problem to more senior personnel. If applying the solution is disruptive to the wider network, you also need to consider the most appropriate time to schedule the reconfiguration work and plan how to notify other network users. When you change a system as part of implementing a solution, make sure you have a data and configuration backup before proceeding and test after each change. If the change does not fix the problem, reverse it and then try something else. If you make a series of changes without recording what you have done, you could turn a minor problem into a major one.



Virtualization and the cloud provide the means to trial changes before updating the production environment. They allow the rapid creation of sandbox environments that simulate the production one.

Verify the Solution

When you apply a solution, validate that it fixes the reported problem and that the system as a whole continues to function normally. In other words, identify the results and effects of the solution. Ensure that you were right and that the problem is resolved. Can the user now log in properly? Is there any way you can induce the problem again?

Before you can consider a problem closed, you should be satisfied in your own mind that you have resolved it, and you should get the customer's acceptance that it has been fixed. Restate what the problem was and how it was resolved, then confirm with the customer that the incident log can be closed.

To fully solve a problem, you should try to eliminate any factors that may cause the problem to recur. For example, if a user plugs their laptop into the wrong network jack, ensure that the jacks are clearly labeled to help users in the future. If a faulty server induces hours of network downtime, consider implementing failover services to minimize the impact of the next incident.

Document Findings, Actions, and Outcomes

Most troubleshooting takes place within the context of a ticket system. This shows who is responsible for any given problem and what its status is. This gives you the opportunity to add a complete description of the problem and its solution, including findings, actions, and outcomes.

This is very useful for future troubleshooting, as problems fitting into the same category can be reviewed to see if the same solution applies. It also helps to analyze IT infrastructure by gathering statistics on what type of problems occur and how frequently. Analyzing support incidents in a lessons learned process can be used to improve network design, adjust standard procedures, and guide investments in appliance and infrastructure upgrades.

The screenshot shows the 'Create a new ticket' modal in the Spiceworks interface. The form is populated with the following information:

- Contact:** James Pengelly
- Related to:** (empty field)
- Summary:** (empty field)
- Description:** (empty field)
- Assigned to:** Unassigned
- CC Users:** (empty field)
- Due Date:** tomorrow, friday, next we
- Due Time:** 5PM, 18:00, etc
- Priority:** Medium
- Category:** (empty dropdown)
- Attachment:** Choose File

A tip at the bottom reads: "TIP: Record the data you need by adding custom attributes." The 'Save' button is highlighted in blue.

*Creating a ticket in the Spiceworks IT Support management tool.
(Screenshot courtesy of Spiceworks.)*

When you complete a problem log, remember that people other than you may come to rely on it. Also, logs may be presented to customers as proof of troubleshooting activity. Write clearly and concisely, checking for spelling and grammar errors.

Module 1

Summary

You should be able to compare and contrast OSI model layers, encapsulation concepts, and the CompTIA troubleshooting methodology and apply them to analyzing and troubleshooting the function of networks and networking components.

Guidelines for Using Basic Network Concepts

Follow these guidelines to make effective use of basic network concepts, the OSI model, and a structured troubleshooting methodology:

- Use characteristics of Physical layer media and devices to plan wiring topologies and identify potential performance issues.
- Use the Data Link layer to plan logical segments to isolate groups of hosts for performance or security reasons.
- At the Network layer, map Data Link segments to logical network IDs and work out rules for how hosts in one network should be permitted or denied access to other networks.
- Evaluate service requirements at the Transport layer to determine which ports a host should expose.
- Use the Session, Presentation, and Application layers to determine performance and security requirements for the services that the network is providing.
- Use the process of identify, theorize, test, plan, implement, verify, and document to structure all troubleshooting activity.
- Prepare for troubleshooting by developing questioning skills and building reference documentation.
- Use top-to-bottom/bottom-to-top/divide-and-conquer with reference to the OSI model to isolate problems to layers.

Module 2

Supporting Cabling and Physical Installations

Module Introduction

At the Physical layer, networks are made from different cabling types and their connectors and transceivers. These establish direct links between nodes in a local segment. At the Data Link layer, nodes in these segments are given a standard means of exchanging data as frames.

As they are closely related, networking products often define standards for both the Physical and Data Link layers. While plenty of products have been used in the past, many cabled networks are now based on the Ethernet standards. Understanding the options and specifications for Ethernet are essential to building and supporting networks of all sizes.

In this module, you will summarize standards for deploying Ethernet over copper and fiber optic media types and identify the tools and techniques required to deploy and troubleshoot Ethernet cabling.

Module Objectives

In this module, you will do the following:

- Summarize Ethernet standards.
- Summarize copper cabling and connector types.
- Summarize fiber optic cabling and connector types.
- Describe physical installation factors for rack-based installations in server rooms and datacenters.
- Deploy and troubleshoot Ethernet cabling.

Lesson 2.1

Ethernet



Exam Objectives Covered

1.5 Compare and contrast transmission media and transceivers.

In this lesson, you will identify the components used in an Ethernet network implementation. A standard provides detailed specifications for Physical layer media and interfaces. The Ethernet standard dominates the wired LAN product market. Large and small networks use Ethernet technology to provide both backbone and end user services. Due to the wide deployment of Ethernet today, you will undoubtedly be required to manage and troubleshoot Ethernet networks.

As you study this lesson, answer the following questions:

- What cable topologies and appliances support the creation of an Ethernet network?
- How do multiple hosts share access to the same media?
- How can I identify what cable speed is specified by a given Ethernet standard?

Network Data Transmission

Network data transfer works by modulating the properties of a transmission medium—electric current, infrared light, or radio waves—to encode a signal. One example of modulation is transitioning between low and high voltage states in an electrical circuit. These voltage pulses can encode symbols, which can be mapped to digital bits—ones and zeros.

Each media type supports a range of possible frequencies. Higher frequencies allow for more data to be transferred per second. The range of frequencies is referred to as the media bandwidth.



The narrow definition of bandwidth is a frequency range measured in cycles per second or Hertz (Hz), but the term is very widely used in data networking to mean the amount of data that can be transferred, measured in multiples of bits per second (bps). Encoding methods mean that, for instance, a signal with 100 MHz frequency bandwidth can transfer much more than 100 Mbps.

Ethernet Standards

Over the years, many protocols, standards, and products have been developed to implement the functions of the Physical and Data Link layers of the OSI model. A standard must define cable and connector specifications and define schemes for modulation and encoding.

The Institute of Electrical and Electronics Engineers (IEEE) **802.3 Ethernet standards** (ieee802.org/3) are very widely used on both LANs and WANs. Ethernet standards provide assurance that network cabling will meet the bandwidth requirements of applications. These Ethernet media specifications are named using a three-part convention, which is often referred to as xBASE-y. This describes the following:

- The speed or bit rate in megabits per second (Mbps) or gigabits per second (Gbps).
- The signal mode (baseband or broadband). All mainstream types of Ethernet use baseband transmissions, so you will only see specifications of the form xBASE-y.
- A designator for the media type.

For example, 10BASE-T denotes an early implementation that works at 10 Mbps (10), uses a baseband signal (BASE), and runs over **twisted pair** copper cabling (-T).


Copper cable is used to transmit electrical signals. The cable between two nodes creates a low voltage electrical circuit between the interfaces on the nodes. There are two main types of copper cable: twisted pair and coaxial (coax). Copper cable suffers from high attenuation, meaning that the signal quickly loses strength over long links. Twisted pair cable is rated to **Category** (or "Cat") standards that define what bandwidth it should support, up to a given distance.

Media Access Control and Collision Domains

Ethernet is a multiple access area network, which means that the available communications capacity is shared between the nodes that are connected to the same media. Media access control (MAC) refers to the methods a network technology uses to determine when nodes can communicate on shared media and to deal with possible problems, such as two devices attempting to communicate simultaneously.

Ethernet uses a contention-based MAC system. Each network node connected to the same media is in the same **collision domain**. When two nodes transmit at the same time, the signals are said to collide, and neither signal can reach its destination. This means that they must be resent, reducing available bandwidth. The collisions become more frequent as more nodes are added, and consequently the effective data rate is reduced.

The Ethernet protocol governing contention and media access is called **Carrier Sense Multiple Access with Collision Detection (CSMA/CD)**. A collision is the state when a signal is present on an interface's transmit and receives lines simultaneously. On detecting a collision, the node broadcasts a jam signal. Each node that was attempting to use the media then waits for a random period (backoff) before attempting to transmit again.



The CSMA/CD media access method. (Images © 123RF.com.)

The collision detection mechanism means that only half-duplex transmission is possible. This means that a node can transmit or receive, but it cannot do both at the same time.

In the 10BASE-T star wiring physical topology, each node is cabled to an Ethernet hub. The hub repeats incoming signals to each connected node. Consequently, every host connected to the same hub is within the same collision domain. However, this 10BASE-T physical topology dates from 1990. You are very unlikely to find it deployed in a modern network.

100BASE-TX Fast Ethernet Standards

The Fast Ethernet standard uses the same CSMA/CD protocol as 10BASE-T but with higher frequency signaling and improved encoding methods, raising the bit rate from 10 Mbps to 100 Mbps. 100BASE-TX refers to Fast Ethernet working over Cat 5 (or better) twisted pair copper cable with a maximum supported link length of 100 meters (328 feet).

100BASE-TX can be implemented with a hub, but the standard was created at a time that switches started to replace hubs as the connection point for end systems. The contention-based access method used by a hub does not scale to large numbers of end systems within the same collision domain. Where a hub works only at the Physical layer, a switch uses information about source and destination addresses carried in layer 2 frames to establish a temporary circuit between two nodes. Unlike a hub, each switch port is a separate collision domain. By eliminating the effect of contention, switches allow for full-duplex transmissions, where a node can transmit and receive simultaneously, and each node can use the full 100 Mbps bandwidth of the cable link to the switch port.

To support compatibility with hosts still equipped with 10 Mbps Ethernet interfaces, Fast Ethernet introduced an autonegotiation protocol to allow a host to choose the highest supported connection parameters (10 or 100 Mbps and half- or full-duplex). 10BASE-T Ethernet specifies that a node should transmit regular electrical pulses when it is not transmitting data to confirm the viability of the link. Fast Ethernet codes a 16-bit data packet into this signal, advertising its service capabilities. This is called a Fast Link Pulse. A node that does not support autonegotiation can be detected by one that does and sent ordinary link integrity test signals, or Normal Link Pulses.

Fast Ethernet would not be deployed on new networks, but you may need to maintain it in legacy installations.

Gigabit Ethernet Standards

Gigabit Ethernet builds on the standards defined for Ethernet and Fast Ethernet to implement rates of 1,000 Mbps (1 Gbps). When installed using Cat 5e or better copper wire, Gigabit Ethernet is specified as 1000BASE-T. Gigabit Ethernet does not support hubs; it is implemented only using switches. The maximum distance of 100 meters (328 feet) applies to cabling between the node and a switch port, or between two switch ports.

Gigabit Ethernet is the mainstream choice for new installations of access networks; that is, cabling to connect client workstations to a local network. The main decision would be whether to use copper or fiber optic cable. Fiber gives better upgrade potential in the future, while copper cable is cheaper to install and far more hosts are installed with network cards that support copper than support fiber.

10 Gigabit Ethernet (10 GbE) multiplies the nominal speed of Gigabit Ethernet by a factor of 10. Because of the higher frequencies required, 10 GbE can only run at reduced distances over unshielded copper cable. Longer runs require higher categories of copper cable with some type of shielding, or the use of fiber optic cable. There are also specifications for 40 Gbps operation.

Specification	Cable	Maximum Distance
10GBASE-T	UTP (Cat 6)	55 meters (180 feet)
	F/UTP (Cat 6A)	100 meters (328 feet)
	S/FTP (Cat 7)	100 meters (328 feet)
40GBASE-T	S/FTP (Cat 8)	30 meters (100 feet)

10/40 GbE Ethernet is not deployed in many access networks, as the cost of 10/40 GbE compatible network adapters and switch transceiver modules is high. It might be used where a company's business requires very high-bandwidth data transfers, such as TV and film production. It is also widely used as backbone cabling, where it supports high-bandwidth links between switches and routers, or between appliances in a datacenter.

Fiber Ethernet Standards

Fiber optic cable uses infrared light signals. The light signals are also not susceptible to interference or noise from other sources and less effected by attenuation. Consequently, fiber optic cable supports higher bandwidth over longer links than copper cable. Fiber optic cabling is divided into single mode (SMF) and multimode (MMF) types, and MMF is categorized by optical mode designations (OM1, OM2, OM3, and OM4).

Ethernet standards over fiber set out the use of different types of cable for 100 Mbps, 1 Gbps, 10 Gbps, and 40/100 Gbps operation. There are variants for long wavelength optics, required for long-distance transmission, and short wavelength optics. Some of the main standards for speeds up to 10 Gbps are listed in the table.

Specification	Optics	Cable	Maximum Distance	Connectors
100BASE-FX	1300 nm	MMF (OM1)	4 km (2.48 miles)	ST, SC, MT-RJ
100BASE-SX	850 nm	MMF (OM1) MMF (OM2)	300 meters (984 feet)	ST, SC, LC
1000BASE-SX	850 nm	MMF (OM1) MMF (OM2) MMF (OM3)	275 meters (902 feet) 550 meters (1804 feet)	ST, SC, LC, MT-RJ
1000BASE-LX	1,300 nm 1,310 nm	MMF (OM1/OM2/OM3) SMF (OS1/OS2)	550 meters (1,804 feet) 5 km (3.1 miles)	SC, LC
10GBASE-SR	850 nm	MMF (OM1) MMF (OM2) MMF (OM3) MMF (OM4)	33 meters (108 feet) 82 meters (269 feet) 300 meters (984 feet) 400 meters (1,312 feet)	SC, LC
10GBASE-LR	1,310 nm	SMF (OS1/OS2)	10 km (6.2 miles)	SC, LC

Fiber is often used for backbone cabling in office networks and for workstations with high-bandwidth requirements, such as video editing. The principal applications of 10 GbE (and better) are the following:

- Increasing bandwidth for server interconnections and network backbones, especially in datacenters and for storage area networks (SANs).
- Replacing existing switched public data networks based on proprietary technologies with simpler Ethernet switches (Metro Ethernet).

Lesson 2.2

Copper Cables and Connectors



Exam Objectives Covered

1.5 Compare and contrast transmission media and transceivers.


Copper wire twisted pair cabling is the most popular choice for access networks in offices. You are likely to work with this network media daily as part of your duties as a network professional. Understanding the characteristics of twisted pair will enable you to properly install and service your networks.

As you study this lesson, answer the following questions:

- Why are wires twisted together in twisted pair cables?
- What is the difference between STP cabling and UTP cabling?
- What speeds and distances are supported by different cable categories?
- How can you tell the difference between RJ11 and RJ45 connectors?
- What are the uses of coax and twinax copper cabling?

Unshielded Twisted Pair Cable

Twisted pair is a type of copper cable that has been extensively used for telephone systems and data networks. One pair of insulated wires twisted together forms a balanced pair. The pair carries the same signal but with different polarity; one wire is positive, and the other is negative. This allows the receiver to distinguish the signal from any noise more strongly. The cable is completed with an insulating outer jacket.



*Twisted pair cable—Each color-coded pair is twisted at a different rate to reduce interference.
(Image by Thuansak Srilao © 123RF.com.)*

The pairs are twisted at different rates to reduce external interference and crosstalk. Crosstalk is a phenomenon whereby one pair causes interference in another as a result of their proximity.

Twisted pair can use either solid or stranded conductor wires. Solid cabling uses a single thick wire per conductor and is used for cables that run behind walls or through ducts. Stranded cabling uses thin filament wires wrapped around one another and is used to make flexible patch cords for connecting computers to wall ports and switch ports to patch panel ports. Copper wire thickness is measured using American Wire Gauge (AWG). Increasing AWG numbers represent thinner wire. Solid cable uses thicker 22 to 24 AWG, while the stranded cable used for patch cords is often 26 AWG. The attenuation of stranded wire is higher than solid wire, so it should not be used for cables over 5 m in length.

Most twisted pair cable used in office networks is **unshielded twisted pair (UTP)**. Modern buildings are often flood wired using UTP cabling. This involves cables being laid to every location in the building that may need to support a telephone or computer.

Shielded and Screened Twisted Pair Cable

Shielded cable is less susceptible to interference and crosstalk. This type of cable is required for some Ethernet standards and may also be a requirement in environments with high levels of interference, such as cabling that is run near motors, generators, or fluorescent lighting.

Shielded cable can be referred to generically as **shielded twisted pair (STP)**, but there are actually several types of shielding:

- Screened cable has one thin outer foil shield around all pairs. Screened cable is usually designated as screened twisted pair (ScTP) or foiled/unshielded twisted pair (F/UTP), or sometimes just foiled twisted pair (FTP).
- Fully shielded cabling has a braided outer screen and foil-shielded pairs and is referred to as shielded/foiled twisted pair (S/FTP). There are also variants with a foil outer shield (F/FTP).
- U/FTP cable has foil-shielded pairs but no outer shield.

Legacy STP cable could be complex to install, as it required bonding each element to ground manually, but modern F/UTP and S/FTP solutions with appropriate cable, connectors, and patch panels reduce this complexity by incorporating bonding within the design of each element.

Cat Cable Standards

The American National Standards Institute (ANSI) and the Telecommunications Industry Association (TIA)/Electronic Industries Alliance (EIA) have created categories of cable standards for twisted pair to simplify selection of a suitable quality cable. These categories, along with other aspects of telecommunications wiring best practices, are defined in the ANSI/TIA/EIA 568 Commercial Building Telecommunications Cabling Standards ([tiaonline.org/standard/tia-568](https://www.tiaonline.org/standard/tia-568)). Similar standards are also maintained by the ISO (ISO/IEC 11801), which refers to categories of components and classes of permanent links (incorporating both cable and termination).


Cat/Class	Cable Type	Ethernet Standard	Max. Distance	Frequency	Connector
5e (Class D)	UTP or F/UTP	1000BASE-T	100 m (328 ft)	100 MHz	RJ45
6 (Class E)	UTP, F/UTP, or U/FTP	1000BASE-T	100 m (328 ft)	250 MHz	RJ45
		10GBASE-T	55 m (180 ft)		
6A (Class Ea)	UTP, F/UTP, U/FTP, or S/FTP	10GBASE-T	100 m (328 ft)	500 MHz	RJ45
7 (Class F)	S/FTP or F/FTP	10GBASE-T	100 m (328 ft)	600 MHz	GG45/ TERA
8/8.1 (Class I)	U/FTP or F/UTP	40GBASE-T	30 m (100 ft)	2,000 MHz	RJ45
8.2 (Class II)	F/FTP or S/FTP	40GBASE-T	30 m (100 ft)	2,000 MHz	GG45/ TERA

Twisted Pair Connector Types

Twisted pair copper cabling uses **Registered Jack (RJ)** connectors for the physical interface. There are many different types of RJ connector, identified by numbers (and sometimes letters). Some are physically different, while others are identical but wired differently for different applications. The most widely used connectors are RJ45 and RJ11.

RJ45 Connectors

RJ45 connectors are used with 4-pair copper cables. The connectors are also referred to as 8P8C, standing for 8-position/8-contact. This means that all eight “potential” wire positions are supplied with contacts, so that they can all carry signals if needed. RJ45 is used for Ethernet twisted pair cabling.



RJ45 jack and plug. (Image © 123RF.com.)




RJ45 plugs have a plastic retaining clip. This is normally protected by a rubber boot. This type of cable construction is also referred to as snagless.



There are also GG45 and TERA connectors, associated with ISO Class F and Class II cabling. GG45 has a similar form factor to RJ45 but has four conductors in the corners. TERA connectors have a completely different form factor.

RJ11 Connectors

The smaller RJ11 connector is used with 2-pair copper cable. An RJ11 connector can support six positions, but only the center two contacts are wired (6P2C). In a telephone system, this pair carries the dial tone and voice circuit. These are also called the Tip and Ring wires after the way older phone plugs were wired. The other pair is usually unused but can be deployed for a secondary circuit. RJ11 connectors are used for telephone systems and to connect analog data modems to a phone jack.



RJ11 jack and 6P2C plug. (Image © 123RF.com.)

Other six position connectors are the same physical form factor but wired to use more pairs. RJ14 is 6P4C, and RJ25 is 6P6C.

Plenum and Riser-Rated Cable

When choosing a copper cable type, basic considerations include the category, unshielded versus shielded, and the appropriate use of solid versus stranded cable. Another consideration is how cable installation must conform to any national or local building regulations.

A **plenum** space is a void in a building designed to carry heating, ventilation, and air conditioning (HVAC) systems. Plenum space is typically a false ceiling, though it could also be constructed as a raised floor. As it makes installation simpler, this space has also been used for communications wiring in some building designs. Plenum space is an effective conduit for fire, as there is plenty of airflow and no fire breaks (such as walls or doors). If the plenum space is used for heating, there may also be higher temperatures. Therefore, building regulations require the use of fire-retardant plenum cable in such spaces. Plenum cable must not emit large amounts of smoke when burned, be self-extinguishing, and meet other strict fire safety standards.

General-purpose (non-plenum) cabling uses PVC (polyvinyl chloride) jackets and insulation. Plenum-rated cable uses treated PVC or fluorinated ethylene propylene (FEP). This can make the cable less flexible, but the different materials used have no effect on bandwidth. Communications cable that is plenum rated under the U.S. National Electrical Code (NEC) is marked CMP. General-purpose cables are marked CMG or CM.

Cabling that passes between two floors is referred to as riser. Conduit for riser cabling must be fire stopped. This means that fire cannot spread through the opening created by the conduit. Riser cabling (in conduit or in spaces such as lift shafts) should also conform to the appropriate fire safety standards. These are similar to the requirements for plenum spaces but not quite as strict. Data cable that is riser rated under the NEC is marked CMR.



You can use plenum-rated cables in place of riser-rated cables, but never use riser-rated cables in place of plenum-rated cables. Both of these typically include a rope or filament that helps support their weight when they're installed vertically.


Coaxial and Twinaxial Cable and Connectors

Coaxial (or coax) cable is made of two conductors that share the same axis. The core conductor is made of solid or stranded copper wire and is enclosed by plastic insulation. A wire mesh wrapped around the plastic constitutes the second conductor. This serves both as shielding from interference.



Coax cable. (Image by destinacigdem © 123RF.com.)

Coax cables are categorized using Radio Grade (RG) designations, which represent the thickness of the core conductor and the cable's characteristic impedance. RG6 is 18 AWG cable with 75 ohm impedance typically used as drop cable for Cable Access TV (CATV) and broadband cable modems. Thinner, more flexible RG59 cable is used for audio/video and closed-circuit television (CCTV). For these applications, coax can be terminated using either a **Bayonet Neill-Concelman (BNC) connector** or an **F-type connector**. BNC uses a twist-and-lock mechanism, while F-connectors are secured by screwing them into place.



BNC connector on the left. (Image created by Krzysztof Burghardt and reproduced under the Creative Commons Attribution ShareAlike 3.0 license.) F-type connector on the right. (Image created by Colin and reproduced under the Creative Commons Attribution ShareAlike 3.0 license.)

Twinaxial (or twinax) is similar to coax but contains two inner conductors. Twinax is used for datacenter interconnects working at 10 GbE (unofficially referred to as 10GBASE-CR) and 40 GbE (40GBASE-CR4). The maximum distance is up to about 5 meters for passive cable types and 10 meters for active cable types. Twinax for 10/40 GbE is terminated using **Direct Attach Copper (DAC)** transceivers. These transceivers can be installed as modules in switch, router, and server appliances.



Direct Attach Copper (DAC) twinax cabling with SFP+ termination. (Image created by Labsy and reproduced under the Creative Commons Attribution ShareAlike 4.0 license.)

Lesson 2.3

Wiring Implementation



Exam Objectives Covered

2.4 Explain important factors of physical installations.

The networking industry has developed a standard model for deploying a structured cabling system. The model is adaptable to both small and large networks. In this topic, you will describe how a distribution system can provision network access throughout a building or site.


As you study this lesson, answer the following questions:

- What is the difference between the MDF and an IDF, and how do they relate to patch panels?
- When should you use stranded core twisted pair cable instead of solid core twisted pair?
- What is the difference between the T568A and T568B standards?
- When you use a punch down tool, which way should the blade be facing?

Structured Cabling System


A structured cabling scheme is a standard way of provisioning cabled networking for computers in an office building. The best known is the ANSI/TIA/EIA 568 Commercial Building Telecommunications Wiring Standard. ANSI/TIA/EIA 568 identifies the following subsystems within a structured cabling system:

- **Work Area**—The space where user equipment is located and connected to the network, usually via a patch cable plugged into a wall port.



Modular wall plate. (Image by Nikolai Lebedev © 123RF.com.)

- **Horizontal Cabling**—Connects user work areas to an **intermediate distribution frame (IDF)**. Horizontal cabling is so-called because it typically consists of the cabling for a single floor and so is made up of cables run horizontally through wall ducts or ceiling spaces. When using copper cabling, the IDF must be within 90 m (295 feet) cabling distance of each wall port. If this is not possible, multiple IDFs must be provisioned. Multiple IDFs on the same floor are linked by horizontal cross connects.



Wiring distribution components. (Images © 123RF.com.)

- **Telecommunications Room**—Room or closet that houses an intermediate distribution frame and networking equipment, such as switches. Essentially, this is a termination point for the horizontal cabling along with a connection to backbone cabling. This wiring closet must be used only for networking equipment (not general storage) and should ideally be secured by a lockable door.
- **Backbone Cabling**—Connects IDFs to a **main distribution frame (MDF)**. Backbone cabling is also referred to as vertical cabling, as it is more likely to run up and down between floors.
- **Entrance Facilities/Demarc**—Special type of telecommunications room marking the point at which external cabling is joined to internal cabling, via the MDF. Entrance facilities are required to join the local exchange carrier's (LEC's) network and for inter-building communications. The **demarcation point** is where the access provider's network terminates and the organization's network begins.



Smaller facilities might not require IDFs. If distance limitations are not exceeded, wall ports can be terminated directly to a single main distribution frame.

T568A and T568B Termination Standards

Twisted pair must be properly terminated. Patch cords are terminated with RJ45 plugs, while structured cabling is terminated to insulation displacement connectors (IDCs) in wall ports and distribution frames. When terminating cable, an organization should use a consistent wiring scheme across all sites.

Each conductor in a 4-pair data cable is color-coded. Each pair is assigned a color (blue, orange, green, or brown). The first conductor in each pair has a predominantly white insulator with strips of the color; the second conductor has an insulator with the solid color. The ANSI/TIA/EIA 568 standard defines two methods for terminating Ethernet cabling: **T568A** and **T568B**. The wiring for both standards is shown in the following figure.



T568A and T568B wiring diagrams. (Images © 123RF.com.)

In T568A, the green pairs are wired to pins 1 and 2, and the orange pairs are wired to pins 3 and 6. In T568B, these pairs swap places, so orange is terminated to pins 1 and 2 and green to 3 and 6. Organizations should try to avoid using a mixture of the two standards. T568A is mandated by the residential cabling standard (TIA 570), but T568B is probably the more widely deployed of the two.

Pin	Wire Color (T568A)	Wire Color (T568B)	10/100 Mbps	1/10/40 Gbps
1	Green/White	Orange/White	Tx+	BixA+
2	Green	Orange	Tx-	BixA-
3	Orange/White	Green/White	Rx+	BixB+
4	Blue	Blue		BixC+
5	Blue/White	Blue/White		BixC-
6	Orange	Green	Rx-	BixB-
7	Brown/White	Brown/White		BixD+
8	Brown	Brown		BixD-




Cat 7 and Cat 8 are so sensitive to noise that the secondary wire in each pair is solid white with no stripe, as the coloring process reduces the effectiveness of the insulation.

Patch Panels


Cable management techniques and tools ensure that cabling is reliable and easy to maintain. Structured copper wiring runs from a wall port in the user's work area to some type of distribution frame in the network closet. At both ends, it is terminated at a punch down block with **insulation-displacement connection (IDC)** terminals. An IDC contains contacts that cut the insulation from a wire and hold it in place. This design allows large numbers of cables to be terminated within a small space.

In data networks, numerous moves, adds, and changes (MACs) would require re-terminating the wiring. To simplify MACs, a distribution frame is normally implemented as a **patch panel**. This has punch down blocks on one side and pre-terminated RJ45 modular ports on the other. This allows incoming and outgoing connections to be reconfigured by changing the patch cable connections, which is much simpler than re-terminating punch down blocks.



IDCs at the rear of a patch panel. (Image by plus69 © 123RF.com.)

The structured cabling running from the work area or forming a backbone is terminated at the back of the patch panel on the IDCs, using either T568A or T568B wiring order. An RJ45 patch cord is used to connect the port to another network port, typically a switch port housed in the same rack. This greatly simplifies wiring connections and is the most commonly installed type of wiring distribution where connections need to be changed often.



Patch panel with pre-wired RJ45 ports. (Image by Svetlana Kurochkina © 123RF.com.)

Structured Cable Installation

Installing structured cable from a bulk spool is referred to as pulling cable because the cable must be pulled, carefully, from the telecommunications closet to the work area. Cable is normally routed through conduits or wall spaces, avoiding excessive bends and proximity to electrical power cables and fittings, such as fluorescent lights, as these could cause interference. The main fixed cable run can be up to 90 m (295 feet). Stranded-wire patch cords can be up to 5 m each (16 feet) and no more than 10 m (33 feet) in overall length. This is because the attenuation of stranded cable is higher than solid cable.

Starting at the patch panel, label the end of the cable with the appropriate jack ID, then run it through to the work area. This is also referred to as a drop, as in most cases you will be dropping the cable from the ceiling space through a wall cavity. If several cables are going to roughly the same place, you can bundle them and pull them together. Leave enough slack at both ends (a service loop) to make the connection and to accommodate future reconnections or changes, cut the cable, and label the other end with the appropriate ID. Electrician's scissors (snips) are designed for cutting copper wire and stripping insulation and cable jackets. Alternatively, there are dedicated **cable stripper** tools that have replaceable blades for different data cable types. Cable-cutting blades should be rounded to preserve the wire geometry. Stripping tools should have the correct diameter to score a cable jacket without damaging the insulation wires.

Termination Tools and Techniques


To terminate a cable, untwist the ends of the wire pairs and place them into the punch down block in the correct order for the wiring configuration (T568A or T568B) you want to use.



You must not untwist the wires too much. Cat 6 is demanding in this respect and requires no more than 0.375" (1 cm) of untwisting.

Fixed cable is terminated using a **punch down tool**. This tool fixes conductors into an IDC. There are different IDC formats (66, 110, BIX, and Krone), and these require different blades. Many punch down tools have replaceable blades, though. Blades are double sided; one side pushes the wire into the terminal while the other side cuts the excess. Make sure the blade marked "cut" is oriented correctly to cut the excess wire.

Alternatively, a block tool terminates a group of connectors in one action. For a 110 format panel, a four position block is suitable for terminating 4-pair data cabling.



*Adding RJ45 terminals to a network cable using a punch down tool.
(Image by dero2084 © 123RF.com.)*

A patch cord is created using a **cable crimper**. This tool fixes a plug to a cable. The tools are specific to the type of connector and cable, though some may have modular dies to support a range of RJ-type plugs.

For shielded and screened cable, termination must be made to shielded IDCs or modular plugs. On an IDC, a metal clip placed over the exposed foil or braided shield bonds the cable to the housing. A shielded modular plug has a metal housing and is not terminated using a standard crimper. There are several different designs, but all follow the principle of connecting the cable shield to a bonding strip.

Lesson 2.4

Fiber Optic Cables and Connectors



Exam Objectives Covered

- 1.5 Compare and contrast transmission media and transceivers.*
- 2.4 Explain important factors of physical installations.*

Fiber optic media can support higher bandwidths over longer distances than copper wire. These advantages make it a popular choice for long-distance telecommunications networks and for reliable, high-speed networking within datacenters. Understanding the characteristics of fiber optic media will help you to support existing installations and upgrades.

As you study this lesson, answer the following questions:

- What advantages do fiber optic cables offer over copper cables?
- What is the difference between single mode and multimode cables?
- What are the different fiber connector types, and what are their typical uses?
- What are duplex, parallel, and multiplexed transmission types?

Fiber Optic Cable Considerations

The electrical signals carried over copper wire are subject to interference and attenuation. Fiber optic signaling uses pulses of infrared light, which are not susceptible to interference, cannot easily be intercepted, and suffer less from attenuation. Consequently, fiber optic cabling supports higher bandwidth over longer cable runs. Fiber optic cabling can be many kilometers long.

A single optical fiber is constructed from three elements:

- Core provides the transmission path, or waveguide, for the light signals.
- Cladding reflects signals back into the waveguide as efficiently as possible. The core and cladding can be made from glass or plastic. The cladding is applied as a thin layer surrounding the core. While made of the same material, the cladding has a different refractive index than the core. The effect of this is to create a boundary that causes the light to bounce back into the core, facilitating the process of total internal reflection that guides the light signal through the core.
- Buffer is a protective plastic coating. It may be of a tight or loose configuration, with the loose format using some form of lubricant between the strand and the sheath.

In basic operation modes, each fiber optic strand can only transfer light in a single direction at a time. Therefore, multiple fibers are often bundled within a cable to allow simultaneous transmission and reception of signals or to provide links for multiple applications.

There are many different outer jacket designs and materials suited for different installations (indoor/plenum, outdoor, underground, undersea, and so on). Kevlar (Aramid) strands and sometimes fiberglass rods (strength members) are often used to protect the fibers from excessive bending or kinking when “pulling” the cable to install it. For exposed outdoor applications, a steel shield (armor) may be added to deter rodents from gnawing the cable.

Single Mode Fiber and Multimode Fiber

Fiber optic cables are specified using the mode, composition (glass/plastic), and core/cladding size; for example, 8.3 micron core/125 micron cladding single mode glass or 62.5 micron core/125 micron cladding multimode plastic. Fiber optic cables fall into two broad categories: single mode and multimode.

- **Single Mode Fiber (SMF)** has a small core (8 to 10 microns) and a long wavelength. It uses a laser to generate a near infrared (1,310 nm or 1,550 nm) light signal. Single mode cables support data rates up to 100 Gbps and cable runs of many kilometers, depending on the quality of the cable and optics. There are two grades of SMF cable; OS1 is designed for indoor use, while OS2 is for outdoor deployment.
- **Multimode Fiber (MMF)** has a larger core (62.5 or 50 microns) and shorter wavelength light (850 nm or 1,300 nm) transmitted in multiple waves of varying length. MMF uses less expensive optics and consequently is less expensive to deploy than SMF. However, it does not support such high signaling speeds or long distances as single mode and so is more suitable for LANs than WANs.



Optical transceivers for SMF are now only slightly more expensive than ones for MMF. Consequently, SMF is often used for short-range applications in datacenters, as well as for long-distance links. SMF still comes at a slight price premium, but it provides better support for 40 Gbps and 100 Gbps Ethernet standards.

MMF is graded by **optical multimode (OM)** categories, defined in the ISO/IEC 11801 standard:

- **OM1/OM2**—62.5-micron cable is OM1, while early 50-micron cable is OM2. OM1 and OM2 are mainly rated for applications up to 1 Gbps and use LED transmitters.
- **OM3/OM4**—These are also 50-micron cable, but manufactured differently, designed for use with 850 nm vertical-cavity surface-emitting lasers (VCSEL), also referred to as laser optimized MMF (LOMMF). A VCSEL is not as powerful as a laser type used for SMF, but it supports higher modulation (transmitting light pulses rapidly) than LED-based optics.

Fiber Optic Connector Types

Fiber optic connectors are available in many different form factors. Some types are more popular for multimode and some for single mode.

Straight Tip

Straight Tip (ST) is an early bayonet-style connector that uses a push-and-twist locking mechanism. ST was used mostly for multimode networks, but it is not widely used for Ethernet installations anymore.




Two ST connectors. (Image by Aleh Datskevich © 123RF.com.)

Subscriber Connector

The **Subscriber Connector (SC)** is a push/pull design, allowing for simple insertion and removal. It can be used for single- or multimode. It is commonly used for Gigabit Ethernet.

Local Connector

The **Local Connector (LC)** (also referred to as Lucent Connector) is a small-form-factor connector with a tabbed push/pull design. LC is similar to SC, but the smaller size allows for higher port density. LC is a widely adopted form factor for Gigabit Ethernet and 10/40 GbE.



*Patch cord with duplex SC format connectors (left) and LC connectors (right).
(Image by YANAWUT SUNTORNIJ © 123RF.com.)*


Fiber Optic Cable Installation

Fiber optic can be installed in the same topology as copper cable using distribution frames and switches. Long-distance cables are typically laid as trunks or rings with repeaters or amplifiers between cable segments to strengthen the signal.

With duplex fiber, strands are installed in pairs, with one strand for transmit (Tx) and one strand for receive (Rx).

Fiber Optic Patch Cords

Patch cables for fiber optic can come with the same connector on each end (LC-LC, for instance) or a mix of connectors (LC-SC, for instance). Duplex patch cords must maintain the correct polarity, so that the Tx port on the transmitter is linked to the Rx port on the receiver and vice versa. The TIA/EIA cabling standard sets out a system of A to B polarity. Each element in the link must perform a crossover, and there must be an odd number of elements, such as two patch cords and a permanent link (three elements).



Fiber patch cord polarity.

Most connectors are keyed to prevent incorrect insertion, but if in doubt, an optical power meter can be used to determine whether an optical signal is being received from a particular fiber.



Transmitted optical signals are visible as bright white spots when viewed through a smartphone camera. This can be used to identify which adapter on an optical interface is transmitting and which fiber patch cord is receiving a signal from the other end of the cable.

Finishing Type

The core of a fiber optic connector is a ceramic or plastic ferrule that holds the glass strand and ensures continuous reception of the light signals. The tip of the ferrule can be finished in several formats. The two most popular are:

- **Ultra Physical Contact (UPC)**—The faces of the connector and fiber tip are polished so that they curve slightly and fit together better.
- **Angled Physical Contact (APC)**—The faces are angled for an even tighter connection. APC cannot be mixed with PC or UPC.


It is important to match the finishing type when you are selecting a connector type. APC finishing is often not supported by the patch panels, transceivers, and switch ports designed for Ethernet.

Also, by convention, cable jackets and connectors use the following color-coding:

Type	Jacket Color	Connector Color
OM1	Orange	Beige
OM2	Orange	Black
OM3/OM4	Aqua	Aqua
SMF PC/UPC	Yellow	Blue
SMF APC	Yellow	Green

Fiber Distribution Panels and Fusion Splicing

A modern build or refurbishment might replace copper wiring with fiber optic cabling. Structured cabling links are installed in a manner similar to copper cabling. However, to avoid the wear and tear damage associated with continually reconnecting fiber optic cables, it's essential not to frequently replace cable runs through conduit. Permanent cables are therefore routed through conduit to wall ports at the client access end, and to a **fiber distribution panel** at the switch end. To complete the connection, fiber patch cables are used to link the wall port to the network interface card (NIC) and the patch panel to the switch port.




Fiber distribution panel. (Image by Aleh Datskevich © 123RF.com.)

Multi-Fiber Push On Connectors

Multi-fiber push-on (MPO) termination allows for low-footprint backbone or trunk cabling. An MPO backbone ribbon cable bundles 12 or more strands terminated to a single compact ferrule. MPO cables are usually prefabricated and not typically field terminated. There are MMF and SMF variants.

MPO is mostly used to aggregate 10 Gbps or 25 Gbps lanes into a 40 Gbps, 100 Gbps, or 400 Gbps parallel optical link. Each lane normally requires two fiber strands (send and receive). A 40 Gbps link comprising 4 x 10 Gbps lanes therefore requires eight strands. MPO can terminate this type of parallel optical link more efficiently than separate LC-terminated strands. An MPO connector capable of carrying 24 or 32 fibers has the same footprint as a duplex LC pair.



Multi-fiber push-on (MPO) patch cord. (Image © 123RF.com.)



Where there are multiple strands within a single cable, the strands are color-coded (TIA/EIA 598) to differentiate them.

Wavelength Division Multiplexing

A duplex fiber channel link uses one transmit lane and one receive lane and requires two fiber strands. Parallel fiber uses bundles of lanes working at 10 Gbps or 25 Gbps to implement 40 Gbps or 100 Gbps links. These channel links require between eight and twenty strands.

Wavelength Division Multiplexing (WDM) is a means of using one or two strands to provision multiple channels.

Bidirectional Wavelength Division Multiplexing

Bidirectional (BiDi) transceivers support transmit and receive signals over the same strand of fiber. This uses WDM to transmit the Tx and Rx signals over slightly shifted wavelengths, such as 1,310 nm for Tx and 1,490 nm for Rx. BiDi transceivers must be installed in opposite pairs, so the downstream transceiver would have to use 1,490 nm for Tx and 1,310 for Rx. **Bidirectional wavelength division multiplexing (BWDM)** links are documented in Ethernet standards (1000BASE-BX and 10GBASE-BX).

Coarse and Dense Wavelength Division Multiplexing

Coarse Wavelength Division Multiplexing (CWDM) supports up to 16 wavelengths and is typically used to deploy four or eight bidirectional channels over either a single fiber strand or unidirectional channels over dual fiber strands (one strand for transmit, the other for receive). **Dense Wavelength Division Multiplexing (DWDM)** provisions greater numbers of channels (20, 40, 80, or 160). This means that there is much less spacing between each channel and that it requires more precise and expensive lasers. CWDM and DWDM transceivers support multi-channel 1 G, 10 G, and 40 G Ethernet links. The transceivers must be installed in opposite pairs.

Lesson 2.5

Physical Installation Factors



Exam Objectives Covered

2.4 Explain important factors of physical installations.

Cabling is only one part of physical installation. There are also important considerations around installing switches, routers, access points, and servers within a site. Power, temperature, humidity, and fire risks can all adversely affect the reliability of network services. There are also security and access control factors to account for. While you might not be responsible for site design at this stage in your career, you should be able to explain the importance of these factors when performing maintenance and upgrades.

As you study this lesson, answer the following questions:

- How do rack systems ensure density and security?
- What considerations must be made for supplying power to run networking equipment?
- What are the risks from environmental factors, and how can they be controlled?

Rack Systems

Networking equipment should be installed within secure areas. Within a building, these can be referred to as telecommunications closets, equipment rooms, or server rooms. A whole facility dedicated to provisioning servers is called a datacenter. All these spaces should be dedicated to appliance and server installation and not used for other kinds of storage. They need physical access controls so that only authorized persons are allowed entry.

Within a telecommunications closet, server room, or datacenter, equipment is installed in racks. A **rack** is a specially configured steel shelving system designed for standard-size equipment. Using a rack allows equipment to be stored more securely and compactly than ordinary desks or shelving would allow for. The concept of installing more computing appliances in a smaller space is referred to as density.

Network appliances and server hardware designed for rack-mounting are EIA standard 19" / 48.26 cm width. Each appliance can be screwed into the rack directly. Nonstandard components, such as a tower server or monitor, can be installed on shelves.



A server rack is a compact and secure installation option for servers and networking appliances.




If there is little need to remove it for upgrades or maintenance, an appliance can be screwed directly into the rack. However, devices are often mounted on rail kits so that they can be slid out of the rack for hardware maintenance and upgrades.

Rack height is measured in “U” units of 1.75" / 4.45 cm. Racks are sold in heights from 8U to 48U. Rack-compatible equipment is designed with a vertical height quoted in U so you can plan exactly how much vertical space you require.

Most racks are designed to be freestanding, though smaller wall-mounted cabinet units are also available. Freestanding racks can be bolted together in rows. There should be about 3 feet (1 meter) clearance aisle for service access and airflow. Multiple rows should be placed back-to-back not front to back to maximize cooling. This is referred to as a hot aisle/cold aisle layout.

Rack-mounted appliances are usually designed with intake fans on the front to draw in cool air and exhaust fans on the back to expel warm air. Some switch models can be configured between **port-side exhaust**, where hot air is expelled on the same side as the port interfaces, and port-side intake. Port-side intake allows a switch to be installed with ports facing the front of the rack, which might be better for some cable management scenarios.




A hot aisle/cold aisle layout ensures that hot air expelled from exhaust vents does not contaminate cool air drawn in through intake vents.

Side panels and blanking plates should cover unused rack slots to improve airflow. Each rack can be installed with lockable doors (front and rear) to prevent unauthorized access to the equipment.

Humidity and Temperature

Environmental controls mitigate the loss of availability through mechanical issues with equipment, such as overheating. Building control systems maintain an optimum working environment for different parts of the building. The acronym **HVAC (Heating, Ventilation, Air Conditioning)** is often used to describe these services. An HVAC uses temperature sensors and moisture detection sensors (to measure humidity).

Servers and appliances are fitted with internal sensors to monitor conditions within the device chassis. These can report problems such as excessive temperatures within the device chassis, fan speeds, component failure, and chassis intrusion to a monitoring system.



CPUID's HWMONITOR app can report temperatures from sensors installed on PC components. (Screenshot used by permission of CPUID.)

Sensors can also be installed to measure ambient environmental conditions for a network rack or enclosure or within a server room or equipment closet. The following environmental factors need monitoring:

- **Temperature**—High temperature will make it difficult for device and rack cooling systems to dissipate heat effectively. This increases the risk of overheating of components within the device chassis and consequent faults.
- **Humidity**—More water vapor in the air risks condensation forming within a device chassis, leading to corrosion and short circuit faults. Conversely, very low humidity increases risks of static charges building up and damaging components.
- **Electrical**—Computer systems need stable power supply, free from outages (power failure), voltage dips (under-voltage events), and voltage spikes and surges. Sensors built into power distribution systems and backup battery systems can report deviations from a normal power supply.
- **Flooding**—There may be natural or person-made flood risks from nearby watercourses and reservoirs or risks from leaking plumbing or fire suppression systems. Electrical systems need to be shut down immediately in the presence of any significant amount of water.

Power Management

All types of network appliances require a stable power supply to operate. Electrical events, such as voltage spikes or surges, can crash computers, switches, and routers, while loss of power from under-voltage events or power failures will cause equipment to fail. An under-voltage event is where the voltage drops briefly, while a power failure is a complete loss of power lasting seconds or more. Power management means deploying systems to ensure that equipment is protected against these events and that network operations can either continue uninterrupted or be recovered quickly.

Power Load and Voltage

The circuits supplying grid power to a rack, network closet, or server room must meet the load capacity of all the installed equipment (plus room for growth). Consequently, the alternating current (AC) circuits to a server room will typically be higher capacity than domestic or office circuits (30 or 60 amps as opposed to 13 amps, for instance). They might also be run at a higher voltage (240 VAC, rather than 120 VAC).

The power supply for each appliance has a wattage rating. For example, a basic switch might be 20 watts, while a 1U server might be 200 watts. Wattage is calculated as $V(\text{olts}) * \text{Current (Amps)}$. To calculate the maximum load for a rack, add up the watts used by each appliance power supply and divide by the circuit voltage. For example, if a rack contains equipment that draws 2,000 watts in total, and the circuit VAC is 240, the amperage is 8.3. A single 30 amp circuit could supply three such racks.



If the circuits were 120 VAC, the amperage would be double. This is why equipment room and datacenter facilities tend to use high voltage circuits.

Power Distribution Units

Each circuit might be run through a **power distribution unit (PDU)**. A PDU has circuitry to “clean” the power signal, provides protection against spikes, surges, and under-voltage events, and can integrate with an uninterruptible power supply (UPS).

On a smaller scale, PDUs are also available as “strip” sockets that can take a higher load than a typical 13 amp rated strip. Such sockets are rack mounted and can be oriented horizontally or vertically to allow for different cabling and layout options. PDUs also often support remote power monitoring functions, such as reporting load and status, switching power to a socket on and off, or switching sockets on in a particular sequence.

Battery Backups and Uninterruptible Power Supplies

If there is loss of power, system operation can be sustained for a few minutes or hours (depending on load) using battery backup. Battery backup can be provisioned at the component level for storage device or array cache. The battery protects any read or write operations cached at the time of power loss.

At the system level, an **uninterruptible power supply (UPS)** will provide a temporary power source in the event of a power failure. UPS runtime may range from a few minutes for a desktop-rated model to hours for an enterprise system. In its simplest form, a UPS comprises a bank of batteries and their charging circuit plus an inverter to generate AC voltage from the direct current (DC) voltage supplied by the batteries. Different UPS models support different power outputs and form factors—from desktop to rack mounted depending on your needs.

Fire Suppression

Health and safety legislation dictates what mechanisms an organization must put in place to detect and suppress fires. Some basic elements of fire safety include the following:

- Well-marked fire exits and an emergency evacuation procedure that is tested and practiced regularly.
- Building design that does not allow fire to spread quickly, by separating different areas with fire-resistant walls and doors.
- Automatic smoke or fire detection systems, as well as alarms that can be operated manually.

Fire suppression systems work on the basis of the fire triangle. The fire triangle works on the principle that a fire requires heat, oxygen, and fuel to ignite and burn. Removing any one of those elements provides fire suppression (and prevention). In the United States (and most other countries), fires are divided by class under the NFPA (National Fire Protection Association) system, according to the combustible material that fuels the fire. Portable fire extinguishers come in several different types, with each type being designed for fighting a particular class of fire. Notably, Class C extinguishers use gas-based extinguishing and can be used where the risk of electric shock makes other types unsuitable.

Premises may also be fitted with an overhead sprinkler system. Wet-pipe sprinklers work automatically, are triggered by heat, and discharge water. Wet-pipe systems constantly hold water at high pressure, so there is some risk of burst pipes and accidental triggering, as well as the damage that would be caused in the event of an actual fire. There are several alternatives to wet-pipe systems that can minimize damage that may be caused by water flooding the room:

- **Dry-pipe**—These are used in areas where freezing is possible; water only enters this part of the system if sprinklers elsewhere are triggered.
- **Pre-action**—A pre-action system only fills with water when an alarm is triggered; it will then spray when the heat rises. This gives protection against accidental discharges and burst pipes and gives some time to contain the fire manually before the sprinkler operates.
- **Halon**—Gas-based systems have the advantage of not short circuiting electrical systems and leaving no residue. Up until a few years ago, most systems used Halon 1301. The use of Halon has been banned in most countries as it is ozone depleting, though existing installations have not been replaced in many instances and can continue to operate legally.
- **Clean agent**—Alternatives to Halon are referred to as “clean agent.” As well as not being environmentally damaging, these gases are considered nontoxic to humans. Examples include INERGEN (a mixture of CO₂, Argon, and Nitrogen), FM-200/HFC-227, and FE-13. The gases both deplete the concentration of oxygen in the area (though not to levels dangerous to humans) and have a cooling effect. CO₂ can be used too, but it is not safe for use in occupied areas.

Lesson 2.6

Cable Troubleshooting



Exam Objectives Covered

5.2 Given a scenario, troubleshoot common cabling and physical interface issues.

5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

Applying a layer-by-layer approach to network troubleshooting can greatly assist with isolating symptoms and causes. In this topic, you will investigate some common issues that can affect cabled networks and identify the tools and techniques that can be used to solve problems at the Physical layer.

As you study this section, answer the following questions:

- How can you distinguish an issue caused by improper cable choice or termination from one caused by environmental factors?
- Which tools can you use to diagnose issues with copper and fiber optic cable and connectors?
- What is the difference between a regular cable tester and a cable certifier?
- Which tool would you use to find the end of a specific cable within a wiring closet?

Specification and Limitations

When troubleshooting a link, you will need to compare the expected performance with the actual current performance. To do this, you must understand how to assess and distinguish speed, throughput, and distance specifications and limitations.

Speed Versus Throughput

At the Physical layer, a signal transmitted over a communications channel consists of a series of events referred to as symbols. A symbol could be something like a pulse of higher voltage in an electrical current or the transition between the peak and the trough in an electromagnetic wave. The number of symbols that can be transmitted per second is called the baud rate. The baud rate is measured in hertz (or MHz or GHz).

At the Data Link layer, the nominal **bit rate**—or bandwidth—of the link is the amount of information that can be transmitted, measured in bits per second (bps), or some multiple thereof. In order to transmit information more efficiently, a signaling method might be capable of representing more than one bit per symbol. This also helps to overcome noise and detect errors. The use of these encoding methods means that the bit rate will be higher than the baud rate. In Ethernet terms, the bit rate is the expected performance of a link that has been properly installed to operate at 10 Mbps, 100 Mbps, 1 Gbps, or better.

The nominal bit rate will not often be achieved in practice. **Throughput** is an average data transfer rate achieved over a period of time excluding encoding schemes, errors, and other losses incurred at the Physical and Data Link layers. Throughput can be adversely affected by link distance and by interference (noise).



Throughput is typically measured at the Network or Transport layer. Often the term “goodput” is used to measure an averaged data transfer rate at the Application layer. This takes account of the effect of packet loss. Throughput is also sometimes measured as packets per second.

As well as bandwidth or throughput and **packet loss**, the speed at which packets are delivered is also an important network performance characteristic. Speed is measured as a unit of time—typically milliseconds (ms)—and is also referred to as **latency**, or delay.



The term “speed” is also used to describe how well or badly a link is performing in terms of throughput but do be aware of the distinction between bit rate and latency.

Distance Limitations, Attenuation, and Interference

Each type of media can consistently support a given bit rate only over a defined distance. Some media types support higher bit rates over longer distances than others. Attenuation and interference enforce distance limitations on different media types.

- **Attenuation** is the loss of signal strength, expressed in decibels (dB). dB expresses the ratio between two measurements; in this case, signal strength at origin and signal strength at destination.
- **Interference** (or noise) is anything that gets transmitted within or close to the channel that isn't the intended signal. This serves to make the signal itself difficult to distinguish, causing errors in data and forcing retransmissions. This is expressed as the signal to noise ratio (SNR).

Cable Issues

When troubleshooting cable connectivity, you are focusing on issues at the Physical layer. At layer 1, a typical Ethernet link for an office workstation includes the following components:

- Network transceiver in the host (end system).
- Patch cable between the host and a wall port.
- Structured cable between the wall port and a patch panel (the permanent link).
- Patch cable between the patch panel port and a switch port.
- Network transceiver in the switch port.

The entire cable path (patch cords plus permanent link) is referred to as a channel link.



Verifying patch cord connections. (Image by Kjetil Kolbjornsrud © 123RF.com.)

Assuming you are investigating link failure (complete loss of connectivity), the first step is to check that the patch cords are properly terminated and connected to the network ports. If you suspect a fault, substitute the patch cord with a known good cable.

If you cannot isolate the problem to the patch cords, test the transceivers. You can use a loopback tool to test for a bad port.



If you don't have a loopback tool available, another approach is to substitute known working hosts (connect a different computer to the link or swap ports at the switch). This approach may have adverse impacts on the rest of the network, however, and issues such as port security may make it an unreliable method.

If you can discount faulty patch cords and bad network ports/NICs, you will need to use tools to test the structured cabling. The solution may involve installing a new permanent link, but there could also be a termination or external interference problem.

Cable Category Issues

When troubleshooting a permanent link, you should verify that the cable type is appropriate to the application. For example, you cannot expect 10 GbE Ethernet to run over an 80 meter Cat 5e link. You may also need to verify that unshielded cable has not been installed where shielded or screened cable would be more suitable. Using an incorrect cable type might result in lower-than-expected speed and/or numerous checksum errors and link resets. Check the identifier printed on the cable jacket to verify the type that has been

When evaluating whether a cable category is suitable for a given use in the network, consider the following factors:

- Cat 5e supports Gigabit Ethernet and could still be an acceptable choice for providing network links for workstations, but most new installations and upgrades would now use Cat 6 or better.



Unlike Ethernet and Fast Ethernet, Gigabit Ethernet uses all four pairs for transmission and is thus more sensitive to crosstalk between the wire pairs.

- Cat 6 can support 10 Gbps, but over a 55 meters maximum distance.
- Cat 6A is an improved specification cable that can support 10 Gbps over 100 meters. Cat 6A cable is bulkier than Cat 5e, and the installation requirements more stringent, so fitting it within pathways designed for older cable can be problematic. TIA/EIA standards recommend Cat 6A for use in healthcare facilities, with Power over Ethernet (PoE) 802.3bt installations, and for horizontal connections to wireless access points.




Cabling is not the only part of the wiring system that must be rated to the appropriate category. For Gigabit Ethernet and better, the performance of connectors becomes increasingly critical. For example, if you are installing Cat 6A wiring, you must also install Cat 6A patch panels, wall plates, and connectors.

- Cat 7 cable is always of a screened/shielded type and is rated for 10 Gbps applications up to 100 meters (328 feet). Cat 7 is not recognized by TIA/EIA but appears in the cabling standards created by the ISO (ISO/IEC 11801). It must be terminated with GG45 or TERA connectors rather than standard RJ45 connectors.
- Cat 8 is intended for use in datacenters only for short patch cable runs that make top-of-rack connections between adjacent appliances. ISO defines two variants; 8.1 (Class I) is equivalent to TIA/EIA Cat 8 and uses RJ45 connectors while 8.2 (Class II) must use outer shielding or screening and GG45 or TERA connectors.

From a safety point of view, you must also ensure that the cable jacket type is suitable for the installation location, such as using plenum-rated cable in plenum spaces and plenum- or riser-rated cable in riser spaces.

Cable Testers

If the cable is not accessible, cable testing tools can also be used to diagnose intermittent connectivity or poor performance issues. A **cable tester** reports detailed information on the physical and electrical properties of the cable. For example, it can test and report on cable conditions, crosstalk, attenuation, noise, resistance, and other characteristics of a cable run. Devices classed as certifiers can be used to test and certify cable installations to a performance category—for example, that a network is TIA/EIA 568 Category 6A compliant. They use defined transport performance specifications to ensure an installation exceeds the required performance characteristics for parameters such as attenuation and crosstalk.



A technician using a cable certifier. (Image by Wavebreak Media © 123RF.com.)



Cable testing tools can be used for troubleshooting and verification. It is best to verify wiring installation and termination just after you have made all the connections. This means you should still have access to the cable runs. Identifying and correcting errors at this point will be much simpler than when you are trying to set up end user devices.

Wire Map Testers and Tone Generators

Fully featured cable testers/certifiers are expensive. A simpler **wire map tester** device can be used to detect improper termination issues. To perform a wire map test, the base unit is connected to one end of the cable and a remote unit to the other. When the test is activated, an LED for each wire conductor lights up in sequence. If an LED fails to light or does not light in sequence, there is a problem with the cable and/or termination. Wire map testers can identify the following problems:

- **Continuity (open)**—A conductor does not form a circuit because of cable damage or because the connector is not properly wired.
- **Short**—Two conductors are joined at some point, usually because the insulating wire is damaged, or a connector is poorly wired.
- **Incorrect pin-out/incorrect termination/mismatched standards**—The conductors are incorrectly wired into the terminals at one or both ends of the cable. The following transpositions are common:
 - **Reversed pair**—The conductors in a pair have been wired to different terminals (for example, from pin 3 to pin 6 and pin 6 to pin 3 rather than pin 3 to pin 3 and pin 6 to pin 6).
 - **Crossed pair (TX/RX transposed)**—The conductors from one pair have been connected to pins belonging to a different pair (for example, from pins 3 and 6 to pins 1 and 2). This may be done deliberately to create a crossover cable, but such a cable would not be used to link a host to a switch.

Another potential cable wiring fault is a split pair. This is where both ends of a single wire in one pair are wired to terminals belonging to a different pair. This type of fault can only be detected by a cable tester that measures crosstalk.

A network **tone generator** (or toner) and probe are used to trace a cable from one end to the other. This may be necessary when the cables are bundled and have not been labeled properly. This device is also known as a Fox and Hound. The tone generator is used to apply a signal on the cable to be traced so that you can use the probe to identify the same cable within a bundle or duct.

Attenuation and Interference Issues

If a cable link is too long, **decibel (dB) loss** (or insertion loss) may mean that the link experiences signal degradation problems with high error rates and retransmissions (frame or packet loss) resulting in reduced speeds and possibly loss of connectivity. Insertion loss is measured in decibels (dB) and represents the ratio of the received voltage to the original voltage.

A dB expresses the ratio between two values using a logarithmic scale. A logarithm is a mathematical tool for performing complex multiplication and division exponential operations as simpler additions and subtractions. The essential point is that a logarithmic scale is nonlinear, so a small change in dB value represents a large change in the performance measured. The following reference points are useful to remember:

- +3 dB means doubling, while -3 dB means halving.
- +6 dB means quadrupling, while -6 dB relates to a quarter.
- +10 dB means 10 times the ratio, while -10 dB is a tenth.



For a longer primer on dB math, view the presentation at internetsociety.org/wp-content/uploads/2017/10/dB-Math.pdf.

The maximum value allowed for insertion loss depends on the link category. For example, Cat 5e at 100 MHz allows up to 24 dB, while Cat 6 allows up to 21.7 dB at 250 MHz. When you are measuring insertion loss itself, smaller values are better (20 dB insertion loss is better than 22 dB, for instance). A cable certifier is likely to report the margin, which is the difference between the actual loss and the maximum value allowed for the cable standard. Consequently, higher margin values are better. For example, if the insertion loss measured over a Cat 5e cable is 22 dB, the margin is 2 dB; if another cable measures 23 dB, the margin is only 1 dB, and you are that much closer to not meeting acceptable link standards. Higher grade or shielded cable may alleviate the problem; otherwise, you will need to find a shorter cable run or install a repeater or additional switch.

Careful cable placement is necessary during installation to ensure that the wiring is not subject to interference from sources such as electrical power cables, fluorescent lights, motors, electrical fans, radio transmitters, and so on. Electromagnetic interference (EMI) is something that should be detected when the cable is installed, so you should suspect either some new source that has been installed recently or some source that was not taken into account during testing (machinery or power circuits that weren't activated when the installation testing took place, for instance). Interference from nearby data cables is also referred to as alien crosstalk.



Radio frequency interference (RFI) is EMI that occurs in the frequencies used for radio transmissions.

Crosstalk Issues

Crosstalk usually indicates a problem with bad wiring (poor quality or damaged or the improper type for the application), a bad connector, or improper termination. Check the cable for excessive untwisting at the ends and for kinks or crush points along its run. Crosstalk is also measured in dB, but unlike insertion loss, higher values represent less noise. Again, the expected measurements vary according to the cable category and application. There are various types of crosstalk that can be measured:

- **Near End (NEXT)**—This measures crosstalk on the receive pairs at the transmitter end and is usually caused by excessive untwisting of pairs or faulty bonding of shielded elements.
- **Attenuation to Crosstalk Ratio, Near End (ACRN)**—This is the difference between insertion loss and NEXT. ACRN is equivalent to a signal-to-noise ratio (SNR). A high value means that the signal is stronger than any noise present; a result closer to zero means the link is likely to be subject to high error rates.
- **Attenuation-to-Crosstalk Ratio, Far End (ACRF)**—Far-end crosstalk (FEXT) is measured on the receive pairs at the recipient end. The difference between insertion loss and FEXT gives ACRF, which measures cable performance regardless of the actual link length.
- **Power Sum**—Gigabit and 10 GbE Ethernet use all four pairs. Power sum crosstalk calculations (PSNEXT, PSACRN, and PSACRF) confirm that a cable is suitable for this type of application. They are measured by energizing three of the four pairs in turn.
- **Alien Crosstalk**—This is signal traffic from cables in close proximity that causes interference to a disturbed or victim cable. This is commonly caused by cinching a cable bundle with ties too tightly and by poorly terminated cabling.



Complete loss of connectivity indicates a break in the cable (or a completely faulty installation), while intermittent loss of connectivity is more likely to be caused by attenuation, crosstalk, or noise.

Fiber Optic Cable Testing Tools

When you are working with fiber optic cabling, it is important to understand that any mismatch between the cables coupled together will result in data loss. This can occur if the fiber cables are not properly aligned, are different sizes, or may have suffered damage (broken/misshaped fiber strands) during transport. If you connect single mode fiber to multimode fiber, you will introduce a catastrophic signal loss of up to 99%. Even connecting fiber cables of the same type but with different diameters can cause a loss of up to 50% of the signal strength.

Whenever a connector is installed on the end of fiber optic cables, a degree of signal loss occurs. This is called insertion loss. In addition, some of the light that is lost is reflected directly back down the cable toward the source. This is called back-reflection or optical return loss (ORL). Ultra Physical Contact (UPC) and Angled Physical Contact (APC) polishing reduce ORL reflections, but remember that mating an APC connector to a non-APC port causes major insertion loss. Because of this, APC connectors are always colored green to keep you from mixing them with non-APC connectors.

Visual Fault Locator

If a break is identified in an installed cable, the location of the break can be found using a **visual fault locator**. There are different models for short and long link distances, and they can be supplied with adapters for different connector types (ST, SC, or LC). The tool shines visible light down the cable and glows brightly at the point where a cable is broken, excessively bent, or improperly spliced.

Dirty Optical Cables

Dirt, dust, or grease in the transmission path will greatly reduce signal strength or block transmission completely. Most commonly, this occurs at a connector. Connectors should be covered with a dust cap when removed, and the surrounding area should be dust free before performing a disconnection. Connectors should be cleaned using solvent designed for fiber optics, taking care not to apply excess solvent. The wet-to-dry method applies a drop of solvent to a lint-free cloth and moves the connector from the wet drop across a dry part. Contamination could also be introduced when a cable is spliced. Ensure splicing equipment is cleaned according to the manufacturer's instructions before every splice operation.



The powerful light sources used by fiber optics are a hazard. Wear appropriate safety goggles, and never look directly at an active transceiver port or the end of a fiber cable. Point a cable at a flat surface to confirm whether visible light is being transmitted, or use a smartphone camera to detect whether infrared light is being transmitted.

Cable Troubleshooting Strategies

Sometimes, cables will fail. This can happen for many different reasons. Your job will be to troubleshoot connection issues and find the root cause. Common issues to network connections include physical damage to the cable, loose connections, interference from other devices, and issues with the network adapter or drivers. Let's look at steps we can take to troubleshoot cable issues:

1. Physical Inspection
 - Check the cable for any visible damage such as cuts, kinks, or severe bends.
 - Ensure that the connectors are not damaged and are securely plugged into the network device and the computer.
2. Reseat the Cable
 - Unplug the cable from both ends and then plug it back in. This can resolve loose connection issues.
3. Verify Drivers
 - If the problem persists, the issue could be the drivers or a physical problem with the network adapter.
 - Open Device Manager on your computer, find your network adapter in the list, and check if it's working properly.
 - If it is not working correctly, you may need to update the drivers or replace the network adapter.

Module 2

Summary

You should be able to summarize the properties of copper and fiber optic media and connectors and match them to an appropriate Ethernet standard for a particular solution. Additionally, you should be able to use appropriate tools to diagnose symptoms and causes of common cable connectivity issues.

Guidelines for Deploying and Troubleshooting Ethernet Cabling

Consider these best practices and guidelines when you are installing and maintaining Ethernet networks:

- Select an Ethernet standard, media type, connectors, and transceivers that meet the requirements for different types of network segment:
 - 1000BASE-T over Cat 5e or Cat 6 for client access/office networks or 10GBASE-T over Cat 6A or better for high-bandwidth requirements.
 - 10GBASE-CR/40GBASE-CR4 twinax or 10GBASE-SR MMF for datacenter applications.
 - 1000BASE-SX or 10GBASE-SR over MMF for backbones.
 - 1000BASE-LX or 10GBASE-LR for site-to-site links.
 - Wave division multiplexing to get more bandwidth from existing fiber.
- Consider the factors that can affect the performance of network media, such as electromagnetic interference and attenuation and whether shielded copper or fiber optic cable will be required to ensure reliable performance.
- Follow the 568 Commercial Building Telecommunications Cabling Wiring Standard to apply a structured cabling design with patch panels to distribute cabling from communications closets to work areas. Use either T568A or T568B termination consistently.
- Use plenum cables in designated plenum spaces of a building to comply with fire codes and use PVC in non-plenum spaces.
- Use appropriate tools to prepare cable and terminate to either punch down blocks or to connectors.
- Document expectations of speed/throughput over link distances to identify when performance is reduced. Verify that cables are being used for their proper application, given factors such as shielding against external interference, crossover, rollover/console, and plenum/riser.
- Test the connectivity path (ports, patch cords, structured links) methodically.
- Use wire mappers and cable testers to identify faults in copper cable. Use power meters and visual fault locators for fiber optic plant.

Module 3

Configuring Interfaces and Switches

Module Introduction

Cabling establishes the links between nodes on the network, but each node also requires a network interface that both connects the cabling and performs the Data Link layer addressing and framing functions necessary for communications.

Also, not many networks are established by directly connecting each end system to every other local system. Cabling and support costs are reduced by using intermediate systems to establish local networks. These intermediate systems are deployed as network appliances such as hubs, bridges, and switches. Installing and configuring these devices will be a regular task for you during your career in network administration.

Module Objectives

In this lesson, you will do the following:

- Deploy networking devices.
- Explain network interfaces.
- Deploy common Ethernet switching features.
- Troubleshoot transceiver and switching issues.

Lesson 3.1

Network Interfaces



Exam Objectives Covered

- 1.5 Compare and contrast transmission media and transceivers.*
- 5.2 Given a scenario, troubleshoot common cabling and physical interface issues.*


As a network technician, you will frequently be involved with installing, configuring, and troubleshooting Ethernet network interfaces and transceivers. A network interface is the means by which a node is connected to the media and exchanges data with other network hosts. This involves both a physical connection to the network media and the use of a data link protocol such as Ethernet to perform addressing and framing. Understanding the properties of these components is essential for selecting appropriate products and configuring them correctly.

As you study this lesson, answer the following questions:

- What is the function of a transceiver?
- At which OSI layers do transceivers and network interfaces operate?
- What is Ethernet framing, and how are network interfaces uniquely addressed?
- What is the purpose of the cyclic redundancy check (CRC)?
- What does the MAC address ff:ff:ff:ff:ff:ff indicate?
- What are the symptoms of transceiver issues?

Network Interface Cards

The transceiver component responsible for physically connecting a host to the transmission medium is implemented in a network interface card/controller (NIC), also referred to as a network adapter. Most Ethernet adapters designed for use with copper cabling now support Gigabit Ethernet. A different kind of adapter would have to be provisioned for a fiber link. Adapters that support 10 GbE or 40 GbE come at a considerable price premium over basic Gigabit models. A NIC may also provision multiple ports on the same card. This allows either connections to different networks or aggregating the separate links into a higher bandwidth channel.



Network interface card (NIC) with two RJ45 ports. (Image © 123RF.com.)


Each Ethernet network interface port has a unique hardware address known as the **Media Access Control (MAC) address**. This may also be referred to as the Ethernet address (EA) or, in IEEE terminology, as the **extended unique identifier (EUI)**. A MAC address is also referred to as a local or physical address.

Modular Transceivers

A network might involve the use of multiple types of cabling. When this occurs, server, switch, and router equipment must be able to terminate different cable and connector types. Enterprise servers, switches, and routers are available with modular, hot-swappable transceivers for different types of fiber optic and copper connections.

SFP/SFP+

Small form-factor pluggable (SFP) transceivers use LC connectors and support Gigabit Ethernet data rates. Enhanced SFP (SFP+) is an updated specification to support 10 GbE but still uses the LC form factor. There are different modules to support the various Ethernet standards and fiber mode type (10GBASE-SR versus 10GBASE-LR, for instance).



Switch with hot-pluggable SFP fiber transceivers. (Image by Zdenek Maly © 123RF.com.)



You will often see the term “MSA” in conjunction with modular transceivers. Multi-source agreement (MSA) is intended to ensure that a transceiver from one vendor is compatible with the switch/router module of another vendor.



There are also transceivers that support the Fibre Channel storage area network (SAN) protocol. These are not compatible with Ethernet switches.

QSFP/QSFP+

Quad small form-factor pluggable (QSFP) is a transceiver form factor that supports 4 x 1 Gbps links, typically aggregated to a single 4 Gbps channel. Enhanced quad small form-factor pluggable (QSFP+) is designed to support 40 GbE by provisioning 4 x 10 Gbps links. QSFP+ is typically used with parallel fiber and multi-fiber push-on (MPO) termination. QSFP+ can also be used with Wavelength Division Multiplexing (WDM) Ethernet standards.



There are also SFP+ and QSFP+ transceivers with Direct Attach Copper (DAC) ports.



WDM transceivers must be installed in matched pairs. The Tx wavelength used on one side must match the Rx wavelength used on the other.

Transceiver Mismatch Issues

Modular transceivers are designed to be used with a specific type of optical fiber. For example, transceivers designed for single mode fiber use laser diodes while multimode fiber transceivers use LEDs or a different type of laser (VCSEL). Different transceivers are designed to work at different optical wavelengths (typically 850 nm, 1300 nm, or 1550 nm) and support different Ethernet standards and bit rates. This means it is important to check the manufacturer's documentation for the interface to ensure the correct fiber type is used, not only for the fiber optic cable, but also for the fiber patch cords used to connect to it at each end. Mismatches between cable, patch cords, and interfaces may lead to significant signal loss.

As well as the fiber mode, there are three main ways of deploying fiber:

- Duplex uses two strands for transmit (Tx) and receive (Rx).
- Parallel uses multiple strands (typically eight or twenty) to implement Tx and Rx channels.
- Wavelength Division Multiplexing uses either a single bidirectional strand or dual unidirectional strands to implement multiple channels, distinguished by wavelengths.

Each type is implemented by a different transceiver model. These might need to be installed in matched pairs. For example, when using BiDi, the Tx wavelength used by one transceiver must connect through to the same wavelength on the other transceiver's Rx port.

Transceiver Signal Strength Issues

Although fiber optic cable does not suffer from attenuation in the same way as copper cable or to the same extent, there will still be some loss of signal strength from one end of the connection to the other. This is due to microscopic imperfections in the structure of the glass fiber and in the smoothness of the edge of the core, leading to some small fraction of the light within the core being scattered or absorbed. Attenuation can be tested using an optical source and optical power meter (or fiber light meter), which may be purchased together as a fiber testing kit.

An **optical link budget**, or loss budget, is the amount of loss suffered by all components along a fiber transmission path. This is calculated using the following parameters:

- **Attenuation**—This is the loss over the length of the cable, based on fiber type and the wavelength used. Single mode has a loss of up to 0.4 dB/km, while multimode can be from 0.8 dB/km to 3 dB/km.
- **Connectors**—Each connector in the path incurs a loss, usually assumed to be 0.75 dB.
- **Splices**—Additional splices in the cable are budgeted at around 1 dB for mechanical and 0.3 dB for fusion.

Typically, an estimated loss budget is calculated when planning the link. The link is tested at deployment to derive an actual value. Differences between these values may reveal an installation fault or some unexpected source of signal loss.



FOA have a loss budget calculator at thefoa.org/tech/ref/Loss_Budget/Loss_Budget.htm.

The loss budget must be less than the power budget. The power budget is calculated from the transceiver transmit (Tx) power and receiver (Rx) sensitivity, which are both typically measured in dB per milliwatt or dBm. For example, if Tx is -8 dBm and Rx is -15 dBm, then the power budget is 7 dB.



dBm measures signal strength against a reference value, where 0 dBm is 1 milliwatt. A negative dBm is typical of Ethernet transceivers, which output less than 1 mw.

If the loss budget is 5 dB, the margin between the power budget and loss budget will be 2 dB. Margin is a safety factor to account for suboptimal installation conditions (such as bends or stress), aging, repair of accidental damage (additional splices), and performance under different thermal conditions (extreme temperatures can cause loss).

If the margin between the transmitter power and link budget is low, the link is less likely to achieve the expected bandwidth. There may be opportunities to improve performance with better or fewer splices, or it may be necessary to use an amplifier to boost the signal. Most outdoor plans would be designed with a margin of at least 5 dB. In a datacenter where conditions are less variable a lower margin might be acceptable.

Ethernet Frame Format

The transceiver implements a link at the Physical layer, but Ethernet interfaces also perform addressing and framing functions at layer 2 of the OSI model. This is referred to as the Data Link layer.

Ethernet encapsulates the payload from higher layer protocols within a protocol data unit (PDU) called a frame. The basic format of an Ethernet frame and **Ethernet headers** is shown in the following figure.



Header fields in an Ethernet frame.

Preamble

The preamble and start frame delimiter (SFD) are used for clock synchronization and as part of the CSMA/CD protocol to identify collisions. The preamble consists of 8 bytes of alternating 1s and 0s with the SFD being two consecutive 1s at the end. This is not technically considered to be part of the frame.

EtherType

The 2-byte EtherType field is usually used to indicate the type of protocol in the frame payload. For example, a frame carrying an IPv4 packet would have an EtherType value of 0x0800; one carrying IPv6 data would be 0x86DD.



You might see the 2-byte field described as the EtherType/Length field. When Ethernet was being developed, there were several alternative frame formats, one of which used the 2-byte field to indicate the frame length. To maintain compatibility, EtherTypes are values of 0x0600 (1536 in decimal) or greater. Anything less than that would be interpreted as the payload length.

Error Checking

The error checking field contains a 32-bit (4-byte) checksum called a **cyclic redundancy check (CRC)** or frame check sequence (FCS). The CRC is calculated based on the contents of the frame; the receiving node performs the same calculation and, if it matches, accepts the frame. There is no mechanism for retransmission if damage is detected nor is the CRC completely accurate at detecting damage; these are functions of error checking in protocols operating at higher layers.


Media Access Control Address Format

The source and destination addresses in a frame are 48-bit Media Access Control (MAC) identifiers. The notation format of this number differs depending on the system architecture. It is often displayed as six groups of two hexadecimal digits with colon or hyphen separators or no separators at all (for example, 00:60:8c:12:3a:bc or 00608c123abc) or as three groups of four hex digits with period separators (0060.8c12.3abc, for instance).

Burned-in Addresses

The IEEE gives each network adapter manufacturer a range of numbers, and the manufacturer hard codes every interface produced with a unique number from their range. This is called the burned-in address or the universal address. The first six hex digits (3 bytes or octets), also known as the organizationally unique identifier (OUI), identify the manufacturer of the adapter. The last six digits are a serial number.

An organization can decide to use locally administered addresses in place of the manufacturers' universal coding systems. This can be used to make MACs meaningful in terms of location on the network, but it adds a significant amount of administrative overhead. A locally administered address is defined by changing the universal/local (U/L) bit from 0 to 1. The rest of the address is configured using the card driver or network management software. It becomes the network administrator's responsibility to ensure that all interfaces are configured with a unique MAC.



Captured Ethernet frame showing the resolved OUI and IG and LG bits in the destination (broadcast) and source addresses. Note that Wireshark uses local/global (L/G) terminology rather than universal/local (U/L). (Screenshot courtesy of Wireshark.)

Broadcast Address

The I/G bit of a MAC address determines whether the frame is addressed to an individual node (0) or a group (1). The latter is used for **broadcast** and **multicast** transmissions. A MAC address consisting entirely of 1s is the broadcast address (ff:ff:ff:ff:ff:ff).

A **unicast** transmission is one sent to an individual host. This is achieved by adding the host's unique MAC address as the destination address. When a frame uses the broadcast address as the destination address, it should be processed by all nodes that receive the frame. These nodes are said to be within the same **broadcast domain**.

Lesson 3.2

Ethernet Switches



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 2.2 Given a scenario, configure switching technologies and features.

Most networks use intermediate systems to reduce cabling costs and complexity. In this topic, you will summarize the functions of hub, bridge, and switch appliances working at the Physical and Data Link layers.

As you study this lesson, answer the following questions:

- A host on a network sends a frame to the hub. Which other devices on the network will see this frame?
- A host on a network sends a frame to a switch. Which other devices on the network will see this frame?
- What are the similarities and differences between a bridge and a switch?
- What are the advantages of using switches instead of hubs?
- What is the difference between a managed and an unmanaged switch?

Hubs


Most Ethernet networks are implemented so that each end system node is wired to a central intermediate system. In early types of Ethernet, this function was performed by a hub. While hubs are no longer widely deployed as standalone appliances, it is important to understand the basic functions they perform.

A hub acts like a multiport repeater so that every port receives transmissions sent from any other port. As a repeater, the hub works only at the Physical layer. Electrically, the network segment looks like a single length of cable. Consequently, every hub port is part of the same shared media access area and within the same collision domain. All node interfaces are **half-duplex**, using the CSMA/CD protocol, and the media bandwidth (10 Mbps or 100 Mbps) is shared between all nodes.



A broadcast transmission is sent to all hosts in the same logical network area. In Ethernet, this is accomplished by using the broadcast MAC address ff:ff:ff:ff:ff:ff. A unicast transmission is addressed to a single host only, using its MAC address. With hubs, all interfaces receive all unicast and broadcast transmissions. Hosts are typically configured to ignore unicast transmissions that are not addressed to them. However, setting an interface to promiscuous mode allows a host to capture (or “sniff”) all unicast transmissions sent via the hub. This is a major security weakness of hubs.

When Ethernet is wired with a hub there needs to be a means of distinguishing the interface on an end system (a computing host) from the interface on an intermediate system (the hub). The end system interface is referred to as **medium dependent interface (MDI)**; the interface on the hub is referred to as MDI crossover (MDIX). This means that the transmit (Tx) wires on the host connect to receive (Rx) wires on the hub.




Hub-based Ethernet communications. (Images © 123RF.com.)

There are no configuration options for a hub. You just connect the device to a power source and then connect the network cables for the hosts that are going to be part of the network segment served by the hub.

Bridges


An Ethernet bridge works at the Data Link layer (layer 2) to establish separate physical network segments while keeping all nodes in the same logical network. This reduces the number of collisions caused by having too many nodes contending for access.



Collision and broadcast domains. (Images © 123RF.com.)

The previous figure shows how a bridge creates separate collision domains. Each hub is a shared access media area. The nodes connected to the hubs share the available bandwidth—a 100 Mbps Ethernet for domain A and a 10 Mbps Ethernet for domain B—because only one node within each collision domain can communicate at any one time. The bridge isolates these segments from each other, so nodes in domain B do not slow down or contend with nodes in domain A. The bridge does allow nodes to communicate with the other collision domain. It does this by forwarding only the appropriate traffic. This creates a single logical network, referred to as a layer 2 broadcast domain.

An Ethernet bridge builds a forwarding database to track which addresses are associated with which of its ports. When the bridge is initialized, the database's **MAC address table** is empty, but information is constantly added as the bridge listens to the connected segments. Entries are flushed out of the table after a period to ensure the information remains current.




Collision and broadcast domains. (Images © 123RF.com.)

If no record of the MAC address exists or the frame is a broadcast or multicast, then the bridge floods the frame to all segments except for the source segment (acting like a hub).

Switches

The problems created by contention can be more completely resolved by moving from a shared Ethernet system to a switched Ethernet. Hubs and bridges are replaced with switches. Gigabit Ethernet and faster can *only* be deployed using switches.

An Ethernet switch performs the same sort of function as a bridge, but in a more granular way and for many more ports than are supported by bridges. Each switch port is a separate collision domain. In effect, the switch establishes a point-to-point **full-duplex** link between any two network nodes. This is referred to as **microsegmentation**.



Switch operation. (Images © 123RF.com.)

Because each port is in a separate collision domain, collisions can occur only if the port is operating in half-duplex mode. This would only be the case if a legacy network card or a hub is attached to it. Even then, collisions affect only the microsegment between the switch port and the connected interface; they do not slow down the whole network. As with a bridge, traffic on all switch ports is in the same broadcast domain unless the switch is configured to implement **virtual LANs (VLANs)**.

Ethernet Switch Types

Ethernet switches from different vendors come in a variety of ranges to support various sizes of networks. While a basic model might feature 12 to 48 ports and little scope for expansion, advanced switches support high-speed interconnections and expandable capacity through plug-in modules plus power supply redundancy, management consoles, and transceivers for fiber optic connectivity.



An example of a workgroup switch. (Image © 123RF.com.)

The market is dominated by Cisco's Catalyst and Nexus platforms (over 55% of sales), but other notable vendors include HP Enterprise, Huawei, Juniper, Arista, Linksys, D-Link, NETGEAR, and NEC.

Ethernet switches can be distinguished using the following general categories:

- **Unmanaged versus managed**—On a SOHO network, switches are more likely to be unmanaged, standalone units that can be added to the network and run without any configuration. The switch functionality might also be built into an Internet router/modem. On a corporate network, switches are most likely to be managed. This means the switch settings can be configured. If a managed switch is left unconfigured, it functions the same as an unmanaged switch does.
- **Stackable**—Switches that can be connected together and operate as a group. The switch stack can be managed as a single unit.
- **Modular versus fixed**—A fixed switch comes with a set number of ports that cannot be changed or upgraded. A modular switch has slots for plug-in cards, meaning it can be configured with different numbers and types of ports.
- **Desktop versus rack-mounted**—Simple unmanaged switches with five or eight ports might be supplied as small freestanding units that can be placed on a desktop. Most larger switches are designed to be fitted to the standard-size racks that are used to hold networking equipment.



*An example of a modular chassis that allows provisioning multiple access switches.
(Image © 123RF.com.)*

Switch Interface Configuration

Configuration of a managed switch can be performed at a command-line interface (CLI). Once you have established a connection to the switch's management interface, you can configure settings for each of the switch port interfaces. These settings control the network link configured for each client device attaching to the switch. Most switch operating systems work in multiple command modes or hierarchies. For example, Cisco IOS has three principal modes:

- **User EXEC mode**—This is a read-only mode where commands can be used to run basic troubleshooting tools. This mode is indicated by the `>` prompt.
- **Privileged EXEC mode**—This allows the user to report the configuration, show system status, reboot or shut down the appliance, and backup and restore the system configuration. This mode is activated using the `enable` command from user EXEC mode. It is denoted by a `#` prompt.
- **Global configuration mode**—This allows the user to write configuration updates. It is activated by using the `configure terminal` command from privileged mode and indicated by a **(config)#** prompt.

Most switch CLIs also support `TAB` and/or use of `?` to list different ways to complete a partial instruction.

Interfaces are identified by type, slot, and port number. For example, GigabitEthernet 0/2 (or G0/2) is port #2 on the first 10/100/1000 slot (or only slot).



Stackable switches precede interface identifiers with a module ID. For example, GigabitEthernet 3/0/2 is the second port on the first slot in the third module in the stack. Note that this numbering does vary between manufacturers. Also, some start from zero and some from one.

Switches normally support a range of Ethernet standards so that older and newer network adapters can all be connected to the same network. In most cases, the port on the switch is set to auto-negotiate speed (10/100/1000) and full- or half-duplex operation. A static configuration can be applied manually if necessary.



If you don't use autonegotiation, you need to manually configure the speed and duplex to match both devices. For best performance, if one end of the connection is hard coded, it's advised to hard code the other end and not rely on autonegotiation.

To configure the first interface, from global config mode, run `interface GigabitEthernet0/1`. This changes the prompt to **(config-if)#**. Some of the main subcommands are the following:

- `shutdown` disables the interface; `no shutdown` enables the interface.
- `speed` and `duplex` are both normally set to `auto` (the default). Using `speed 100` and `duplex half` would apply a static configuration.
- `switchport` configures switching mode characteristics. Interfaces connected to computer devices are usually set to `switchport mode access`. `switchport port-security` allows configuration of various security mechanisms.

Once done, run `exit`. To make changes persistent, run `do copy running-config startup-config`.

```
Press RETURN to get started!

Switch>enable
Switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#interface GigabitEthernet0/1
Switch(config-if)#no shutdown
Switch(config-if)#switchport mode access
Switch(config-if)#exit
Switch(config)#do copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
Switch(config)#do sh int G0/1
GigabitEthernet0/1 is up, line protocol is up (connected)
  Hardware is Lance, address is 0001.6487.9619 (bia 0001.6487.9619)
  BW 1000000 Kbit, DLY 1000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s
```

Cisco IOS switch interface configuration commands. (Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)



`copy` is a privileged mode command. `do copy` allows you to run the command from within config mode.



You can use the `range` command to configure a number of interfaces simultaneously. For example, `interface range GigabitEthernet0/1-24` enters configuration mode for all 24 interfaces in module 0.

Lesson 3.3

Switch Port Configuration



Exam Objectives Covered

2.2 Given a scenario, configure switching technologies and features.

5.3 Given a scenario, troubleshoot common issues with network services.

Switch ports have a number of additional configuration parameters that solve common network infrastructure problems, such as making more bandwidth available, supplying power to devices, and building a loop free mesh or partial mesh topology.

As you study this lesson, answer the following questions:


- What are the benefits of link aggregation?
- Why would you consider configuring switches in your network to handle jumbo frames?
- What are the options and benefits of integrating PoE devices in your network?
- What protocol can be used to prevent looping and broadcast storms in your network, and how does it work?

Link Aggregation and NIC Teaming

Link aggregation means combining two or more separate cabled links into a single logical channel. For example, a single network adapter and cable segment might support 1 Gbps; bonding this with another adapter and cable segment gives a link of 2 Gbps. Link aggregation can also be used in an uplink between two switches or between a switch and a router or between two routers.



*From the host end, this can also be called **NIC teaming**; at the switch end, it can be called port aggregation and is referred to by Cisco as an EtherChannel. The term “bonding” is also widely substituted for “aggregation.”*



A server node uses NIC teaming to create a 4 Gbps channel link from four 1 Gbps ports to a workgroup switch, while the workgroup switch bonds its uplink transceivers to create a 20 Gbps channel to a router.

Link aggregation can also provide redundancy; if one link is broken, the connection is still maintained by the other. It is also often cost-effective; a four-port Gigabit Ethernet card might not match the bandwidth of a 10 GbE port but will cost less.



This configuration is fully redundant only if the business function does not depend on the full speed of the bonded link. If one port fails, and the link drops to 1 Gbps, but that bandwidth is insufficient, there is not full redundancy.

Link aggregation is typically implemented using the IEEE 802.3ad/802.1ax standard. 802.3ad bonded interfaces are described as a link aggregation group (LAG). 802.3ad also defines the **Link Aggregation Control Protocol (LACP)**, which can be used to detect configuration errors and recover from the failure of one of the physical links.

On a Cisco switch, the following commands configure LACP to group the first four Gigabit interfaces into a single channel with the ID 2:

```
interface range GigabitEthernet0/1-4
channel-group 2 mode active
```

The following commands configure the 10G interfaces into a channel group with ID 1. In this example, the 10G interfaces are on a different module than the Gigabit interfaces:

```
interface range 10GigabitEthernet1/1-2
channel-group 1 mode passive
```

On the router/layer 3 switch, the channel should be set to active:

```
interface range 10GigabitEthernet0/1-2  
channel-group 1 mode active
```



Optionally, both sides can be configured as active. However, if both sides are set to passive, no bonded channel will be created. The channel ID on each side does not have to match, but it is easier to manage the connection if it is the same on both switches.

Maximum Transmission Unit

A standard Ethernet frame has a maximum length of 1,518 bytes, excluding the preamble. Each frame has an 18-byte header:

- 6-byte destination and source MAC address fields.
- 4-byte error checking field.
- 2-byte EtherType field.

The maximum size of the data payload is 1,500 bytes. This upper limit of the payload is also referred to as the **maximum transmission unit (MTU)**.

In circumstances where data payloads can be very high, a 1500-byte MTU means using a lot of frames. A **jumbo frame** is one that supports a data payload of up to 9,216 bytes. This reduces the number of frames that need to be transmitted, which can reduce the amount of processing that switches and routers need to do. It also reduces the bandwidth requirement somewhat, as fewer frame headers are being transmitted. The benefits of jumbo frames are somewhat disputed, however.

When implementing jumbo frames, it is critical that all hosts and appliances (switches and routers) along the communications path be able and configured to support them. It is also vital to ensure that each device supports the same MTU. Also, it can be complex to calculate the MTU if any additional headers are used (for IPSec, for instance).


Jumbo frame support can be configured using the command `mtu 9018`, where 9,018 is the required size. On some appliances, this must be configured for the whole system; on others, it can be configured on a per-interface basis.

Spanning Tree Protocol

Large networks make use of multiple switches configured in a mesh or partial mesh topology to implement redundant links. Multiple paths are part of good network design as they increase resilience; if one link fails, then the network can remain operational by forwarding frames over a different path. However, Ethernet has no concept of a “time to live” value for frames, so layer 2 broadcast and flooded traffic could continue to loop through a network with multiple paths indefinitely.

The **Spanning Tree Protocol (STP)** is a means for the bridges or switches to organize themselves into a hierarchy and block loops. The switch at the top of the hierarchy is the root. The switch with the lowest ID, comprising a priority value and the MAC address, will be selected as the root.

Each switch then determines the shortest path to the root bridge by exchanging information with other switches. This STP information is packaged as bridge protocol data unit (BPDU) multicast frames. Different **port roles** are assigned to the interfaces participating in the spanning tree. A port that forwards “up” to the root, possibly via intermediate switches, is identified as a root port. Ports that can forward traffic “down” through the network with the least cost are identified as designated ports. A port that would create a loop is identified as a blocking or non-designated port. Subsequently, bridges exchange Topology Change Notifications if devices are added or removed, enabling them to change the status of forwarding/blocked ports appropriately.



Spanning tree configuration. (Images © 123RF.com.)

This image shows the minimum configuration necessary to prevent loops in a network with three bridges or switches. The root bridge has two designated ports (DP) connected to Bridge A and Bridge B. Bridges A and B both have root ports (RP) connected back to the interfaces on the root bridge. Bridges A and B also have a connection directly to one another. On Bridge A, this interface is active and traffic for Bridge B can be forwarded directly over it. On Bridge B, the interface is blocked (BP) to prevent a loop, and traffic for Bridge A must be forwarded via the root bridge.

Spanning Tree Protocol Configuration

If a switch supports spanning tree, it should operate by default without configuration. An administrator can (and should) set the priority value to predetermine **root bridge selection**. The root will usually be part of a high-bandwidth backbone or core switch group; performance will suffer if a switch on a low-bandwidth segment becomes root. You can use the `show spanning-tree` command to report the current configuration. Using `spanning-tree id root primary` and `spanning-tree id root secondary` assign main and backup priority values to the chosen switches.

```

RootBridge#show spanning-tree
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    24577
             Address     00D0.FF49.D04D
             This bridge is the root
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24577  (priority 24576 sys-id-ext 1)
             Address     00D0.FF49.D04D
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 20

Interface          Role Sts Cost          Prio.Nbr Type
-----
Gi1/0/23           Desg FWD 4           128.23  P2p
Gi1/0/22           Desg FWD 4           128.22  P2p

```

Viewing spanning tree configuration on a Cisco switch. This switch has been designated the root.
(Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)

The following table shows the different **port states** associated with spanning tree operation.

State	Forwards Frames?	Learns MACs?	Notes
Blocking	No	No	The port drops all frames other than BPDUs.
Listening	No	No	The port is listening for BPDUs to detect loops.
Learning	No	Yes	The port discovers the topology of the network and builds the MAC address table.
Forwarding	Yes	Yes	The port works as normal.
Disabled	No	No	The port has been disabled by the administrator.


```
BridgeB#show spanning-tree
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority    24577
            Address    00D0.FF49.D04D
            Cost      4
            Port      23(GigabitEthernet1/0/23)
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Bridge ID    Priority    32769 (priority 32768 sys-id-ext 1)
            Address    00D0.589B.6401
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
            Aging Time 20
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Gil/0/21	Altn	BLK	4	128.21	P2p
Gil/0/23	Root	FWD	4	128.23	P2p

*One of this switch's interfaces would make a loop, so it has been put in the blocking state.
(Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)*

When all ports on all switches are in forwarding or blocking states, the network is converged. When the network is not converged, no communications can take place. Under the original 802.1D standard, this made the network unavailable for extended periods—tens of seconds—during configuration changes. STP is now more likely to be implemented as 802.1D-2004/802.1w or Rapid STP (RSTP). The rapid version creates outages of a few seconds or less. In RSTP, the blocking, listening, and disabled states are aggregated into a discarding state.

Power Over Ethernet

Power over Ethernet (PoE) is a means of supplying electrical power from a switch port over ordinary data cabling to a connected powered device (PD), such as a VoIP handset, IP camera, or wireless access point. Powering these devices through a switch is more efficient than using a wall-socket AC adapter for each appliance. It also allows network management software to control the devices and apply schemes, such as making unused devices go into sleep states and power capping.

PoE is defined in two IEEE standards (now both rolled into 802.3-2018):

- **802.3af**—Power is supplied as 350mA@48V and limited to 15.4 W output. Given that some of this dissipates over the length of cable, it supports PDs that require up to about 13 W.
- **802.3at (PoE+)**—Supplies at 30 W, with a maximum current of 600 mA. This can support PD requirements of up to about 25 W.
- **802.3bt (PoE++)**—Supplies at 60 W (Type 3) or 90 W (Type 4), with up to 51 W and 71 W usable power, respectively.

PoE switches are referred to as endspan (or endpoint) power sourcing equipment (PSE). On a Cisco switch, the command `power inline auto max 15000` enables a port for PoE and sets a maximum output of 15,000 mW (or 15 W).

When a device is connected to a port on a PoE switch, the switch goes through a detection phase to determine whether the device is PoE enabled. If not, it does not supply power over the port and, therefore, does not damage non-PoE devices. If so, it determines the device's power consumption and sets the supply voltage level appropriately.



If a switch does not support PoE, a device called a power injector (or midspan) can be

Lesson 3.4

Switch Troubleshooting



Exam Objectives Covered

- 5.2 Given a scenario, troubleshoot common cabling and physical interface issues.
- 5.3 Given a scenario, troubleshoot common issues with network services.
- 5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

Ethernet switches and network adapters introduce the potential for issues at the Data Link layer and can reveal subtle cabling problems and interference at the Physical layer. Diagnosing and resolving problems gets more complex as you work up through the network stack. You need to assimilate your knowledge of both cabling types and Ethernet framing with awareness of status indicators and commands for network equipment to resolve these issues.

As you study this lesson, answer the following questions:

- How can you use the physical and logical topology to isolate a problem to a particular area of the network?
- What information can you obtain from network device status lights?
- Given symptoms of a particular problem, what commands should you use to gather information about an Ethernet switch's configuration?
- What are the symptoms and causes of a network loop?

Hardware Failure Issues

When you are using the CompTIA Network+ troubleshooting model, it is wise to rule out physical hardware failure and Data Link layer issues before diagnosing a Network layer or application issue.

Power Issues

Like any computer system, networks require stable power to operate properly. Power anomalies, such as surges and spikes, can damage devices, under-voltage events (very brief power loss) can cause systems to lock up or reboot, while power failures will down everything, including the lights. Enterprise sites have systems to protect against these issues. Uninterruptible power supplies (UPSs) can keep servers, switches, and routers running for a few minutes. This provides time to either switch in a secondary power source (a generator) or shut down the system gracefully, hopefully avoiding data loss. Most power problems will have to be escalated to an electrician or the power company, depending on where the fault lies.

Hardware Failure Issues

If power is not the issue, consider other components that might have experienced hardware failure, including host network adapters, switch/router/modem appliances, and the cabling between them. Complete hardware failure is relatively uncommon, so if you can rule out power and cabling problems, then for a network adapter, verify that the driver is working correctly. The easiest thing to do is to replace the driver (in Windows, use Device Manager to do this). For a network appliance, use status LEDs to confirm operation and check that things such as plug-in cards and modules are seated correctly. You should also consider overheating as a potential cause of hardware issues. Make sure there is good airflow around the intake and outlet vents. Check that fans and internal components are not clogged with dust and that systems are not exposed to direct sunlight.

At the Data Link layer, most wired hosts connect to the network via a switch. If you suspect a device such as a switch, analyze the topology of your network. You should be able to view those users who are suffering from the problem, identify which part of the network is affected, and identify the problem bridging or switching device.

When you have narrowed the problem to a device, you must determine what the nature of the problem is. It is always worth resetting the switch to see if that resolves the problem. Often, restarting network devices can clear transitory errors.



Do be aware that restarting a switch, router, or server can be very disruptive to the rest of the network. Identify how to mitigate potential impacts and seek authorization for your plan before proceeding. Also, remember that a restart will apply the startup configuration. Any unsaved changes in the running configuration will be discarded.

Port Status Indicators

When you are troubleshooting a suspected layer 1 or layer 2 problem, check the LED status indicators on the NIC at one end and the switch/router port at the other. You will need the vendor documentation to interpret the LEDs. There may be two LEDs for status and for activity, or the LED might use a mode button to show different information. On a switch port, the following LED link states are typical:

- **Solid green**—The link is connected, but there is no traffic.
- **Flickering green**—The link is operating normally (with traffic). The blink rate indicates the link speed.
- **No light**—The link is not working, or the port is shut down.
- **Blinking amber**—A fault has been detected (duplex mismatch, excessive collisions, or redundancy check errors, for instance).
- **Solid amber**—The port is blocked by the spanning tree algorithm, which works to prevent loops within a switched network.

Switch Show Commands

If you can isolate the issue to a single host and then rule out cable, transceiver, and bad port issues at the Physical layer, bear in mind that the Data Link configuration might not be working.

In privileged mode, a variety of **show commands** can be used to display the current configuration of a switch. There are usually many `show` commands, but two of particular importance are as follows:

- `show config` displays the switch's configuration. The startup configuration (`show startup-config`) could be different from the running configuration (`show running-config`). If there has been some undocumented change to the switch, using these commands and comparing the output may reveal the source of a problem.
- `show interface` lists the state of all interfaces or the specified interface. An interface has a line status (up if a host is connected via a good cable) and a protocol status (up if an Ethernet link is established). `show interface` will also report configuration details and traffic statistics if the link is up/up.

```
Switch#show int g0/1
GigabitEthernet0/1 is up, line protocol is up (connected)
  Hardware is Lance, address is 0001.6487.9619 (bia 0001.6487.9619)
  BW 1000000 Kbit, DLY 1000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:08, output 00:00:05, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue :0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    956 packets input, 193351 bytes, 0 no buffer
    Received 956 broadcasts, 0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    0 watchdog, 0 multicast, 0 pause input
    0 input packets with dribble condition detected
    2357 packets output, 263570 bytes, 0 underruns
    0 output errors, 0 collisions, 10 interface resets
    0 babbles, 0 late collision, 0 deferred
    0 lost carrier, 0 no carrier
    0 output buffer failures, 0 output buffers swapped out
```

*Viewing interface configuration on a Cisco switch.
(Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)*

If an interface is not up/up, you need to diagnose the cause from the state:

- **Down/down**—There is no link. This is typically because no host is attached, but it could also be caused by a speed mismatch.
- **Administratively down/down**—The interface has been disabled using the `shutdown` command. Use `no shutdown` to bring it up.

- **Down/error disabled**—The interface is disabled due to some error state. This is typically either a spanning tree or port security violation issue.
- **Up/down (suspended)**—The port is part of a link aggregation group, and the channel has not been negotiated. Use `show etherchannel` to investigate the cause. Both sides should use the same speed, duplex, and link control type, and use the same number of ports. When using LACP, at least one side must be active.

If the line and protocol status is down/down, check whether autonegotiation of speed and duplex settings is configured and whether it is failing. In most cases, this will be because either the adapter or the switch port has been manually configured. If a host is set to a fixed configuration and the switch is set to auto-negotiate, the switch will default to 10 Mbps/half-duplex because the host will not negotiate with it! So, if the host is manually configured to 100 Mbps/full-duplex, the link will fail. Setting both to auto-negotiate will generally solve the problem. A speed mismatch will cause the link to fail, while a duplex mismatch will slow the link down (it will cause high packet loss and late collisions).

Interface Error Counters

Interface status commands will also report whether any collisions are being generated. Collisions might occur if the duplex setting on the switch port and host is mismatched or if a legacy hub device or host NIC is connected to a switch. Other types of interface errors might indicate a misconfiguration problem at the Data Link layer or interference at the Physical layer.

Increasing Interface Counters

An interface might change rapidly or “flap” between up and down states, making the problem harder to observe and diagnose. Interface counters record the number of events over time. This allows you to diagnose issues with an interface that is up but that is unreliable or performing poorly.

- **Link state**—Measures whether an interface is working (up) or not (down). You would configure an alert if an interface goes down so that it can be investigated immediately. You may also want to track the uptime or downtime percentage so that you can assess a link’s reliability over time.
- **Resets**—The number of times an interface has restarted over the counter period. Interfaces may be reset manually or could restart automatically if traffic volume is very high or a large number of errors are experienced. Anything but occasional resets should be closely monitored and investigated.
- **Discards/drops**—An interface may discard incoming and/or outgoing frames for several reasons, including checksum errors, mismatched MTUs, packets that are too small (runt) or too large (giants), high load, or permissions—the sender is not on the interface’s access control list (ACL) or there is some sort of virtual LAN (VLAN) configuration problem, for instance. Each interface is likely to class the type of discard or drop separately to assist with troubleshooting the precise ca

Cyclic Redundancy Check Errors

A cyclic redundancy check (CRC) is calculated by an interface when it sends a frame. A CRC value is calculated from the frame contents to derive a 32-bit value. This is added to the header as the frame check sequence. The receiving interface uses the same calculation. If it derives a different value, the frame is rejected. The number of CRC errors can be monitored per interface.

CRC errors are usually caused by interference. This interference might be due to poor quality cable or termination, attenuation, mismatches between optical transceivers or cable types, or some external factor.

Runt Frame Errors

A **runt** is a frame that is smaller than the minimum size (64 bytes for Ethernet). A runt frame is usually caused by a collision. In a switched environment, collisions should only be experienced on an interface connected to a legacy hub device and there is a duplex mismatch in the interface configuration (or possibly on a misconfigured link to a virtualization platform). If runts are generated in other conditions, suspect a driver issue on the transmitting host.

Giant Frame Errors

A **giant** is a frame that is larger than the maximum permissible size for the receiving interface. There are two likely causes of giant frames:

- **Jumbo frames**—A host might be configured to use jumbo frames, but the switch interface is not configured to receive them. This type of issue often occurs when configuring storage area networks (SANs) or links between SANs and data networks. The MTU value in the `show interface` output will indicate whether jumbo frames are accepted on a particular port.
- **Ethernet trunks**—A trunk link carries traffic between switches or between a switch and a router. Trunk links often use 802.1Q framing to carry virtual LAN (VLAN) information. If one switch interface is configured for 802.1Q framing, but the other is not, the frames will appear too large to the receiver, as 802.1Q adds 4 bytes to the header, making the maximum frame size 1522 bytes.



An Ethernet frame that is slightly larger (up to 1600 bytes) is often referred to as a baby giant.

MAC Address Table

A switch learns MAC addresses by reading the source address when a frame is received on a port. The address mapping for that port is cached in a MAC address table. The address table is implemented as content addressable memory (CAM), a special type of memory optimized for searching, rather than random access. Consequently, the MAC address table is often also referred to as the CAM table. Entries remain in the MAC address table for a period of time before being flushed. This ensures problems are not encountered when network cards (MAC addresses) are changed.

If a MAC address cannot be found in the MAC address table, then the switch acts like a hub and transmits the frame out of all the ports, except for the source port. This is referred to as flooding.

Knowing the MAC addresses associated with a particular interface is often important for troubleshooting. You can query the MAC address table of a switch to find the MAC address or addresses associated with a particular port using a command such as:

```
show mac address-table
```

```
Switch#show mac address-table
```

```
Mac Address Table
```

Vlan	Mac Address	Type	Ports
----	-----	-----	-----
1	0090.0cd5.1001	DYNAMIC	Fa0/24
10	0000.0c11.d48e	DYNAMIC	Fa0/3
10	0090.0cd5.1001	DYNAMIC	Fa0/24
10	0090.2105.8582	DYNAMIC	Fa0/2
10	00d0.d3db.9ca9	DYNAMIC	Fa0/1
20	0005.5e32.c45b	DYNAMIC	Fa0/2
20	000c.cf54.66d4	DYNAMIC	Fa0/1
20	0060.5c30.9777	DYNAMIC	Fa0/3
20	0090.0cd5.1001	DYNAMIC	Fa0/24

```
Switch#
```

*Displaying dynamic entries in the MAC address table of a Cisco switch.
(Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)*

Network Loop and Broadcast Storm Issues

A **network loop** is where flooded frames circulate the network perpetually. Because switches flood broadcasts out all ports, these frames will go down one link to the next switch, which will send the broadcast back up the redundant link, and back to the originating switch. As this repeats, the switches start to see source MAC addresses associated with multiple ports and so clear the MAC address table mapping, which causes them to start flooding unicast traffic too.

Without intervention, this loop will continue indefinitely, causing a **broadcast storm**. A broadcast storm will cause network utilization to go to near maximum capacity and the CPU utilization of the switches to jump to 80% or more. This makes the switched segment effectively unusable until the broadcast storm stops. A broadcast storm may quickly consume all link bandwidth and crash network appliances.

If there is a loop, spanning tree should shut down the port. This will isolate the problem to a segment of the network. Inspect physical ports that correspond to the disabled interfaces for looped connections. At the patch panel, this could mean a patch cable that connects two ports on the same switch. On the office floor, it could mean a patch cable between two wall ports. Check the switch for log events related to MAC address flapping.

If a broadcast storm occurs on a network where spanning tree is already enabled, you should investigate the following potential causes:

- Verify compatible versions of Spanning Tree Protocol or Rapid Spanning Tree Protocol are enabled on all switches.
- Verify the physical configuration of segments that use legacy equipment, such as Ethernet hubs.
- Investigate networking devices in the user environment and verify that they are not connected as part of a loop. Typical sources of problems include unmanaged desktop switches and VoIP handsets.

Power Over Ethernet Issues

Power over Ethernet (PoE) uses data cabling to run lightweight Powered Device (PD) appliances, such as Voice over IP (VoIP) handsets, IP cameras, and wireless access points.

Cabling Issues

Cabling for PoE+ must be Cat 5e or better, but standards typically recommend the use of Cat 6A. Drawing power down the cable generates more heat. If this heat is not dissipated, it can affect data rates. Thermal performance is improved by using pure copper cabling with thicker conductors. A thin conductor will generate more heat through resistance. Shielded cabling is capable of dispersing heat more efficiently.



Conductor thickness is measured as American Wire Gauge (AWG). Remember that smaller numbers mean thicker wires, so 23 AWG cable will have superior PoE performance to 24 AWG cable.

Incorrect Standard

A PD should be able to negotiate the correct mode and power output with the switch. However, this process can fail with some devices that only support the first PoE standard, especially if the switch interface is enabled for high power PoE++ Type 4 PDs. The switch and PD must negotiate a compatible mode:

- Alternative A delivers power with data over pairs 1/2 and 3/6. This is compatible with 10/100 and 10/100/1000 links.
- Alternative B delivers power over the 10/100 spare pairs (4/5 and 7/8). This is not compatible with Gigabit Ethernet.
- Four-pair delivers power over all pairs. This is required by PoE++ Type 3 and Type 4 PDs. This is compatible with 10/100/1000 and also supports 10G.

Power Budget Exceeded

Each switch has a total **power budget** for all ports. This will typically be around 300–400 watts (30,000–40,000 milliwatts). If the power requirements of all connected devices exceed the budget, some will not be activated, or there might be intermittent resets. You can use the `show power inline` command to report the power budget and power consumption. If the power budget is exceeded, you will typically need to provision another switch, though it is also possible to use power injector devices to remove the load of selected PDs from the switch.

```
Switch#show power inline
```

```
Available:370.0(w)  Used:30.0(w)  Remaining:340.0(w)
```

Interface	Admin	Oper	Power (Watts)	Device	Class	Max
Fa0/1	auto	on	10.0	Switch 7960	3	15.4
Fa0/2	auto	on	10.0	Switch 7960	3	15.4
Fa0/3	auto	on	10.0	Switch 7960	3	15.4
Fa0/4	auto	off	0.0	n/a	n/a	15.4
Fa0/5	auto	off	0.0	n/a	n/a	15.4

Summary of PoE-enabled ports. The “switch” devices listed here are actually Voice over IP (VoIP) handsets. (Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)



Actual power consumption can fluctuate quite widely. For example, a camera with pan-tilt-zoom controls will use more power when its motor is active.

Module 3

Summary

You should be able to identify the features of network devices operating at layers 1 and 2 and determine their appropriate placement on the network. You should be able to deploy and troubleshoot Ethernet switches with appropriate port configurations.

Guidelines for Deploying Switched Networks

Follow these guidelines to deploy switches and other networking devices:

- Identify a switch model that will meet current and future needs, taking into account port density, management features, and any requirement to use Power over Ethernet.
- Configure appropriate port settings for high-bandwidth hosts, such as link aggregation and jumbo frames.
- Identify any hosts that require nonstandard configuration, such as disabling autonegotiation of speed and duplex settings.
- Create a management plan for legacy hub and bridge appliances to ensure they do not impact overall network performance.
- Enable spanning tree to prevent loops around redundant circuits and ensure the selection of an appropriate root bridge.
- Use status indicators and networking device commands to verify system and interface configurations.

Module 4

Configuring Network Addressing

Module Introduction

The Physical and Data Link layers covered in the previous course modules establish local links between nodes. At the Network layer—layer 3—these individual networks can be connected together into a network of networks, or internetwork.

In this module, you will identify the addressing and data delivery methods of the Internet Protocol (IP). IP is at the heart of most modern networks and consequently one of the most important topic areas for a network professional to understand and apply.

IP is implemented on network hosts using a wide variety of configuration interfaces and tools. You must be confident about selecting an appropriate tool to use to complete a particular support or troubleshooting task.

This module also introduces IPv6 addressing concepts and highlights some key differences between IPv6 and IPv4.

Module Objectives

In this module, you will do the following:

- Explain IPv4 addressing schemes.
- Explain IPv4 forwarding.
- Configure IP networks and subnets.
- Use appropriate tools to test a host's IP configuration.
- Explain IPv6 addressing schemes.
- Troubleshoot IP networks and hosts.

Lesson 4.1

Internet Protocol Basics



Exam Objectives Covered

1.4 Explain common networking ports, protocols, services, and traffic types.

The Transmission Control Protocol/Internet Protocol (TCP/IP) suite consists of complementary protocols and standards that work together to provide the functionality of the vast majority of modern networks. The Internet Protocol (IP) stands at the heart of this protocol suite, providing logical addressing and packet forwarding between different networks. In this topic, you will start to investigate the characteristics of IP by examining the structure of IPv4 packets, the format of IPv4 addresses, the differences between unicast, broadcast, multicast, and anycast addressing schemes.

As you study this lesson, answer the following questions:

- What headers does IPv4 use for addressing and forwarding?
- What is the purpose of using logical network addressing at layer 3?
- How does addressing between layer 2 and layer 3 work?
- How does a host address multiple destination hosts efficiently?




There are two versions of IP; version 4 is more widely adopted and is the version discussed in the following few lessons. IPv6 introduces a much larger address space and different means of configuring clients and is discussed later in the module.

IPv4 Datagram Header

Ethernet works at the Physical and Data Link layers of the OSI model (layers 1 and 2). Ethernet, and other layer 1/layer 2 products, have no concept of multiple networks or of logical subdivisions within a network. This function is implemented at the Network layer (layer 3). As a layer 3 protocol, the Internet Protocol (IP) provides logical network addressing and forwarding.

The **Internet Protocol (IP) header** contains fields to manage the logical addressing and forwarding function. In IPv4, the header contains two fields for the 32-bit source and destination addresses, plus a number of other fields to support forwarding functions.



IPv4 header.

The Version field indicates the version of Internet Protocol in use (4), while the Length fields indicate the size of the header and the total packet size (including the payload). The maximum theoretical size is 65,535 bytes, but actual packets would typically be much smaller to avoid fragmentation when transported as the payload of Ethernet frames, which usually have 1,500-byte MTUs.

The Protocol field describes what type of data is encapsulated in the payload so that the receiving host knows how to process it. For most packets, the **IP protocol type** value in the Protocol field will indicate a Transmission Control Protocol (TCP/6) segment or a User Datagram Protocol (UDP/17) datagram, which work at the Transport layer. The values assigned to protocol types, such as 6 for TCP and 17 for UDP, are managed by IANA.



Those are the values in decimal. You are also likely to see them in their hex forms (0x06 and 0x11). Both formats ultimately represent 8-bit binary values (00000110 and 00010001).


Some Network layer protocols run directly on IP. These IP protocol types include the following:

- Internet Control Message Protocol (ICMP/1) is used for status messaging and connectivity testing.
- Internet Group Messaging Protocol (IGMP/2) is used with multicasting.
- Generic Routing Encapsulation (GRE/47) is used to tunnel packets across an intermediate network. This is used (for example) in some virtual private network (VPN) implementations.
- Encapsulating Security Payload (ESP/50) and Authentication Header (AH/51) are used with the encrypted form of IP (IPSec).
- Enhanced Interior Gateway Routing Protocol (EIGRP/88) and Open Shortest Path First (OSPF/89) are protocols used by routers to exchange information about paths to remote networks.

Layer 2 Versus Layer 3 Addressing and Forwarding

When designing or supporting an IP network, it is important to understand how the logical addressing scheme at the Network layer maps to forwarding at the Data Link layer. Forwarding at layer 3 is referred to as routing, while forwarding at layer 2 is described as switching.

The following figure illustrates how both switching and routing components might be used in a typical network. The whole network is connected to the wider Internet via the WAN interface on the router. The router's other interfaces are used to divide the network into three logical subnetworks. These subnets are mapped to layer 2 segments, each implemented using a switch.



Network placement of routers and switches. (Images © 123RF.com.)

Nodes within each subnet can address one another directly (they are in the same broadcast domain), but they can only communicate with nodes in other subnets via the router.


Within each subnet, nodes use Media Access Control (MAC) addresses to forward frames to one another, using a mechanism to translate between layer 3 IP addresses and layer 2 MAC addresses.

The Network layer can also accommodate forwarding between different types of layer 1/layer 2 networks. The private zone is implemented using Ethernet, but the link between the router's public interface and the ISP might use a different technology, such as digital subscriber line (DSL).

Address Resolution Protocol

When two hosts communicate over an Ethernet network using IP, an IP address identifies each Network layer host interface. However, transmission of data must take place at the Physical and Data Link levels using the local or hardware/MAC address of the interface. The TCP/IP suite includes the **Address Resolution Protocol (ARP)** to perform the task of resolving an IPv4 address to a hardware MAC address.

When both sending and receiving hosts are within the same broadcast domain or subnet, local address resolution takes place using ARP requests and ARP replies, as shown in the figure:



ARP requests and replies. (Images © 123RF.com.)


If the destination address is on a different subnet or on a remote network, then the local host must use its default gateway to forward the packet. Therefore, it must determine the MAC address of the default gateway using ARP.

The router also uses ARP messaging for its Ethernet interfaces. ARP messaging is only used with Ethernet, however. A router's public interface might use a different type of framing and local addressing.

Unicast and Broadcast Addressing

When an IP host wants to send a packet to a single recipient, it uses a unicast packet, addressed to the IP address of the destination host. If, however, the local host needs to communicate with multiple hosts, it can use a different scheme.

One means of addressing multiple hosts is to perform a broadcast. A broadcast can be performed by sending a packet to the network or subnet's broadcast address. The broadcast address is the last address in any IP network or subnet.




Unicast and broadcast addressing.

All hosts that share the same broadcast address receive the packet. They are said to be in the same layer 3 broadcast domain. Broadcast domain boundaries are established at the Network layer by routers. Routers do not forward broadcasts, except in some specially configured circumstances.

As with unicast traffic, IP packets must be delivered to hosts using layer 2 MAC addresses. At layer 2, broadcasts are delivered using the group MAC address (ff:ff:ff:ff:ff:ff). This means that there is also a broadcast domain scope at layer 2. With legacy devices such as hubs and bridges, every port on all physically connected nodes is part of the same layer 2 broadcast domain. This is also the case with a basic or unmanaged switch. By default, a switch floods broadcasts out of every port except the source port.

Multicast and Anycast Addressing

While the majority of IPv4 traffic is unicast or broadcast, other addressing schemes are used in special circumstances.



Multicast and anycast addressing.

Multicast Addressing

IP multicasting allows one host on the Internet (or private IP network) to send content to other hosts that have identified themselves as interested in receiving the originating host's content. Multicast packets are sent to a destination IP address from a special range configured for use with that multicast group. In IPv4, the range 224.0.0.0 through to 239.255.255.255 is reserved for multicast addressing.

The intent to receive multicasts from a particular host is signaled by joining a multicast group. The **Internet Group Management Protocol (IGMP)** is typically used to configure group memberships and IP addresses.

At layer 2, multicasts are delivered using a special MAC address, comprised of the prefix 01-00-53, with the remainder expressing the multicast group IP address in hex notation. To deliver this frame only to members of the multicast group, the switch must be capable of IGMP snooping. If the switch is not multicast capable, it will treat the multicast MAC address like a broadcast and flood the multicast transmissions out of all ports.

Anycast Addressing

Anycast means that a group of hosts is configured with the same IP address. When a router forwards a packet to an anycast group, it uses a prioritization algorithm and metrics to select the host that is “closest” (that will receive the packet and be able to process it the most quickly). This allows the service behind the IP address to be provisioned more quickly and reliably. It allows for load balancing and failover between the server hosts sharing the IP address.



There isn't an anycast address range. Hosts participating in an anycast group are configured with the same unicast address. Anycast forwarding is handled by routers, typically using a dynamic routing protocol, such as Border Gateway Protocol (BGP).

Lesson 4.2

IP Version 4 Addressing



Exam Objectives Covered

1.7 Given a scenario, use appropriate IPv4 network addressing.

The core function of IP is to facilitate the creation of a group of logically distinct but interconnected networks, referred to as an internetwork. This means that some packets addressed to hosts on remote networks must be forwarded via one or more of the intermediate systems that establish paths between networks.

In this topic, you will identify the basic principles by which IPv4 distinguishes local and remote hosts and networks.

As you study this lesson, answer the following questions:

- What is the format of an IPv4 address?
- How can I convert between binary and decimal representations of IPv4 addresses?
- What is the purpose of a network mask, and what format can these masks take in IPv4?
- What is the purpose of subnetting?
- What formula is used to calculate the number of hosts per subnet?

IPv4 Address Format

Networks in an internetwork must have a way of uniquely identifying each logical network and each individual host within those networks. At the Data Link layer, an interface is identified by using a MAC or hardware address. This type of address can be used only for local delivery of frames. At the Network layer, IP source and destination addresses are used to forward packets to the proper destination. An IP address provides two pieces of information:

- **The network number (Network ID)**—This number is common to all hosts on the same IP network.
- **The host number (Host ID)**—This number identifies a host within an IP network.

You can use the same columnar method to convert from decimal to binary. For example, the number 51 can be converted as follows:

	2 ⁷	2 ⁶	2 ⁵	2 ⁴	2 ³	2 ²	2 ¹	2 ⁰
	128	64	32	16	8	4	2	1
51 =	0	+0	+32	+16	+0	+0	+2	+1
	0	0	1	1	0	0	1	1

Decimal to binary conversion.

If all the bits in an octet are set to 1, the number obtained is 255 (the maximum possible value). Similarly, if all the bits are set to 0, the number obtained is 0 (the minimum possible value). Therefore, theoretically an IPv4 address may be any value between 0 . 0 . 0 . 0 and 255 . 255 . 255 . 255. However, some addresses are not permitted or are reserved for special use.

Network Masks

An IP address represents both a network ID and a host ID. In IPv4, a 32-bit **network mask** (or netmask) is used to distinguish these two components within a single IPv4 address. The mask conceals the host ID portion of the IP address and thereby reveals the network ID portion.

The mask and the IPv4 address are the same number of bits. Wherever there is a binary 1 in the mask, the corresponding binary digit in the IPv4 address is part of the network ID. The 1s in the mask are always contiguous. For example, this mask is valid:


```
11111111 11111111 11111111 00000000
```

But the following string is not a valid mask:

```
11111111 00000000 11111111 00000000
```

The network ID portion of an IP address is revealed by ANDing the mask to the IPv4 address. When two 1s are ANDed together, the result is a 1. Any other combination produces a 0.


For example, to determine the network ID of the host IPv4 address 198 . 51 . 100 . 1 with a mask of 255 . 255 . 255 . 0, the dotted decimal notation of the IP address and mask must first be converted to binary notation. The next step is to AND the two binary numbers. The result can be converted back to dotted decimal notation to show the network ID (198 . 51 . 100 . 0). The only difference between the host IP address and the network ID lies in the last octet, which is not masked.



ANDing IP address and subnet mask to derive the network ID.

Instead of the dotted decimal mask 255 . 255 . 255 . 0, this network can be identified using prefix or slash notation. The prefix is simply the number of bits set to 1 in the mask. The network can therefore be referred to as 198 . 51 . 100 . 0 / 24.

The default masks align with octet boundaries. This means that the values in the dotted decimal mask will be 255 or 0. For example, the default 24-bit mask is as follows:



Network ID and host ID portions when using a 24-bit mask.

An 8-bit mask is 255.0.0.0, and a 16-bit mask is 255.255.0.0.

A longer network portion, such as 255.255.255.0, allows for more network IDs within the overall internetwork but with fewer available host addresses per network. Each /24 network has only 254 possible host addresses. Conversely, the short netmask 255.0.0.0 allows for millions of hosts per network but only 126 possible network addresses.

Because the 1s in a mask are always contiguous, each octet in decimal in an IPv4 mask will always be one of the following.


Octet Mask Bits	Binary Octet	Decimal Equivalent
1	10000000	128
2	11000000	192
3	11100000	224
4	11110000	240
5	11111000	248
6	11111100	252
7	11111110	254
8	11111111	255

Try to memorize these values to make converting masks between binary and decimal formats easier. For example, if the mask has 14 bits, you can work out the octets as 8 bits plus 6 bits. Therefore, a /14 network has the following mask:


11111111 11111100 00000000 00000000 255 252 0 0

Host Address Ranges

The IP network 198 . 51 . 100 . 0 /24 allows for 254 possible host IDs. The host ID portion is 8 bits long:



Host address range for a /24 network.




Host address range for a /28 network.

The purpose of subnetting is to create layer 3 broadcast domain segments with fewer hosts. The trick with subnet design is to fit the scheme to the requirements for number of subnetworks and number of hosts per subnet. Each bit added to the mask approximately halves the number of available host addresses.

Default Gateway


When two end system hosts attempt to communicate via IPv4, the protocol compares the source and destination address in each packet against the netmask. If the masked portions of the source and destination IP addresses match, then the destination interface is assumed to be reachable via the local layer 2 network.



Matching source and destination network IDs.

In the figure, the first 28 bits of the source and destination address are the same. Therefore, IP concludes the destination IPv4 address is on the same IP network or subnet and tries to deliver the packet locally.

If the masked portion does not match, as in the following figure, IP assumes the packet must be routed to another IP network or subnet:



Different source and destination network IDs.

When the destination IPv4 address is on a different IP network or subnet, the host forwards the packet to its **default gateway**, rather than trying to deliver it locally. The default gateway is a router configured with a path to remote networks.

The router determines what to do with the packet by performing the same comparison between the source and destination address and netmask. The router then uses its routing table to determine which interface it should use to forward the packet. If no suitable path is available, the router drops the packet and informs the host that it could not be delivered.

If the message is destined for yet another network, the process is repeated to take it to the next stage, and so on.

Paths to other IP networks can be manually configured in the routing table or learned by a dynamic routing protocol. Dynamic routing protocols allow routers to share information about known networks and possible paths to them. This information allows them to choose the best routes to any given destination and select alternate routes if one of these is unavailable.



A default gateway router's interface IP can be any usable host ID, but by convention it is normally set to either the first or last usable host address.

Broadcast Addresses

A broadcast can be performed by sending a packet to the network or subnet's broadcast address. The broadcast address is the last address in any IP network, or put another way, the address in any IP network where all the host bits are set to 1.

For example, if the subnet mask is 255 . 255 . 255 . 240, the last four digits of the last octet in the IP address is the host ID portion. If these digits are set to all 1s, that is the last possible address before the next subnet ID, and therefore the network broadcast address:

11000110 198	00110011 51	01100100 100	0001	1111 31
11000110 198	00110011 51	01100100 100	0010	1111 47
11000110 198	00110011 51	01100100 100	0011	1111 63
11111111 255	11111111 255	11111111 255	1111	0000 240
Network ID			Subnet ID	Broadcast

Broadcast addresses for three subnets.

IP Interface Configuration in Windows

Each host adapter must be allocated an appropriate IP address and subnet mask, plus the IP address of the default gateway (router) for its network. Typically, a host is also configured with the addresses of domain name system (DNS) servers that can resolve IP address to names, making identification of hosts and services simpler.

These IP configuration values can be assigned statically or dynamically. Configuring large numbers of hosts with a valid static address is a complex management task. Most hosts are configured to obtain an address automatically, using a service called the Dynamic Host Configuration Protocol (DHCP).

Under Windows, each Ethernet adapter is assigned a name. In early Windows versions, the first adapter was named "Local Area Connection," but recent versions just use the label "Ethernet." Additional adapters are identified as "Ethernet2," "Ethernet3," and so on. A new name can be applied if necessary. The IP configuration for each adapter interface is often set using the GUI Properties dialog accessed via the Network Connections applet or Windows Settings app. However, you can also configure interfaces using `netsh` commands.

```
netsh interface ip set address "Ethernet" dhcp

netsh interface ip set address "Ethernet" static
198.51.16.17 255.255.255.252 198.51.100.30
```

The second example configures a host with the IP address `198.51.16.17/28` and identifies the default gateway as `198.51.100.30` (the last host address in the subnet). In Windows, all changes to the network interface configuration are persistent, meaning that they continue to apply when the system is rebooted.

You can also use `netsh` to report the IP configuration (`netsh interface ip show config`, for example).

`netsh` is implemented in the legacy command prompt interface. Script-based configuration is now more likely to use PowerShell cmdlets. The `Get-NetAdapter` and `Get-NetIPAddress` cmdlets can be used to query the existing configuration. A new configuration can be applied using `New-NetIPAddress`, or an existing one can be modified using `Set-NetIPAddress`.

IP Interface Configuration in Linux

In Linux, Ethernet interfaces are classically identified as `eth0`, `eth1`, `eth2`, and so on, although some network packages now use different schemes, such as `en` prefixes. In Linux, you need to distinguish between the running configuration and the persistent configuration. The persistent configuration is the one applied after a reboot or after a network adapter is reinitialized.

The method of applying an IP configuration to an adapter interface is specific to each distribution. Historically, the persistent configuration was applied by editing the `/etc/network/interfaces` file and bringing interfaces up or down with the `ifup` and `ifdown` scripts. Many distributions now use the NetworkManager package, which can be operated using a GUI or the `nmcli` tools. Alternatively, a network configuration might be managed using the `systemd-networkd` configuration manager. Additionally, recent distributions of Ubuntu use `netplan` to abstract some of this underlying complexity to configuration files written in **YAML Ain't Markup Language (YAML)**. The YAML configuration files are rendered by either `systemd-networkd` or NetworkManager.

Lesson 4.3

IP Version 4 Subnetting



Exam Objectives Covered

1.7 Given a scenario, use appropriate IPv4 network addressing.

Organizations with large networks need to divide those networks up into smaller segments to improve performance and security. A network segment is represented at the Network layer by a subnet. Additionally, IPv4 uses a system of public versus private addressing that determines how hosts and networks can connect over the Internet. Understanding the features of these addressing schemes will be essential to your career in network support.

- What is the original classful IPv4 addressing scheme, and how is it relevant to modern networks?
- Which address ranges are available for use on the public Internet, and which are reserved for private networks or other purposes?
- How can classless addressing summarize network addresses and allow for networks with different-sized subnets?

Classful Addressing

So far, we have considered IP network and subnet IDs that are defined by network masks. This is referred to as classless addressing. A **classful addressing** scheme was employed in the 1980s, before the use of netmasks to identify the network ID portion of an address was developed. Classful addressing allocates a network ID based on the first octet of the IP address.

	Number of Networks	Number of Hosts per Network	First Octet of Address Range
Class A 0??????? ???????? ???????? ???????? Network ID Host ID	126	16,777,214	1-126
Class B 10??????? ???????? ???????? ???????? Network ID Host ID	16,384	65,534	128-191
Class C 110????? ???????? ???????? ???????? Network ID Host ID	2,097,152	254	192-223

Class A, Class B, and Class C networks.

Class A network addresses support large numbers of hosts—over 16 million. However, there are only 126 Class A network addresses. There are 16,000 Class B networks, each containing up to about 65,000 hosts. Finally, Class C networks support only 254 hosts each, but there are over two million of them.

When considering classful addressing, you need to be able to identify the address class from the first octet of the IP address. This table shows how to identify an address class from the first octet of the IP address in decimal.

First Octet	Class
1–126	Class A
128–191	Class B
192–223	Class C



IP ranges 0.0.0.0/8 and 127.0.0.0/8 are also part of Class A but are reserved for special uses. 0.0.0.0/8 means “this” network, and 127.0.0.0/8 is used for loopback addressing.

While routers have performed classless addressing for years, the class terminology is still used in some contexts. Even under classless addressing, the old classes are often used as names for the netmasks that align to whole octet boundaries:

- Class A: 255.0.0.0 (/8)
- Class B: 255.255.0.0 (/16)
- Class C: 255.255.255.0 (/24)

Classful addressing is also important because it established some IP address ranges that cannot be used for ordinary host addressing or for addressing over the Internet.

Public Versus Private Addressing

A **public IP address** is one that can establish a connection with other public IP networks and hosts over the Internet. The allocation of public IP addresses is governed by IANA and administered by regional registries and ISPs. Hosts communicating with one another over a LAN could use a public addressing scheme but will more typically use private addressing.

Private IP addresses can be drawn from one of the pools of addresses defined in **RFC 1918** as non-routable over the Internet:

- 10.0.0.0 to 10.255.255.255 (Class A private address range).
- 172.16.0.0 to 172.31.255.255 (Class B private address range).
- 192.168.0.0 to 192.168.255.255 (Class C private address range).

Any organization can use private addresses on its networks without applying to a registry or ISP, and multiple organizations can use these ranges simultaneously. Internet access can be facilitated for hosts using a private addressing scheme in two ways:

- Through a router configured with a single valid public IP address or a block of valid public IP addresses; the router translates between the private and public addresses using a process called Network Address Translation (NAT).
- Through a proxy server that fulfills requests for Internet resources on behalf of clients. The proxy server itself must be configured with a public IP address on the external-facing interface.

Other Reserved Address Ranges

There are two additional classes of IP address (D and E) that use the values above 223.255.255.255:

- Class D addresses (224.0.0.0 through 239.255.255.255) are used for multicasting.
- Class E addresses (240.0.0.0 through 255.255.255.255) are reserved for experimental use and testing.

Loopback Addresses

While nominally part of Class A, the range 127.0.0.0 to 127.255.255.255 (or 127.0.0.0/8) is reserved. This range is used to configure a **loopback address**, which is a special address typically used to check that TCP/IP is correctly installed on the local host. The loopback interface does not require a physical interface to function. A packet sent to a loopback interface is not processed by a network adapter but is otherwise processed as normal by the host's TCP/IP stack. Every IP host is automatically configured with a default loopback address, typically 127.0.0.1. On some hosts, such as routers, more than one loopback address might be configured. Loopback interfaces can also be configured with an address from any suitable IP range, as long as it is unique on the network. A host will process a packet addressed to a loopback address regardless of the physical interface on which it is received.



Most hosts are also configured with a Domain Name System (DNS) host name. The loopback address is associated with the name `localhost`. The name `localhost` can be substituted for the numeric loopback address.

Other

A few other IPv4 address ranges are reserved for special use and are not publicly routable:

- **0.0.0.0/8**—Used when a specific address is unknown. This is typically used as a source address by a client seeking a DHCP lease.
- **255.255.255.255**—Used to broadcast to the local network when the local network address is not known.
- **100.64.0.0/10, 192.0.0.0/24, 192.88.99.0/24, 198.18.0.0/15**—Set aside for a variety of special purposes.
- **192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24**—Set aside for use in documentation and examples.

IPv4 Address Scheme Design

The following factors must be weighed when planning an IPv4 network addressing scheme:

- The number of IP networks and subnetworks required must be considered.
- The number of hosts per subnet that must be supported must be considered.
- The network ID must be from a valid public or a private range (not from the loopback, link-local reserved range, multicast range, or reserved/experimental range, for instance).
- The network and/or host IDs cannot be all 1s in binary—this is reserved for broadcasts.
- The network and/or host ID cannot be all 0s in binary; 0 means “this network.”
- Each host ID must be unique on the IP network or subnet.
- The network ID must be unique on the Internet (if you are using a public addressing scheme) or on your internal system of internetworks (if you are using a private addressing scheme).

When you are performing subnet calculations, try to think in terms of the number of mask bits. It helps to remember that each power of 2 is double the previous one:

2^2	2^3	2^4	2^5	2^6	2^7	2^8
4	8	16	32	64	128	256

Also memorize the decimal values for the number of bits set to 1 in an octet within a mask:

1	2	3	4	5	6	7	8
128	192	224	240	248	252	254	255

In the following example, the network designer is subnetting the network address 172.30.0.0/16. The process of designing the scheme is as follows:

1. Work out how many subnets are required (remembering to allow for future growth), then round this number up to the nearest power of 2.

For example, if you need 12 subnets, the next nearest power of 2 is 16. The exponent is the number of bits you will need to add to your default mask. For example, 16 is 2^4 (2 to the power of 4), so you will need to add 4 bits to the network prefix. In dotted decimal format, the subnet mask becomes 255.255.240.0.

2. Work out how many hosts each subnet must support and whether there is enough space left in the scheme to accommodate them.

For example, the network address is in the /16 range, and you are using 4 bits for subnetting, so you have $32 - 20 = 12$ bits for hosts in each subnet. The number of hosts per subnet can be expressed using the formula $2^n - 2$, where n is the number of bits you have allocated for the host ID. 12 bits is enough for 4,094 hosts in each subnet.



You subtract 2 because each subnet's network address and broadcast address cannot be assigned to hosts.



Wherever a 1 appears in the binary mask, the corresponding digit in the IP address is part of the network or subnet address. When you are planning what your mask will be, remember this rule. Allocate more bits in the mask if you need more subnets. Allocate fewer bits in the mask if you need more hosts per subnet.

Just for comparison, if you have a /24 (or Class C) network address and try to allocate 16 subnets, there will be enough space left for only 14 hosts per subnet ($2^4 - 2$).

3. Work out the subnets. The easiest way to find the next subnet ID is to deduct the least significant octet in the mask (240 in this example) from 256. This gives the next subnet ID, which, in full, is 172.30.16.0/20.


The subsequent subnet IDs are all the lowest subnet ID higher than the one before—32, 48, 64, and so on.

4. Work out the host ranges for each subnet. Take the subnet address and add a binary 1 to it for the first host. For the last host, take the next subnet ID and deduct two binary digits from it. For the 172.30.16.0/20 subnet, this is 172.30.16.1 and 172.30.31.254, respectively. Repeat for all subnets.

Classless Inter-Domain Routing

Classless addressing was designed to solve two major problems of the classful addressing scheme as more and more networks joined the Internet through the early 1990s. The first was that network addresses, specifically, Class B addresses, were becoming very scarce and the second was near exponential growth in Internet routing tables. As well as support subnet addressing, classless addressing can be used to summarize network prefixes and reduce the complexity of Internet routing tables.


Classless Inter-Domain Routing (CIDR) uses bits normally assigned to the network ID to mask the complexity of the subnet and host addressing scheme within that network. CIDR is also sometimes described as supernetting.



Using CIDR/supernetting to summarize four /24 networks as one /22 network.


For example, rather than allocate a Class B (or /16) network address to a company, several contiguous Class C (or /24) addresses could be assigned. Four /24 network addresses gives 1,016 hosts. However, this would mean complicated routing with many entries in the routing tables to represent four IP networks at the same location. Using CIDR collapses these routing entries into one single entry. If the network addresses assigned to a company were 198.51.100.0 through to 198.51.103.0 and you wanted to view this as one network, you need to allocate two bits from the network address to summarize the four networks. This makes the supernet prefix /22 or the subnet mask 255.255.252.0.

The ANDing process is still used to determine whether to route. If the ANDed result reveals the same network ID as the destination address, then it is the same network. In this next example, the first IP addresses belong to the supernet, but the second is on a different company's network:



Comparing two addresses with /22 prefixes. ANDing reveals that they are separate networks.

Routers external to the network just use this /22 prefix, so the complexity of the LAN subnets is hidden and doesn't need to clog up their routing tables. The LAN's internal routers use the /24 prefix or even multiple prefixes to create subnets of different sizes.




CIDR public and private route advertisements. (Images © 123RF.com.)

Variable Length Subnet Masks

The IPv4 address space is close to being exhausted, making it difficult for ISPs to allocate public addresses to the companies that want them. To mitigate this, more efficient methods of allocating IP addresses must be used. Supernetting simplifies the information Internet routers need to locate IP networks. A complementary classless addressing technique, called **variable length subnet masking (VLSM)**, allows a network designer to allocate ranges of IP addresses to subnets that match the predicted need for numbers of subnets and hosts per subnet more closely. Without VLSM, you have to allocate subnetted ranges of addresses that are the same size and use the same subnet mask throughout the network. This typically means that some subnets have many wasted IP addresses or additional routing interfaces must be installed to connect several smaller subnets together within a single building or department.

VLSM allows different length netmasks to be used within the same IP network, allowing more flexibility in the design process.

For this example, consider a company with three sites, each with differing network sizes and IP address requirements. There are also subnets representing point-to-point WAN links between the routers.



VLSM requirements for host addresses. (Images © 123RF.com.)

VLSM design usually proceeds by identifying the subnets with the most hosts and organizing the scheme in descending order. As with any subnet calculations, it helps to remember that each power of 2 is double the previous one:

2^2	2^3	2^4	2^5	2^6	2^7	2^8
4	8	16	32	64	128	256


1. In the example, the largest requirement is for 80 hosts. 2^6 has a maximum of 64 values, which is not enough, so the nearest match in the table is 2^7 . This tells us that we need 7 bits for host addressing. This actually allows for 126 host addresses once the network and broadcast addresses have been accounted for ($2^7 - 2$). Using 7 bits makes the prefix /25 (32 minus 7).

2. The next requirement is technically met by a 5-bit host address space, but as this allows for exactly 30 addresses, there would be no room for growth. Using 6 bits might be safer, but for this scenario, we will choose the closest match and adopt the /27 prefix.
3. The next three requirements are for 8, 12, and 12 hosts. These all require 4 bits, which gives 14 usable addresses.
4. The routers use point-to-point links, so no more than two addresses will ever be required. This can be met by selecting a /30 prefix.

The final VLSM design is summarized in the following table:

Office/ Subnet	Required Number of IP Addresses	Mask Bits	Actual Number of IP Addresses	Prefix
Main Office 1 (Router A)	80	7	126	/25
Main Office 2 (Router A)	30	5	30	/27
Main Office 3 (Router A)	8	4	14	/28
Branch Office (Router B)	12	4	14	/28
Branch Office (Router C)	12	4	14	/28
Router A – Router B	2	2	2	/30
Router A – Router C	2	2	2	/30
Router B – Router C	2	2	2	/30

All subnets except for Main Office 2 have room for growth.




VLSM design address space utilization.

In fact, if you analyze the final design, you will find that there are 36 unused addresses at the end of the range. Consequently, there would have been space to use a /26 prefix for the group of 30 hosts.

The actual IP address ranges generated by the VLSM design are shown in this table.

Office	Subnet	Subnet Mask	Useable Host Address Range	Broadcast Address
Main Office 1 (Router A)	198.51.100.0/25	255.255.255.128	1–126	127
Main Office 2 (Router A)	198.51.100.128/27	255.255.255.224	129–158	159
Main Office 3 (Router A)	198.51.100.160/28	255.255.255.240	161–174	175
Branch Office (Router B)	198.51.100.176/28	255.255.255.240	177–190	191
Branch Office (Router C)	198.51.100.192/28	255.255.255.240	193–206	207
Router A – Router B	198.51.100.208/30	255.255.255.252	209–210	211
Router A – Router C	198.51.100.212/30	255.255.255.252	213–214	215
Router B – Router C	198.51.100.216/30	255.255.255.252	217–218	219

The VLSM network topology can be summarized by this diagram:



VLSM design topology diagram. (Images © 123RF.com.)

Lesson 4.4

IP Troubleshooting Tools



Exam Objectives Covered

5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

TCP/IP command-line utilities enable you to gather information about how your systems are configured and how they communicate over an IP network. When used for troubleshooting, these utilities can provide information about communication issues and their causes.

As you study this lesson, answer the following questions:

- What command-line tools report a host's IP configuration? Are there different tools for use in Windows and Linux?
- Is there a tool to verify whether one host can contact another host?
- How can I report information about how a host maps an IP address to a MAC address?

ipconfig Tool

On a Windows host, the **ipconfig command** is widely used for basic configuration reporting and support tasks. `ipconfig` can be used as follows:

- `ipconfig` without any switches will display the IP address, subnet mask, and default gateway (router) for all network interfaces to which TCP/IP is bound.
- `ipconfig /all` displays complete TCP/IP configuration parameters for each interface, including whether the Dynamic Host Configuration Protocol (DHCP) is enabled for the interface and the interface's hardware (MAC) address.
- `ipconfig /renew interface` forces a DHCP client to renew the lease it has for an IP address.
- `ipconfig /release interface` releases the IP address obtained from a DHCP server so that the interface(s) will no longer have an IP address.
- `ipconfig /displaydns` displays the Domain Name System (DNS) resolver cache.
- `ipconfig /flushdns` clears the DNS resolver cache.
- `ipconfig /registerdns` registers the host with a DNS server (if it supports dynamic updates).

```

PS C:\Windows\system32> ipconfig /all

Windows IP Configuration

Host Name . . . . . : PC10
Primary Dns Suffix . . . . . : corp.515support.com
Node Type . . . . . : Mixed
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : corp.515support.com

Ethernet adapter Ethernet:

Connection-specific DNS Suffix . : corp.515support.com
Description . . . . . : Microsoft Hyper-V Network Adapter
Physical Address. . . . . : 00-15-5D-00-65-31
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . : Yes
IPv6 Address. . . . . : fd00:2413:6d1c:30:997b:634e:5b90:7e(Preferred)
Link-local IPv6 Address . . . . . : fe80::997b:634e:5b90:7e%9(Preferred)
IPv4 Address. . . . . : 10.1.24.101(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : Wednesday, August 4, 2021 12:11:31 AM
Lease Expires . . . . . : Thursday, August 12, 2021 12:11:30 AM
Default Gateway . . . . . : fe80::215:5dff:fe00:6510%9
                          10.1.24.254
DHCP Server . . . . . : 10.1.16.1
DHCPv6 IAID . . . . . : 67114333
DHCPv6 Client DUID. . . . . : 00-01-00-01-27-E6-CC-0C-00-15-5D-00-65-31
DNS Servers . . . . . : 10.1.16.1
NetBIOS over Tcpip. . . . . : Enabled

PS C:\Windows\system32>

```

*Identifying the current IP configuration with ipconfig.
(Screenshot used with permission from Microsoft.)*



There are also `/release6` and `/renew6` switches for use with DHCPv6 (a DHCP server supporting IPv6).

ifconfig and ip Tools

On a Linux host, when it comes to managing the running configuration, you need to distinguish between legacy and current command packages. **ifconfig** is part of the legacy `net-tools` package. Use of these commands is deprecated on most modern Linux distributions. `ifconfig` can still safely be used to report the network interface configuration, however.

```

lamp@lamp:~$ ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
    inet 172.16.0.201 netmask 255.255.255.0  broadcast 172.16.0.255
    inet6 fe80::215:5dff:fe00:6517 prefixlen 64  scopeid 0x20<link>
    ether 00:15:5d:00:65:17 txqueuelen 1000  (Ethernet)
    RX packets 4042  bytes 589111 (589.1 KB)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 7788  bytes 2885069 (2.8 MB)
    TX errors 0  dropped 0  overruns 0  carrier 0  collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    inet6 ::1 prefixlen 128  scopeid 0x10<host>
    loop txqueuelen 1000  (Local Loopback)
    RX packets 5244  bytes 413133 (413.1 KB)
    RX errors 0  dropped 0  overruns 0  frame 0
    TX packets 5244  bytes 413133 (413.1 KB)
    TX errors 0  dropped 0  overruns 0  carrier 0  collisions 0

```

ifconfig output.

net-tools has been replaced by the iproute2 package. These tools can interface properly with modern network configuration manager packages. As part of the iproute2 package, the **ip command** has options for managing routes as well as the local interface configuration. The basic reporting functionality of `ifconfig` (show the current address configuration) is performed by running `ip addr`; to report a single interface only, use `ip addr show dev eth0`. The `ip link` command shows the status of interfaces, while the `ip -s link` reports interface statistics.

```
lamp@lamp:~$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 00:15:5d:00:65:17 brd ff:ff:ff:ff:ff:ff
    inet 172.16.0.201/24 brd 172.16.0.255 scope global dynamic eth0
        valid_lft 6026sec preferred_lft 6026sec
    inet6 fe80::215:5dff:fe00:6517/64 scope link
        valid_lft forever preferred_lft forever
```

ip a command output.

The `ip link set eth0 up|down` command is used to enable or disable an interface, while `ip addr add|delete` can be used to modify the IP address configuration. These changes are not persistent and apply only to the running configuration, unless run as part of a startup script.

arp Tool

The Address Resolution Protocol (ARP) is used by hosts to determine which MAC address is associated with an IP address on the local network. ARP queries are sent as broadcasts. ARP broadcasts can generate considerable traffic on a network, which can reduce performance. To optimize this process, the results of an ARP broadcast are cached in an ARP table. If the entry is used within the timeout period, the entry is held in the cache for a few minutes before it is deleted.

The **arp command** can be used to perform functions related to the ARP table cache. You would use this to diagnose a suspected problem with local addressing and packet delivery.

- `arp -a` (or `arp -g`) shows the ARP cache contents. You can use this with `IPAddress` to view the ARP cache for the specified interface only. The ARP cache will not necessarily contain the MAC addresses of every host on the local segment. There will be no cache entry if there has not been a recent exchange of frames.
- `arp -s IPAddress MACAddress` adds an entry to the ARP cache. Under Windows, `MACAddress` needs to be entered with hyphens between each hex byte.
- `arp -d *` deletes all entries in the ARP cache; it can also be used with `IPAddress` to delete a single entry.

```
PS C:\Windows\system32> arp -a

Interface: 10.1.24.101 --- 0x9
    Internet Address      Physical Address      Type
    10.1.24.254           00-15-5d-00-65-10    dynamic
    10.1.24.255           ff-ff-ff-ff-ff-ff    static
    224.0.0.22            01-00-5e-00-00-16    static
    224.0.0.251           01-00-5e-00-00-fb    static
    224.0.0.252           01-00-5e-00-00-fc    static
    239.255.255.250       01-00-5e-7f-ff-fa    static
    255.255.255.255       ff-ff-ff-ff-ff-ff    static
```

Output from the arp command showing network (IP) addresses mapped to physical (MAC) addresses. Host interfaces are learned (dynamic), while broadcast and multicast interfaces are configured statically. (Screenshot used with permission from Microsoft.)

In Linux, the `ip neigh` command shows entries in the local ARP cache (replacing the old `arp` command).



Don't confuse the ARP cache with a MAC address table. ARP cache is maintained by layer 3 hosts and routers to map IP addresses to MAC addresses. A switch's MAC address table contains the MAC addresses that the switch has seen on each of its ports.

ping Tool

The **Internet Control Message Protocol (ICMP)** is used to report errors and send messages about the delivery of a packet. ICMP messages are generated under error conditions in most types of unicast traffic but not for broadcast or multicast packets.

ICMP can also be used to test and troubleshoot connectivity issues on IP networks. The **ping command** sends a configurable number and size of ICMP request packets to a destination host. `ping` is implemented on both Windows and Linux hosts. `ping` can be used to perform a basic connectivity test that is not dependent on the target host running any higher-level applications or services.

Basic ping Usage

A basic connectivity test is performed by running `ping IPAddress`, where `IPAddress` is an IPv4 or IPv6 address.

If the probe is successful (as in the first attempts shown in the screen capture), the output shows the message "Reply from IPAddress" and the time it takes for the server's response to arrive. The millisecond measures of round-trip time (RTT) can be used to diagnose latency problems on a link.

```

PS C:\Windows\system32> ping 127.0.0.1 -n 1

Pinging 127.0.0.1 with 32 bytes of data:
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128

Ping statistics for 127.0.0.1:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Windows\system32> ping 10.1.24.101 -n 1

Pinging 10.1.24.101 with 32 bytes of data:
Reply from 10.1.24.101: bytes=32 time<1ms TTL=128

Ping statistics for 10.1.24.101:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Windows\system32> ping 10.1.24.254 -n 1

Pinging 10.1.24.254 with 32 bytes of data:
Reply from 10.1.24.254: bytes=32 time<1ms TTL=64

Ping statistics for 10.1.24.254:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Windows\system32> ping 203.0.113.33 -n 1

Pinging 203.0.113.33 with 32 bytes of data:
Reply from 203.0.113.33: bytes=32 time=2ms TTL=60

Ping statistics for 203.0.113.33:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 2ms, Maximum = 2ms, Average = 2ms

```

Using ping in Windows. (Screenshot used with permission from Microsoft.)

The Time to Live (TTL) IP header field is reduced by one every time a packet is forwarded by a router (referred to as a hop). The TTL output field in the ping command shows the value of the counter when the packet arrived at its destination.



To work out the number of hops it took, you need to know the initial value. Different operating systems and OS versions use different default values. For example, if you ping a remote host from a Windows 10 host and the TTL value in the output is 52, then you know the packet took 12 hops (64–52) to reach its destination.

ping Error Messaging

If ping probes are unsuccessful, one of two messages are commonly received:

- **Destination host unreachable**—There is no routing information (that is, the local computer does not know how to get to that IP address). This might be caused by some sort of configuration error on the local host, such as an incorrect default gateway, by a loss of connectivity with a router, or by a routing configuration error.
- **No reply (Request Timed Out.)**—The host is unavailable or cannot route a reply to your computer. Requests time out when the TTL is reduced to 0 because the packet is looping (because of a corrupted routing table), when congestion causes delays, or when a host does not respond.

ping Switches

ping can be used with several switches. You can use a host name or fully qualified domain name rather than an IP address to test name resolution. When pinging by name, use `-4` or `-6` to force the tool to query the IPv4 host record or IPv6 host record respectively. Also, `-t` continues to ping the host until interrupted (by pressing `Ctrl+C`).

ping has different syntax when used under Linux. By default, the command executes until manually halted, unless run with the number of packets set by the `-c` switch.

Lesson 4.5

IP Version 6



Exam Objectives Covered

1.8 Summarize evolving use cases for modern network environments.

The previous topics focused on IP version 4 (IPv4), which is still the mainstream version of the protocol. In this topic, you will explain IP version 6 (IPv6) addressing. As a network professional, you should be aware of the limitations of IPv4 and the increasing adoption of IPv6. You need to understand the characteristics of IPv6, as well as how it can interoperate with existing IPv4 implementations.

As you study this section, answer the following questions:

- Why is IPv6 needed?
- What is the format of an IPv6 address?
- How can an IPv6 address be simplified?
- What are the two parts of an IPv6 address?
- What allows IPv6 hosts to communicate over an IPv4 network?

IPv4 Versus IPv6

In IPv4, the addressing scheme is based on a 32-bit binary number. 32 bits can express 2^{32} unique addresses (in excess of four billion). However, the way in which addresses have been allocated has been inefficient, leading to waste of available addresses. Inefficiencies in the addressing scheme and unceasing demand for more addresses mean that the available IPv4 public address supply is exhausted.

IP version 6 (IPv6) is designed to mitigate address exhaustion. Its 128-bit addressing scheme has space for 340 undecillion unique addresses. Even though only a small part of the scheme can currently be allocated to hosts, there is still enough address space within that allocation for every person on the planet to own approximately 4,000 addresses. As well as coping with the growth in ordinary company networks and Internet access subscribers, IPv6 is designed to meet the demands of billions of personal and embedded devices with Internet connectivity.

An IPv6 packet consists of two or three elements: the main header, which is a fixed length (unlike in IPv4), one or more optional extension headers, and the payload. As with an IPv4 header, there are fields for the source and destination addresses and the version (0110 or 0x06 for IPv6). Some of the other header fields are as follows:

Field	Explanation
Traffic Class	Describes the packet's priority.
Flow Label	Used for quality of service (QoS) management, such as for real-time streams. This is set to 0 for packets not part of any delivery sequence or structure.
Payload Length	Indicates the length of the packet payload, up to a maximum of 64 KB; if the payload is bigger than that, this field is 0, and a special Jumbo Payload (4 GB) option is established.
Next Header	Used to describe what the next extension header (if any) is, or where the actual payload begins.
Hop Limit	Replaces the TTL field in IPv4 but performs the same function.

Extension headers replace the Options field in IPv4. There are several predefined extension headers to cover functions such as fragmentation and reassembly, security (IPSec), source routing, and so on.

IPv6 Address Format

An IPv6 address contains eight 16-bit numbers, with each double-byte expressed as 4 hex digits. For example, consider the following binary address:

```
0010 0000 0000 0001 : 0000 1101 1011 1000 :
0000 0000 0000 0000 : 0000 0000 0000 0000 :
0000 1010 1011 1100 : 0000 0000 0000 0000 :
1101 1110 1111 0000 : 0001 0010 0011 0100
```

This binary value can be represented in hex notation as:

```
2001:0db8:0000:0000:0abc:0000:def0:1234
```

Using **canonical notation**, the hex notation can be compressed further. Where a double byte contains leading 0s, they can be ignored. In addition, one contiguous series of 0s can be replaced by a double colon place marker. Thus, the prior address would become:

```
2001:db8::abc:0:def0:1234
```

You can only use double colon compression once in a given address. For example, `2001:db8::abc::def0:1234` is not valid as it is unclear which of the following two addresses is represented:

```
2001:db8:0000:0abc:0000:0000:def0:1234
```


```
2001:db8:0000:0000:0abc:0000:def0:1234
```

Where IPv6 addresses are used as part of a URL (web address), because both formats use colon delimiters to mean different things, the IPv6 address must be contained within brackets. For example:

```
https://[2001:db8::abc:0:def0:1234]/index.htm.
```

IPv6 Network Prefixes

An IPv6 address is divided into two parts: the first 64 bits are used as a network ID, while the second 64 bits designate a specific interface. Unlike in IPv4, the interface address (or host ID portion) is always the same 64-bit length.



*In IPv6, the interface identifier is always the last 64 bits.
The first 64 bits are used for network addressing.*

Network addresses are written using classless notation, where */nn* is the length of the network prefix in bits. Within the 64-bit network ID, as with IPv4 netmasks, the length of any given network prefix is used to determine whether two addresses belong to the same IP network. For example, if the prefix is */48*, then if the first 48 bits of an IPv6 address were the same as another address, the two would belong to the same IP network. This means that a given organization's network can be represented by a global routing prefix 48 bits long, and they then have 16 bits left in the network ID to subnet their network. For example,

`2001:db8:3c4d::/48`

would represent a network address, while:

`2001:db8:3c4d:0001::/64`

would represent a subnet within that network address.

Like IPv4, IPv6 can use unicast, multicast, and anycast addressing. Unlike IPv4, there is no broadcast addressing.

IPv6 Unicast Addressing

As with IPv4, an IPv6 unicast address identifies a single network interface. IPv6 unicast addressing is scoped; a scope is a region of the network. Global scopes provide the equivalent of public addressing schemes in IPv4, while link local schemes provide private addressing.

IPv6 Global Addressing

Globally scoped unicast addresses are routable over the Internet and are the equivalent of public IPv4 addresses. The parts of a global address are as follows:

- The first 3 bits (001) indicate that the address is within the global scope. Most of the IPv6 address space is unused. The scope for globally unique unicast addressing occupies just 1/8th of the total address space. In hex, globally scoped unicast addresses will start with a 2 (0010 in binary) or 3 (0011).
- The next 45 bits are allocated in a hierarchical manner to regional registries and from them to ISPs and end

- The next 16 bits identify site-specific subnet addresses.
- The final 64 bits are the interface ID.



IPv6 global unicast address format.

Interface ID/EUI-64

The 64-bit interface ID can be determined by using two techniques.


One is by using the interface's MAC address. This is known as a MAC-derived address or interface identifier. As a MAC address is 48 bits, a translation mechanism allows driver software to create a 64-bit interface ID from these 48 bits. Formally, this is called an Extended Unique Identifier-64 (EUI-64).

Two changes occur to derive the EUI-64 interface ID from an interface's MAC address. First, the digits `fffe` are added in the middle of the MAC address. Second, the first 8 bits, or 2 hex digits, are converted to binary, and the 7th bit (or U/L bit) is flipped (from 0 to 1 or 1 to 0). For example, the MAC address `00608c123abc` would become the EUI-64 address `02608cfffe123abc`, which (when expressed in double bytes) becomes `0260:8cff:fe12:3abc`, or (without the leading 0) `260:8cff:fe12:3abc`.

In the second technique, referred to as privacy extensions, the client device uses a pseudorandom number for the interface ID. This is known as a temporary interface ID or token. There is some concern that using interface identifiers would allow a host to be identified and closely monitored when connecting to the Internet, and using a token mitigates this to some degree.

IPv6 Link Local Addressing

Link local addresses span a single subnet (they are not forwarded by routers). Nodes on the same link are referred to as neighbors. The link local range is `fe80::/10`. Link local addresses start with a leading `fe80`, with the next 54 bits set to 0, and the last 64 bits are the interface ID.



IPv6 link local unicast address format.

The equivalent in IPv4 is Automatic Private IP Addressing (APIPA) and its `169.254.0.0` addresses. However, unlike IPv4, an IPv6 host is always configured with link local addresses (one for each link), even if it also has a globally unique address.

A link local address is also appended with a **zone index** (or scope id) of the form %1 (Windows) or %eth0 (Linux). This is used to define the source of the address and make it unique to a particular link. For example, a given host may have links to a loopback address, Ethernet, and a VPN. Each of these links may use the same link local address, so each is assigned a zone ID to make it unique. Zone indices are generated by the host system, so where two hosts communicate, they may be referring to the link using different zone IDs.



While it is relatively uncommon for an interface to have more than one IPv4 address, in IPv6 it is typical for an interface to have multiple addresses.

IPv6 Multicast and Anycast Addressing

A multicast address is used to send a packet from a single source to multiple network interfaces. All IPv6 routers must support multicast. The parts of a multicast address are subdivided as follows:

- The first 8 bits indicate that the address is within the multicast scope (1111 1111 or ff).
- The next 4 bits are used to flag types of multicast if necessary; otherwise, they are set to 0.
- The next 4 bits determine the scope; for example, 1 is node-local (to all interfaces on the same node), and 2 is link local.
- The final 112 bits define multicast groups within that scope.

The Multicast Listener Discovery (MLD) protocol allows nodes to join a multicast group and discover whether members of a group are present on a local subnet.

Broadcast addresses are not implemented in IPv6. Instead, hosts use an appropriate multicast address for a given situation. The well-known multicast addresses are ones reserved for these types of broadcast functionality. They allow an interface to transmit to all interfaces or routers on the same node or local link.

In IPv4, IP address resolution to a specific hardware interface is performed using ARP. ARP uses inefficient broadcasts and requires every node to process its messages, whether they are relevant to the node or not. IPv6 replaces ARP with the **Neighbor Discovery (ND) Protocol**.

Each unicast address for an interface is configured with a corresponding solicited-node multicast address. It has the prefix `ff02::1:ff` plus the last 24 bits of the unicast address. The solicited-node address is used by ND to perform address resolution. It greatly reduces the number of hosts that are likely to receive ND messages (down to one in most cases) and is therefore much more efficient than the old ARP broadcast mechanism.

IPv6 can also use anycast addressing, though as with IPv4, this is implemented by a routing protocol rather than having a special range of addresses. Anycast interfaces are those configured with the same IPv6 global unicast address.

IPv4 and IPv6 Transition Mechanisms

A network is likely to have to run both IPv4 and IPv6 in some or all segments and for connectivity with internetworks. This compatibility requirement can be implemented using dual stack hosts, using a tunneling mechanism, or by using a type of address translation.

Dual Stack

Dual stack hosts and routers can run both IPv4 and IPv6 simultaneously and communicate with devices configured with either type of address. Most modern desktop and server operating systems implement dual stack IP. Most modern dual stack systems will try to initiate communications using IPv6 by default.



Most services are addressed using names rather than IP addresses. This means that the preference for IPv6 over IPv4 or the availability of either addressing method depends on the Domain Name System (DNS) records for the network.

```
PS C:\Windows\system32> ipconfig /all

Windows IP Configuration

Host Name . . . . . : PC10
Primary Dns Suffix . . . . . : corp.515support.com
Node Type . . . . . : Mixed
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : corp.515support.com

Ethernet adapter Ethernet:

Connection-specific DNS Suffix . : corp.515support.com
Description . . . . . : Microsoft Hyper-V Network Adapter
Physical Address. . . . . : 00-15-5D-00-65-31
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . : Yes
IPv6 Address. . . . . : fd00:2413:6d1c:30:997b:634e:5b90:7e(Preferred)
Link-local IPv6 Address . . . . . : fe80::997b:634e:5b90:7e%9(Preferred)
IPv4 Address. . . . . : 10.1.24.101(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : Wednesday, August 4, 2021 12:11:31 AM
Lease Expires . . . . . : Thursday, August 12, 2021 12:11:30 AM
Default Gateway . . . . . : fe80::215:5dff:fe00:6510%9
                          10.1.24.254
DHCP Server . . . . . : 10.1.16.1
DHCPv6 IAID . . . . . : 67114333
DHCPv6 Client DUID. . . . . : 00-01-00-01-27-E6-CC-0C-00-15-5D-00-65-31
DNS Servers . . . . . : 10.1.16.1
NetBIOS over Tcpip. . . . . : Enabled

PS C:\Windows\system32>
```

Dual stack IP in Windows. (Screenshot used with permission from Microsoft.)

Tunneling

As an alternative to dual stack, **tunneling** can be used to deliver IPv6 packets across an IPv4 network. Tunneling means that IPv6 packets are inserted into IPv4 packets and routed over the IPv4 network to their destination. Routing decisions are based on the IPv4 address until the packets approach their destinations, at which point the IPv6 packets are stripped from their IPv4 carrier packets and forwarded according to IPv6 routing rules. This carries a high protocol overhead and is not nearly as efficient as operating dual stack hosts.

In 6to4 automatic tunneling, no host configuration is necessary to enable the tunnel. 6to4 addresses use the prefix `2002::: / 16`. 6to4 has been widely replaced by an enhanced protocol called IPv6 Rapid Deployment (6RD). With 6RD, the `2002::/16` prefix is replaced by an ISP-managed prefix, and there are various other performance improvements.

Microsoft provides support for tunneling by Windows hosts using its Teredo protocol. Teredo tunnels IPv6 packets as IPv4-based UDP messages over port 3544. Teredo requires compatible clients and servers. The open source Miredo package implements the Teredo for UNIX/Linux operating systems.

Another option for tunneling is Generic Routing Encapsulation (GRE). GRE allows a wide variety of Network layer protocols to be encapsulated inside virtual point-to-point links. This protocol has the advantage that because it was originally designed for IPv4, it is considered a mature mechanism and can carry both v4 and v6 packets over an IPv4 network.

NAT64

A third approach to transitioning from IPv4 to IPv6 is to use Network Address Translation (NAT). This is a well-known process for rewriting network addresses as they pass routing boundaries. With **NAT64**, an IPv6 host addresses an IPv4 host using the prefix `64:ff9b::: / 96` plus the 32-bit IPv4 destination address. When the packet reaches the gateway router, it strips the prefix and forwards the packet using IPv4 headers. Replies from the IPv4 host are directed to the IPv6 host by tracking connections using Transport layer port numbers.

Common IPv6 Address Prefixes

Use the following table to help you recognize some of the commonly used classes of IPv6 address by prefix notation or leading hex digits.

Type	Prefix	Leading Hex Characters
Global unicast	<code>2000::/3</code>	2 3
Link local unicast	<code>fe80::/10</code>	fe80
Multicast	<code>ff00::/8</code>	ff
Multicast (link local)	<code>ff02::/16</code>	ff02::1 (all nodes) ff02::2 (all routers) ff02::1:2 (DHCP)
Solicited-node	<code>ff02::1:ff00:0/104</code>	ff02::1:ff
Unspecified	<code>::/128</code>	0::0
Loopback	<code>::1/128</code>	::1
Documentation/ Examples	<code>2001:db8::/32</code>	2001:db8

Globally unique unicast addresses are also widely referred to as /48s.

The 0000::/8 block (that is, IPv6 addresses where the first bits are 0000 0000) is reserved for special functions. Within this block, there are two special addresses defined:

- **Unspecified address (0:0:0:0:0:0:0:0)**—A host that has not obtained a valid address. This is often expressed as ::.
- **Loopback address (0:0:0:0:0:0:0:1)**—Used for testing (for the host to send a packet to itself). This is often expressed as ::1.

Lesson 4.6

IP Troubleshooting



Exam Objectives Covered

5.3 Given a scenario, troubleshoot common issues with network services.

While some support scenarios require a top-to-bottom or bottom-to-top approach, in practical terms a lot of troubleshooting activity starts with the Network layer. Connectivity tests and configuration information can identify issues within this layer or inform your decision of whether to move up or down layers to pinpoint the cause of a problem.

As you study this lesson, answer the following questions:

- What are the symptoms of different IP misconfiguration scenarios?
- How can an incorrect subnet mask cause IP communication issues?
- What are the symptoms of an incorrect default gateway?

IP Configuration Issues

Troubleshooting IP configuration issues takes place at the Network layer. If you can rule out a problem at the Physical and Data Link layers, the next thing to check is basic addressing and protocol configuration.

If a host cannot perform neighbor discovery to contact any other nodes on the local network, first use `ipconfig` (Windows) or `ip` or `ifconfig` (Linux) to verify the host configuration.

Incorrect IP Address

Each end system host must have the same netmask as its neighbors and an IP address that produces a valid, unique host address within that subnet. A neighbor in this sense is another host in the same layer 2 broadcast domain. For example, if the subnet is `192.168.1.0/24`, consider the following host address configurations:

- Host A: IP: `192.168.1.10`, Mask: `255.255.255.0`
- Host B: IP: `192.168.1.11`, Mask: `255.255.255.0`
- Host C: IP: `192.168.0.21`, Mask: `255.255.255.0`

Host A and Host B have valid configurations, but Host C has an address in a different subnet (`192.168.0.0` compared to `192.168.1.0`). Hosts A and B will try to use the default gateway to forward packets to Host C. Host C is unlikely to be able to communicate on the network at all.

When you encounter non-default masks, it can be slightly more difficult to identify valid host ranges. For example, if the subnet address is $198.51.100.16/28$, consider the following host address configurations:

- Host A: IP: 198.51.100.10, Mask: 255.255.255.240
- Host B: IP: 198.51.100.11, Mask: 255.255.255.240
- Host C: IP: 198.51.100.21, Mask: 255.255.255.240

The network prefix boundary lies within the last octet, so you cannot rely on the first three octets alone. However, if you convert the addresses to binary, you will find that Host C is in a different subnet.

Also, remember that neither the network address nor the broadcast address can be used as a host address.

Incorrect Subnet Mask

Another issue that might arise if a subnet mask is incorrect is that the host can receive communications but misroutes its replies, thinking that the hosts communicating with it are on a different subnet. The replies may still get through, although they may go via the default gateway (router), rather than directly.

- Host A: IP: 192.168.1.10, Mask: 255.255.255.0
- Host B: IP: 192.168.1.11, Mask: 255.255.255.0
- Host C: IP: 192.168.1.21, Mask: 255.255.255.240

Because it is using a longer prefix than it should, Host C will think it needs to route to a different subnet to communicate with Hosts A and B. This will cause packets to go via the router, placing unnecessary load on it.

The other scenario for an incorrect mask is where the mask is shorter than it should be:

- Host A: IP: 192.168.1.10, Mask: 255.255.255.0
- Host B: IP: 192.168.1.11, Mask: 255.255.255.0
- Host C: IP: 192.168.1.21, Mask: 255.255.0.0
- Host D: IP: 192.168.0.10, Mask: 255.255.255.0

In this case, the problem will not be obvious if Hosts A, B, and C are attached to the same switch, as they will be able to use ARP messaging and receive replies. However, Host C will not be able to contact Host D, as it thinks that Host D is on the same local network, whereas in fact it needs to route messages for $192.168.0.0/24$ via the default gateway.



In this scenario, the router might send ICMP redirect status messages to Host C.

Duplicate IP and MAC Address Issues

Two systems could end up with the same IP address because of a configuration error; perhaps both addresses were statically assigned, or one was assigned an address that was part of a DHCP scope. If Windows detects a duplicate IP address, it will display a warning and disable the IP. Linux does not typically disable an interface with a duplicate IP address, but it may log a warning. If there are two systems with duplicate IPs, a sort of race condition will determine which receives traffic. Obviously, this is not a good way for the network to be configured, and you should identify and fix the machines. To do this, obtain the MAC addresses of both interfaces using `ping` and then `arp -a` to examine the ARP cache table. On Linux, you can use the `arping` tool (`arping -D`) to report duplicate replies. Once identified, configure each host to use a unique address.

A duplicate MAC address will cause a problem similar to a duplicate IP address. Both hosts will contend to respond to ARP queries, and communications could be split between them or reach only one of the hosts. Duplicate MAC addresses are unlikely to arise unless the network uses locally administered addressing.




Issues with MAC addressing can be a sign that someone is attempting to perform a spoofing attack. Spoofing attacks are discussed later in the course.

To diagnose MAC address issues, use the `arp` utility to verify the MAC addresses recorded for each host and `ipconfig` or `ip neigh` to check the MAC address assigned to the interface. Also check the MAC address and ARP tables on any switches and routers involved in the communications path. You can use a protocol analyzer to examine ARP traffic and identify which IP hosts are attempting to claim the same MAC address.

IP Forwarding Issues

If the address configuration on the local host seems to be correct, you can complete a series of connectivity tests using `ping` to determine the likely location and scope of a fault.



A general ping sequence for identifying connectivity issues.

1. Ping the loopback address (`ping 127.0.0.1`) to verify TCP/IP is installed and loaded correctly. If this fails, reinstall the network protocol stack.
2. Ping the IP address of the local host to verify it was added correctly and to verify that the network adapter is functioning properly. If you cannot ping the host's own address, there might have been a configuration error, or the network adapter or adapter driver could be faulty.
3. Ping the IP address of the default gateway to verify it is up and running and that you can communicate with another host on the local network.
4. Ping the IP address of other hosts on the same subnet to test for local configuration or link problems.

If a local host cannot be pinged and the error is destination unreachable, then verify the IP configuration does not contain an incorrect IP address or netmask. If these are correct but pings still time out, suspect either a security issue (such as a switch port security issue) or a problem at the Data Link or Physical layer.

5. Ping the IP address of a remote host to verify you can communicate through the router. If a remote IP address cannot be contacted, check the default gateway parameter on the local host to rule out an incorrect default gateway issue. If the gateway is configured correctly and you can ping the router, you need to start investigating the routing infrastructure.



When performing tests using ping, always be aware that ICMP could be blocked by a firewall or other security software, especially when pinging remote hosts.



This methodical approach is suitable when you cannot diagnose the cause of a problem, or when you are verifying a new or changed IP configuration. In practice, you might start with a simple ping test to a remote host to identify or reproduce the problem.

Module 4

Summary

You should be able to use an appropriate IP addressing scheme to plan a subnetted network and be able to configure and troubleshoot host addressing.

Guidelines for Planning IP Addressing Schemes

Follow these guidelines to configure subnets and use appropriate IP addressing schemes:

- Ensure good understanding of IPv4 addressing concepts to facilitate network design and support:
 - The use of 32-bit IPv4 addresses and netmasks or network prefixes to identify networks and subnets within networks.
 - The role of the Address Resolution Protocol (ARP) in mapping layer 3 to layer 2 IP:MAC addresses.
 - The importance and uses of unicast, broadcast, multicast, and anycast addressing schemes.
 - The impact of legacy classful addressing features on address selection and usage, especially with regard to private versus public ranges and loopback, class D, and class E ranges.
 - How features of classless addressing such as supernetting and VLSM allow for better routing design and address space utilization.
- Work out a logical topology of subnets to create broadcast domain segments that meet requirements for performance, security, and Physical/Data Link network technologies.
- Allocate more bits to the netmask to create more subnets with fewer hosts per subnet, or fewer bits to the netmask to create fewer subnets with more hosts per subnet.
- Ensure that each host is configured with an appropriate IP address/subnet mask and default gateway for its subnet.
- Use netsh or PowerShell to configure IP address properties in Windows. The ipconfig tool can be used to quickly report the adapter configuration.
- Use ip or the legacy ifconfig command to report adapter configuring in Linux.
- Use the arp and ping utilities to troubleshoot issues with local addressing and connectivity.

- Use methodical troubleshooting techniques to diagnose common problems:
 - View the local configuration and cache to identify address configuration issues, such as incorrect IP, incorrect mask, incorrect gateway, or duplicate IP/MAC address.
 - Ping loopback, local, and then remote to determine connectivity and problem scope.
- Ensure good understanding of IPv6 addressing concepts to facilitate network design and support:
 - The use of 128-bit IPv6 addresses and with network prefixes and 64-bit interface identifiers.
 - The use of local and global unicast plus multicast addressing schemes.
 - Dual stack, tunneling, and NAT64 mechanisms to transition from IPv4 to IPv6 networks.

Module 5

Configuring Routing and Advanced Switching

Module Introduction

Now that you are aware of the basic concepts of IP addressing and forwarding, you can start identifying ways that paths between networks are implemented. Routers work at layer 3 to aggregate information about neighboring networks and forward packets along an appropriate path to their final destination.

While configuring routing infrastructure is often a senior job role, you should understand basic concepts and be able to apply them to solve common issues.

Module Objectives

In this module, you will do the following:

- Compare and contrast routing concepts.
- Compare and contrast dynamic routing concepts.
- Install and troubleshoot routers.
- Explain tiered switching architecture.
- Explain virtual LANs.

Lesson 5.1

Routing Technologies



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 2.1 Explain characteristics of routing technologies.
- 5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

Routers support internetworks of all sizes and types. It is critical for a network professional to understand the process a router applies to make forwarding decisions. In this lesson, you will examine routing tables and how they are used to select a forwarding path. You will also learn to use command line tools to report and test the routing configuration.

As you study this lesson, answer the following questions:

- What is the function of a routing table?
- What information is contained in routing table entries?
- What are directly connected, static, and default routes?
- What are the uses of route, show route, tracert, and traceroute tools?

Routing Tables and Path Selection

Most end system workstation and server computers are configured with a single network adapter connected to only one network. Although potentially capable of routing, they are not equipped with the necessary interfaces and knowledge of the location of other networks. A router is a multihomed intermediate system with links and network topology information to facilitate forwarding packets between subnets or around an internetwork.

Information about the location of other IP networks and hosts is stored in the **routing table**. Each entry in the routing table represents an available path to a destination network or host. The following main parameters define a routing entry:

- **Protocol**—The source of the route. Paths can be configured statically or learned by exchanging information with other routers via a dynamic routing protocol.
- **Destination**—Routes can be defined to specific hosts but are more generally directed to network IDs. The most specific (longest) destination prefix will be selected as the forwarding path if there is more than one match.
- **Interface**—The local interface to use to forward a packet along the chosen route. This might be represented as the IP address of the interface or as a layer 2 interface ID.
- **Gateway/next hop**—The IP address of the next router along the path to the destination.

```

RCOREX(config)#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 10.254.0.254 to network 0.0.0.0

10.0.0.0/8 is variably subnetted, 28 subnets, 2 masks
D    10.0.0.0/30 [90/3584] via 10.0.0.14, 00:05:41, GigabitEthernet1/0/4
    [90/3584] via 10.0.0.10, 00:05:41, GigabitEthernet1/0/3
C    10.0.0.4/30 is directly connected, GigabitEthernet1/0/2
C    10.0.0.8/30 is directly connected, GigabitEthernet1/0/3
C    10.0.0.12/30 is directly connected, GigabitEthernet1/0/4
D    10.0.0.16/30 [90/28416] via 10.0.0.102, 00:19:26, Port-channel1
D    10.0.0.20/30 [90/3072] via 10.0.0.6, 00:19:26, GigabitEthernet1/0/2
D    10.0.0.24/30 [90/3072] via 10.0.0.10, 00:19:26, GigabitEthernet1/0/3
D    10.0.0.28/30 [90/3072] via 10.0.0.14, 00:19:26, GigabitEthernet1/0/4
C    10.0.0.100/30 is directly connected, Port-channel1
D    10.10.1.0/24 [90/25625856] via 10.0.0.6, 00:19:26, GigabitEthernet1/0/2
D    10.10.2.0/24 [90/25625856] via 10.0.0.6, 00:19:26, GigabitEthernet1/0/2
D    10.11.1.0/24 [90/25625856] via 10.0.0.6, 00:19:26, GigabitEthernet1/0/2
...
C    10.254.0.252/30 is directly connected, GigabitEthernet1/0/1
S*   0.0.0.0/0 [1/0] via 10.254.0.254

```

Routing table for a Cisco router showing directly connected subnets (C) and routes learned from the EIGRP dynamic protocol (D). There is also a static route identifying the gateway of last resort/default gateway. (Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)



An IPv4 host route has a /32 prefix, while an IPv6 host route has a /128 prefix. Directly connected hosts are often defined by host routes. Anycast addressing also requires that the anycast address be entered in routing tables as a host route.

Static and Default Routes

Routing table entries fall into four general categories:

- **Directly connected routes**—For subnets for which the router has a local interface.
- **Remote routes**—For subnets and IP networks that are not directly attached.
- **Host routes**—To a specific IP address. A host route has a /32 (IPv4) or /128 (IPv6) prefix.
- **Default route**—To use when an exact match for a network or host route is not found.

Directly Connected Routes

The IP network or subnet for each active router interface is automatically added to the routing table. These are known as **directly connected routes**.

Static Routes

A **static route** is manually added to the routing table and only changes if edited by the administrator. Configuring static routing entries can be useful in some circumstances, but it can be problematic if the routing topology changes often, as each route on each affected router needs to be updated manually.




Static routes can be configured either as non-persistent or persistent/permanent. A non-persistent route is removed from the routing table if the router is rebooted. A non-persistent route might be added as a troubleshooting action, for instance. If a static route is not reachable, it will be disabled.

Default Route

A **default route** is a special type of static route that identifies the next hop router for a destination that cannot be matched by another routing table entry. The destination address 0.0.0.0/0 (IPv4) or ::/0 (IPv6) is used to represent the default route. The default route is also described as the gateway of last resort. Most end systems are configured with a default route (pointing to the default gateway). This may also be the simplest way for an edge router to forward traffic to an ISP's routers.

Routing Table Example

As examples of directly connected, static, and default route entries that might appear in a routing table, consider the following example of three routers connected in a series:



Routing tables for three routers connected in a series. (Images © 123RF.com.)

First, consider the routing table of router A:

- The router is directly connected to networks 10.0.1.0/24 (via interface G0) and 10.0.2.0/24 (via interface G1).
- The router has been configured with static routes to 10.0.3.0/24 and 10.0.4.0/24, both of which are reachable via interface G1.

Router B has been configured in the same way, but here the networks 10.0.2.0/24 and 10.0.3.0/24 are directly connected, and the paths to 10.0.1.0/24 and 10.0.4.0/24 are configured as static entries.

Router C has been configured differently. It is directly connected to 10.0.3.0/24 and 10.0.4.0/24, but the only static route configured is for 0.0.0.0/0. This is a default route. While the router has no specific knowledge of networks 10.0.1.0/24 and 10.0.2.0/24, it will forward packets for these destinations over its G0 interface.


Packet Forwarding

When a router receives a packet, it reads the destination address in the packet and looks up a matching destination network IP address and prefix in its routing table. If there is a match, the router will forward the packet out of one of its interfaces by encapsulating the packet in a new frame:

- If the packet can be delivered to a directly connected network via an Ethernet interface, the router uses ARP (IPv4) or Neighbor Discovery (ND in IPv6) to determine the Data Link layer address of the destination interface.
- If the packet can be forwarded via a gateway over an Ethernet interface, it inserts the next hop router's MAC address as the destination address in a new frame and uses the MAC address of the outgoing interface as the source address.
- If the packet can be forwarded via a gateway over another type of interface (leased line or DSL, for instance), the router encapsulates the packet in an appropriate frame type.
- If the destination address cannot be matched to a route entry, the packet is either forwarded via the default route or dropped (and the source host is notified that it was undeliverable).

Hop Count

If the packet is forwarded via a gateway, this process is repeated at each router to deliver the packet through the internetwork. Each router along the path counts as one **hop**. For example, in the network shown in the figure, Host A takes one hop to communicate with LOCAL_SRV via a directly connected interface on the LAN router. Note that the switches do not count as hops. Host B takes multiple hops (nine) to communicate with REMOTE_SRV, with traffic routed via two ISP networks. Also, observe the alternative routes that could be taken. Do any have a lower hop count?



Measuring hop count through an internetwork. (Images © 123RF.com.)

Time to Live

At each router, the **Time to Live (TTL)** header field is decreased by at least 1. This could be greater if the router is congested. The TTL is nominally the number of seconds a packet can stay on the network before being discarded. While TTL is defined as a unit of time (seconds), in practice, it is interpreted as a maximum hop count. When the TTL is 0, the packet is discarded. This prevents badly addressed packets from permanently circulating the network.



In IPv6, the field is named *Hop Limit* to formalize the fact that it is a counter not a timer.

Fragmentation

IP provides best-effort delivery of an unreliable and connectionless nature. Delivery is not guaranteed, and a packet might be lost, delivered out of sequence, duplicated, or delayed. It is possible that due to limitations in the underlying network, IP may fragment the packet into more manageable pieces to fit within the maximum transmission unit (MTU) of the Data Link protocol frame.

In IPv4, the ID, Flags, and Fragment Offset IP header fields are used to record the sequence in which the packets were sent and to indicate whether the IP datagram has been split between multiple frames for transport over the underlying Data Link protocol. For example, the MTU of an Ethernet frame is usually 1,500 bytes. An IP datagram larger than 1,500 bytes would have to be fragmented across more than one Ethernet frame. A datagram passing over an internetwork might have to be encapsulated in different Data Link frame types, each with different MTUs.


Most systems try to avoid IP **fragmentation**. IPv6 does not allow routers to perform fragmentation. Instead, the host performs path MTU discovery to work out the MTU supported by each hop and crafts IP datagrams that will fit the smallest MTU.

Router Configuration

Routers serve both to link physically remote networks and subdivide autonomous IP networks into multiple subnets. Router placement is primarily driven by the IP networks and subnets that have been created:

- Hosts with addresses in the same subnet or IP network must not be separated by a router.
- Conversely, hosts with addresses in different subnets or IP networks must be separated by a router.

The figure shows a simplified example of a typical network configuration. An edge router/firewall provides access to the Internet. Traffic between the local subnets is controlled by a separate internal router.



Network placement of edge and internal routers. (Images © 123RF.com.)

As a router appliance does not have a screen or keyboard, it is configured locally either via a serial connection known as a console port or (more usually) remotely over the network by using a protocol such as Secure Shell (SSH). SSH can be used to communicate with the router via the IP address of any configured interface. However, as any given physical interface could suffer a hardware fault or be temporarily unavailable for various reasons, it is considered best practice to create a virtual interface, known as a loopback interface, in the router's operating system and assign it an IP address for use in remotely managing the router. This is a way of giving the router an internal IP address, not connected to any physical network, that is therefore not reliant on a specific network link being available.

```

vyos@vyos:~$ conf
[edit]
vyos@vyos# set protocols rip interface eth0
[edit]
vyos@vyos# set protocols rip interface eth1
[edit]
vyos@vyos# set protocols rip redistribute connected
[edit]
vyos@vyos# commit && save && exit
Saving configuration to '/config/config.boot'...
Done
exit
vyos@vyos:~$ show ip rip
Codes: R - RIP, C - connected, S - Static, O - OSPF, B - BGP
Sub-codes:
      (n) - normal, (s) - static, (d) - default, (r) - redistribute,
      (i) - interface

      Network          Next Hop          Metric From          Tag Time
R(n) 10.0.0.1/32       10.0.2.254        2 10.0.2.254         0 02:57
C(r) 10.0.0.2/32       0.0.0.0           1 self               0
R(n) 10.0.1.0/24       10.0.2.254        2 10.0.2.254         0 02:57
C(i) 10.0.2.0/24       0.0.0.0           1 self               0
C(i) 10.0.3.0/24       0.0.0.0           1 self               0

```

Configuring RIP on a VyOS-based software router. The host can be configured at a local terminal or from a remote computer over Secure Shell (SSH).

Having placed the router at an appropriate point in the network, connected its cabling, and established a management session, the principal configuration tasks are as follows:

- Apply an IP configuration to each interface.
- Configure one or more routing protocols and/or static routes so that the router can serve its function.

Routing Table Tools

Various command line tools are available to report the routing configuration on routers and end systems.

show route

The `show ip route`, `show ipv6 route`, or similar **show route command** will output the active routing table. As well as destination, gateway, AD/metric, and interface, the output will show the source of the route, identified as a letter code (C = connected, S = static, R = RIP, B = BGP, D = EIGRP, O = OSPF, and so on).

show arp

As with any IP host, a router keeps a cache of IP addresses that have been resolved to MAC addresses via the Address Resolution Protocol (ARP). Inspecting the ARP cache table is useful for discovering duplicate IP addresses, IP misconfigurations, and routing protocol misconfigurations. To view the cache, use `show arp` or `show ip arp`.

Each of the router's interfaces has a separate ARP cache. If an entry is listed as incomplete, the router has sent an ARP request but has not received a reply. This indicates that it expects a host with that IP to be present on that network but that the host is not available.

route and ip route

The **route** command is used to view and modify the routing table of end system Windows and Linux hosts.

Apart from loopback addresses and the local subnet, the routing table for an end system generally contains a single entry for the default route. The default route is represented as the destination 0.0.0.0/0. Any traffic that is not addressed to the local subnet is sent over this default route. In Windows, to show the routing table, run `route print`.

```
PS C:\Users\Jaime> route print
=====
Interface List
12...00 15 5d 01 66 65 .....Microsoft Hyper-V Network Adapter
1.....Software Loopback Interface 1
=====

IPv4 Route Table
=====
Active Routes:
Network Destination        Netmask          Gateway           Interface        Metric
0.0.0.0                    0.0.0.0          10.11.2.5         10.11.2.102      15
10.11.2.0                  255.255.255.0    On-link           10.11.2.102      271
10.11.2.102                255.255.255.255  On-link           10.11.2.102      271
10.11.2.255                255.255.255.255  On-link           10.11.2.102      271
127.0.0.0                  255.0.0.0        On-link           127.0.0.1        331
127.0.0.1                  255.255.255.255  On-link           127.0.0.1        331
127.255.255.255            255.255.255.255  On-link           127.0.0.1        331
224.0.0.0                  240.0.0.0        On-link           127.0.0.1        331
224.0.0.0                  240.0.0.0        On-link           10.11.2.102      271
255.255.255.255            255.255.255.255  On-link           127.0.0.1        331
255.255.255.255            255.255.255.255  On-link           10.11.2.102      271
=====

IPv6 Route Table
=====
Active Routes:
If Metric Network Destination      Gateway
12 271 ::/0 fe80::215:5dff:fe01:9903
1 331 ::1/128 On-link
12 271 fde2:d61d:4a8c:112::/64 On-link
12 271 fde2:d61d:4a8c:112:2d58:bb60:66b0:7244/128 On-link
```

*IPv4 and IPv6 routing tables for a Windows host. For IPv4, the host uses 10.11.2.5 as its default gateway. IPv6 default gateway uses the router's link local interface address.
(Screenshot used with permission from Microsoft.)*

In Linux, the `route` command is part of the older, deprecated package of tools. You can use `ip route show` and `ip route add` to achieve the same ends.

```
(root@kali)-[~]
# ip route show
default via 10.11.2.5 dev eth0 proto dhcp src 10.11.2.66 metric 100
10.11.2.0/24 dev eth0 proto kernel scope link src 10.11.2.66 metric 100
```

Output of Linux `ip route show` command.

Tracert and Traceroute

The **traceroute** tool allows you to test the whole path between two nodes with a view to isolating the node or link that is causing the problem.

traceroute


`traceroute` is supported on Linux and router operating systems, such as Cisco IOS. `traceroute` uses UDP probe messages by default. The command issues a UDP probe for port 33434 with a TTL of 1. The first hop should reduce the TTL to zero and respond with an ICMP Time Exceeded message. The command then increments the port number and TTL by one and sends a second probe, which should reach the second hop router. This process is repeated until the end node is reached, which should reply with an ICMP Port Unreachable response.

The output shows the number of hops, the IP address of the ingress interface of the router or host (that is, the interface from which the router receives the probe), and the time taken to respond to each probe in milliseconds (ms). If no acknowledgment is received within the timeout period, an asterisk is shown against the probe. Note that while this could indicate that the router interface is not responding, it could also be that the router is configured to drop packets with expired TTLs silently.

`traceroute` can be configured to send ICMP Echo Request probes rather than UDP by using `traceroute -I`. The `traceroute -6` or `traceroute6` commands are used for IPv6 networks.

tracert

On a Windows system, the same function is performed using the `tracert` command. `tracert` uses ICMP Echo Request probes by default. The command issues an Echo Request probe with a TTL of 1. The first hop should reduce this to zero and respond with a Time Exceeded response. `tracert` then increments the TTL by one each time to discover the full path.



```
PS C:\Windows\system32> tracert 203.0.113.33

Tracing route to 203.0.113.33 over a maximum of 30 hops

  0  <1 ms    <1 ms    <1 ms   10.1.24.254
  1  <1 ms    <1 ms    <1 ms   10.1.128.253
  2   1 ms     *         1 ms   198.51.100.30
  3   1 ms     1 ms     1 ms   198.51.100.253
  4   2 ms     2 ms     1 ms   203.0.113.33

Trace complete.
```

*Using `tracert` in Windows to plot the path from a host in the UK to CompTIA's web server.
(Screenshot used with permission from Microsoft.)*

`tracert` can be used with several switches, which must precede the target IP address or host.

You can use the `-d` switch to suppress name resolution, `-h` to specify the maximum number of hops (the default is 30), and `-w` to specify a timeout in ms (the default is 4,000). If, after increasing the value, destinations are then reachable, you probably have a bandwidth issue to resolve. When used with host names (rather than IP addresses), `tracert` can be forced to use IPv6 instead of IPv4 by adding the `-6` switch.

```
tracert -6 www.microsoft.com
```

Lesson 5.2

Dynamic Routing Technologies



Exam Objectives Covered

2.1 Explain characteristics of routing technologies.

Complex networks need to exchange routing information rapidly to prevent outages, making static routing updates impractical in most cases. This issue is solved by implementing dynamic routing protocols. In this topic, you will compare these protocols to understand the features that make them more or less appropriate for different networks.

As you study this lesson, answer the following questions:

- What is the difference between static and dynamic routing?
- What network link characteristics are used to determine route selection?
- What are the most common routing protocols? Which protocol is best for each situation?

Dynamic Routing Protocols

A **dynamic routing** protocol uses an algorithm and metrics to build and maintain a routing information base. This database stores information about the networks to which the router is connected. Where there are multiple paths, the algorithm and metrics prioritize one over the others. This information can be shared with the router's neighbors. A learned route is one that was communicated to a router by another router. A router can add learned routes from one or more routing protocols to its IP routing table.

Topology and Metrics

The algorithms used for path selection can be categorized according to the topology and metrics that they use to build and update a routing information base and prioritize optimal (or least-cost) paths. Most algorithms are classed as either **distance vector** or as **link state**. Some protocols use a hybrid of different methods to perform path selection more efficiently.

For each protocol that it runs, the router maintains a routing information base of routes discovered by that protocol. These databases are separate to the IP routing table used to determine the forwarding path. The routing protocol's database might contain more than one route to the same destination prefix. In this case, a metric is calculated to determine which path will be selected for use in the IP routing table. The path with the lowest cost metric is preferred.

The type of algorithm determines which factors are used to calculate the metric. For example, distance vector protocols use the number of hops to the destination as the metric. The route with the fewest hops is the least-cost path and will be selected.

Convergence

Convergence is the process whereby routers running dynamic routing algorithms agree on the network topology. Routers must be capable of adapting to changes such as newly added networks, router or router interface failures, link failures, and so on. Routers must be able to communicate changes to other routers quickly to avoid sinkholes and loops. A sinkhole means that a packet is discarded without notification back to the source; a loop causes a packet to be forwarded around the network until its TTL expires.

A network where all the routers share the same topology is described as steady state. The time taken to reach steady state is a measure of a routing protocol's convergence performance.



A flapping interface is one that frequently changes from online to offline and offline to online. Similarly, route flapping refers to a router changing the properties of a route it is advertising quickly and often. Flapping can cause serious convergence problems.

Autonomous Systems


As well as the algorithm used to determine the network topology, routing protocols can be classified according to the way they deal with administrative boundaries. A network under the administrative control of a single owner is referred to as an **autonomous system (AS)**. An Interior Gateway Protocol (IGP) is one that identifies routes within an AS. An Exterior Gateway Protocol (EGP) is one that can advertise routes between autonomous systems. An EGP includes a field to communicate the network's autonomous system ID and allows network owners to determine whether they can use paths through another organization's network.

Routing Information Protocol

The **Routing Information Protocol (RIP)** is a distance vector routing protocol. RIP only considers a single piece of information about the network topology—the next hop router to reach a given network or subnet (vector). It considers only one metric to select the optimal path to a given destination network—the one with the lowest hop count (distance). While RIP is no longer widely deployed, it is useful to review how it works to help to understand the function of more advanced and widely used protocols.


RIP sends regular updates (typically every 30 seconds) of its entire routing database to neighboring routers. It can also send triggered updates whenever changes occur. When a router receives an update from a neighbor, it adds unknown routes to its own routing table, increases the hop count by one, and identifies the originator of the update as the next hop to the specified networks.

In the following figure, RIP has been used to propagate route information between three routers connected in a chain. Router A learns about networks 10.0.3.0/24 and 10.0.4.0/24 from Router B. It adds one to the hop count metric of these routes. Router B learns about 10.0.1.0/24 from Router A and about 10.0.4.0/24 from Router C. Router A and Router C do not exchange any information directly. The distance vector process by which Router A learns about Router C's networks is often referred to as "routing by rumor."




Routers connected in a series exchanging distance vector path updates.

The following example illustrates a mesh topology where there are multiple paths between networks. Router A has two possible paths to network 10.0.3.0/24, which it learns from Router B and Router C. It can forward a packet out of its G1 interface over network 10.0.2.0/24, which will take one hop to reach the destination. It could also forward the packet out of its G2 and reach the destination via Router C and then Router B. This takes two hops and so is not used as the preferred route.



Routers connected in a mesh topology exchanging distance vector path updates.

If Router A's G1 link goes down, those entries will be removed from the routing table, and the alternative routes via 10.0.4.0/24 will be selected:



Network	Interface	Metric	Via	Source
10.0.1.0/24	G0	0		Connected
✗ 10.0.2.0/24	G1	0		Connected
✗ 10.0.2.0/24	G2	2	10.0.4.0	RIP
✗ 10.0.3.0/24	G1	1	10.0.2.0	RIP
10.0.3.0/24	G2	2	10.0.4.0	RIP
10.0.4.0/24	G2	0		Connected
✗ 10.0.5.0/24	G1	1	10.0.2.0	RIP
10.0.5.0/24	G2	1	10.0.4.0	RIP

Routers connected in a mesh topology exchanging distance vector path updates.

To help prevent looping, the maximum hop count allowed is 15. Consequently, this limits the maximum size of a RIP network, since networks that have a hop count of 16 or higher are unreachable.

Enhanced Interior Gateway Routing Protocol

The Interior Gateway Routing Protocol (IGRP) was developed by Cisco to provide a routing protocol for routing within a domain or autonomous system. Limitations in IGRP, such as lack of support for classless addressing, led to the development of **Enhanced IGRP (EIGRP)**. There are versions for IPv4 and IPv6.

Like RIP, EIGRP is a distance vector protocol because it relies on neighboring routers to report paths to remote networks. Unlike RIP, which is based on a simple hop count metric, EIGRP uses a metric composed of administrator weighted elements. The two default elements are bandwidth and delay:

- **Bandwidth**—Applies a cost based on the lowest bandwidth link in the path.
- **Delay**—Applies a cost based on the time it takes for a packet to traverse the link. This metric is most important if the route is used to carry time-sensitive data, such as voice or video. Delay is calculated as the cumulative value for all outgoing interfaces in the path.

Where RIP sends periodic updates of its entire routing information base, EIGRP sends a full update when it first establishes contact with a neighbor and thereafter only sends updates when there is a topology change. This is more efficient and less disruptive to large networks, giving it the best convergence performance in many scenarios. EIGRP does use regular hello messaging to confirm connectivity with its neighbors. Unlike RIP, EIGRP maintains a topology table alongside its routing information base. The topology table is used to prevent loops while also supporting a greater number of maximum hops than RIP (nominally up to 255).

EIGRP is a default IP protocol, which means that it is encapsulated directly in IP datagrams, rather than using TCP or UDP. It is tagged with the protocol number 88 in the Protocol field of the IP header. Updates are transmitted using multicast addressing.


Open Shortest Path First

A distance vector algorithm relies on directly connected neighbors for information about remote networks. By contrast, a link state algorithm allows a router to store the complete network topology and assess the least-cost paths from this topology database.

Open Shortest Path First (OSPF) is suited to large organizations with multiple redundant paths between networks. Where EIGRP is a flat routing system, OSPF is hierarchical. Networks and their connected hosts and routers within an autonomous system are grouped into OSPF areas. Routers within a given area share the same topological database of the networks they serve. Routers that can connect to multiple areas are known as area border routers. A backbone (always called Area 0) is created by the collection of border routers. This backbone is only visible to the border routers and invisible to the routers within a specific area.

In a given area, routers exchange OSPF hello messages, both as a form of a keep-alive packet and in order to acquire neighbors with which to exchange routing information. Neighbors share Link State Advertisement (LSA) updates to build a consistent link state database (LSDB) that represents the network topology of the area. The router applies an algorithm called shortest path first (SPF) to analyze the LSDB and add least-cost, loop free routes to its routing table. This use of a topology table of the whole network to select routes is the key difference between link state and distance vector algorithms.

The small, frequent updates used by OSPF lead to more rapid convergence and more efficiently support larger networks. The use of areas to subdivide the network minimizes the amount of routing traffic that must be passed around the network as a whole, further improving convergence performance. However, link state algorithms can be more expensive to implement because they require more CPU and memory resource.



Typical OSPF topology. (Images © 123RF.com.)

Messages are sent as multicasts using OSPF's own datagram format. This is tagged as protocol number 89 in the IP datagram's Protocol field. There are various packet types and mechanisms to ensure sequencing and reliable delivery and to check for errors. OSPF also supports plaintext or cryptographic authentication.

Border Gateway Protocol

The **Border Gateway Protocol (BGP)** is designed to be used between routing domains in a mesh internetwork and as such is used as the routing protocol on the Internet, primarily between ISPs.

EIGRP and OSPF are used for communications between routers within a single routing domain, referred to as an autonomous system (AS). BGP is primarily used for routing between autonomous systems. An AS is designed to hide the complexity of private networks from the public Internet. If all Internet locations had to be propagated to all Internet routers, the routing tables would become too large to process. Edge routers for each AS exchange only as much network-reachability information as is required to access other autonomous systems (the AS path), rather than networks and hosts within each AS. BGP prioritizes stability and can be slow to converge. Autonomous system numbers (ASN) are allocated to ISPs by IANA via the various regional registries.

BGP works with classless network prefixes called Network Layer Reachability Information (NLRI). Path selection is based on multiple metrics, including hop count, weight, local preference, origin, and community.

BGP works over TCP on port 179.

Route Selection

If a router has multiple entries to similar networks in its routing table, it must determine which route to prefer. The first determining factor is that longer prefixes are preferred over shorter ones. This is referred to as longest prefix match. For example, a routing table contains the following two entries:

```
198.51.100.0/24 g0
```

```
198.51.100.0/28 g1
```

If the router receives a packet for 198.51.100.1, the packet will be routed via g1, as that has the longer and more specific prefix.

Each routing protocol supported by the router can add a single route for any given destination prefix to the routing table. This means that there might be more than one route with an identical length prefix in the routing table. Each routing protocol uses its metric to determine the least-cost path for routes with identical prefix lengths. However, as routing protocols use different methods to calculate the metric, it cannot be used to compare routes from different protocols in the overall IP routing table. Instead, an **administrative distance (AD)** value is used to express the relative trustworthiness of the protocol supplying the route. Default AD values are coded into the router but can be adjusted by the administrator if necessary.

Source	AD
Local interface/Directly connected	0
Static route	1
BGP	20
EIGRP	90
OSPF	110
RIP	120
Unknown	255

This means, for example, that given identical prefix lengths, a static route will be preferred to anything other than directly connected networks and that a route discovered by EIGRP would be preferred to one reported by OSPF. The value of 255 for unknown routes means that they will not be used.

Conversely, a static route with a high AD could be defined to function as a backup if a learned route update fails. In normal circumstances, the router will prefer the learned route because it has a lower AD.

Lesson 5.3

Network Address Translation



Exam Objectives Covered

2.1 Explain characteristics of routing technologies.

As well as understanding the different types of routing algorithms, you must be able to install routing devices to an appropriate place in the network. This lesson will help you to understand the role of routers placed at the network edge or perimeter. On IPv4 networks, this role involves the use of Network Address Translation (NAT) to manage communications between public and private address schemes.

As you study this lesson, answer the following questions:

- Why is there a requirement for network address translation?
- What is the difference between static NAT and dynamic NAT?
- What is port forwarding?
- What is the difference between NAT and PAT?

Edge Routers

Edge routers, placed at the network perimeter, are typified by distinguishing external (Internet-facing) and internal interfaces. These routers can perform framing to repackage data from the private LAN frame format to the WAN Internet access frame format. The customer's router is referred to as the customer edge (CE), while the service provider's router is referred to as the provider edge (PE).



An integrated services router. This type of device combines DSL Internet access with Ethernet switch, Wi-Fi, and VoIP for a "one box" solution for remote sites and branch offices. (Image © 123RF.com.)

Routers designed to service medium to large networks are complex and expensive appliances. They feature specialized processors to handle the routing and forwarding processes, and memory to buffer data. Most routers of this class will also support plug-in cards for WAN interfaces. Another important feature is support for different methods of configuring site-to-site virtual private networks (VPNs).



An advanced services router. This type of device provides network edge connectivity over Carrier Ethernet networks. (Image © 123RF.com.)


Network Address Translation Types

An edge router must facilitate communications between the private and public network. To communicate on the Internet, a host must use an address from a public IP range. As these are in short supply for IPv4, various mechanisms have been devised to reduce the need for public addresses. **Network Address Translation (NAT)** is an example of one of these mechanisms.

NAT is primarily deployed as a service translating between a private (or local) addressing scheme used by hosts on the LAN and a public (or global) addressing scheme used by an Internet-facing device. To perform this role, NAT is configured on a border device, such as a router, proxy server, or firewall. NAT is not a security mechanism; security is provided by the router/firewall's access control list (ACL).

In a static NAT configuration, a 1:1 mapping is made between the private (inside local) network address and the public (inside global) address. If the destination network is using NAT, it is described as having outside global and outside local addressing schemes.

Static NAT is useful in scenarios where an inbound connection to a host must be supported. For example, you might position a web server behind a firewall running NAT. The firewall performs 1:1 address translation on the web server's IP address. This means that external hosts do not know the true IP address of the web server, but they can communicate with it successfully.




Network Address Translation (NAT). (Images © 123RF.com.)



Where the NAT device performs forwarding over selected ports only, this can be referred to as port forwarding.


A single static mapping is not very useful in most scenarios. Under dynamic NAT, the NAT device exposes a pool of public IP addresses. To support inbound and outbound connections between the private network and the Internet, the NAT service builds a table of public to private address mappings. Each new session creates a new public-private address binding in the table. When the session is ended or times out, the binding is released for use by another host.



Defining a one-to-one NAT rule on an OPNsense router/firewall. This rule maps hosts using a 172.200.0.0/24 addressing scheme to 10.200.0.0/24 addresses. (Screenshot used with permission from OPNsense.)

Port Address Translation

Basic NAT supports multiple simultaneous connections but is still limited by the number of available public IP addresses. Smaller companies may only be allocated a single or small block of addresses by their ISPs. In such cases, a means for multiple private IP addresses to be mapped onto a single public address would be useful. This function is provided by **Port Address Translation (PAT)**. This can be referred to as Network Address Port Translation (NAPT), NAT overloading, one-to-many NAT, many-to-one NAT, or NAT masquerade.



Port Address Translation (PAT). (Images ©123RF.com.)

PAT works by allocating each new connection an ephemeral Transport layer port ID. For example, say two hosts (10.0.0.101 and 10.0.0.102) initiate a web connection at the same time. The PAT service creates two new port mappings for these requests (10.0.0.101:61101 and 10.0.0.102:61102) in its state table. It then substitutes the private IP for the public IP and forwards the requests to the public Internet. It performs a reverse mapping on any traffic returned using those ports, inserting the original IP address and port number, and forwarding the packets to the internal hosts.

Lesson 5.4

Firewalls



Exam Objectives Covered

1.2 Compare and contrast networking appliances, applications, and functions.

When a private network is connected to public ones, traffic between the private and public zones needs to be subject to rules. This filtering role is performed by a firewall. Firewalls apply access controls to ensure authorized use of the network. They perform a filtering function to analyze the properties of connection requests and then allow, deny, and/or log them as appropriate. While you may not be installing and configuring these devices at this stage in your career, it is important that you understand their use on the network edge.

As you study this lesson, answer the following questions:

- How is a packet filtering firewall different from a circuit-level gateway?
- Why is a packet filtering firewall a stateless device?
- What devices and software can implement firewall functionality?

Firewall Uses and Types

The basic function of a firewall is traffic filtering. The firewall processes traffic according to rules; traffic that does not conform to a rule that allows it access is blocked.

There are many types of firewalls and many ways of implementing a firewall. One distinction can be made between firewalls that protect a whole network (one that is placed inline in the network and inspects all traffic that passes through) and firewalls that protect a single host only (one that is installed on the host and inspects only that traffic addressed to that host). A further distinction can be made about what parts of a packet a particular firewall technology can inspect and operate on.

Packet Filtering Firewalls

Packet filtering describes the earliest type of firewall. All firewalls can still perform this basic function. A packet filtering firewall is configured by specifying rules in a network access control list (ACL). Each rule defines a specific type of data packet and the appropriate action to take when a packet matches the rule. An action can be either to deny (block or drop the packet, and optionally log an event) or to accept (let the packet pass through the firewall). A packet filtering firewall works at layer 3 of the OSI model to inspect the headers of IP packets. This means that rules can be based on the information found in those headers:

- **IP Filtering**—Accepting or denying traffic based on its source and/or destination IP address.

- Protocol ID/type (TCP, UDP, ICMP, routing protocols, and so on).
- **Port Filtering/Security**—Accepting or denying a packet based on source and destination Transport layer port numbers.



Port numbers are contained in TCP or UDP headers (layer 4) rather than the IP datagram header, but packet filtering firewalls are still almost always described as working at layer 3. They can inspect only port numbers and not any other layer 4 header information.

ACLs might be designed to control only inbound traffic or both inbound and outbound traffic. This is also often referred to as “ingress” and “egress” traffic or filtering. Controlling outbound traffic is useful because it can block applications that have not been authorized to run on the network and defeat malware, such as backdoors. Ingress and egress traffic is filtered using separate ACLs.

A packet filtering firewall is stateless. This means that it does not preserve information about the connection between two hosts. Each packet is analyzed independently with no record of previously processed packets. This type of filtering requires the least processing effort, but it can be vulnerable to attacks that are spread over a sequence of packets. A stateless firewall can also introduce problems in traffic flow, especially when some sort of load balancing is being used or when clients or servers need to make use of dynamically assigned ports.

Stateful Inspection Firewalls

A circuit-level stateful inspection firewall addresses these problems by maintaining stateful information about the session established between two hosts (including malicious attempts to start a bogus session). Information about each session is stored in a dynamically updated state table. A stateful firewall operates at layer 5 (Session) of the OSI model. When a packet arrives, the firewall checks it to confirm whether it belongs to an existing connection. If it does not, it applies the ordinary packet filtering rules to determine whether to allow it. Once the connection has been allowed, the firewall allows traffic to pass unmonitored, in order to conserve processing effort.

Firewall: Diagnostics: States

Select rule

StatesActions

Q

Search

All


☐	Int	Dir	Proto	Source	Nat	Destination	State	Rule	Commands
☐	all	←	icmp	198.51.100.29:53210		198.51.100.30:53210	0:0	let out anything fro...	
☐	all	→	tcp	10.1.128.253:443		10.1.24.101:49885	FIN_WAIT_2:FIN_W...	anti-lockout rule	
☐	all	→	tcp	172.16.0.201:80		203.0.113.44:49692	TIME_WAIT:TIME_W...	Allow web access (u...	
☐	all	←	tcp	203.0.113.44:49692		172.16.0.201:80	TIME_WAIT:TIME_W...	let out anything fro...	
☐	all	→	tcp	172.16.0.201:80		203.0.113.44:49694	TIME_WAIT:TIME_W...	Allow web access (u...	
☐	all	←	tcp	203.0.113.44:49694		172.16.0.201:80	TIME_WAIT:TIME_W...	let out anything fro...	
☐	all	→	tcp	172.16.0.201:80		203.0.113.44:49696	TIME_WAIT:TIME_W...	Allow web access (u...	
☐	all	←	tcp	203.0.113.44:49696		172.16.0.201:80	TIME_WAIT:TIME_W...	let out anything fro...	
☐	all	→	tcp	172.16.0.201:80		203.0.113.44:49698	TIME_WAIT:TIME_W...	Allow web access (u...	
☐	all	←	tcp	203.0.113.44:49698		172.16.0.201:80	TIME_WAIT:TIME_W...	let out anything fro...	
☐	all	→	tcp	172.16.0.201:80		203.0.113.44:49700	FIN_WAIT_2:FIN_W...	Allow web access (u...	

State table in the OPNsense firewall appliance.
(Screenshot used with permission from OPNsense.)

Firewall Selection and Placement

Some types of firewalls are better suited for placement at network or segment borders; others are designed to protect individual hosts. The selection of a network firewall model will largely depend on the volume of traffic it has to process. A single firewall can represent a network bottleneck if it is not able to handle the required traffic volume.

An appliance firewall is a stand-alone hardware firewall that performs only the function of a firewall. The functions of the firewall are implemented on the appliance firmware. This is also a type of network-based firewall and monitors all traffic passing into and out of a network segment. This type of appliance could be implemented with routed interfaces or as a layer 2/virtual wire “transparent” firewall.



*Status dashboard for the OPNsense open source security platform.
(Screenshot courtesy of OPNsense.)*

A router firewall is similar, except that the functionality is built into the router firmware. Most SOHO Internet router/modems have this type of firewall functionality, though they are typically limited to supporting a single subnet within the home network. An enterprise-class router firewall would be able to support far more sessions than a SOHO one.

Lesson 5.5

Enterprise Network Topologies



Exam Objectives Covered

1.6 Compare and contrast network topologies, architectures, and types.

Ethernet, switching, and IP routing are the principal technologies used to implement cabled local networks. There are many types and sizes of network, however, and many different ways of designing cabling and forwarding to suit the requirements of large and small organizations and budgets. While you might not be responsible for network design at this stage of your career, it is important that you be able to identify the components and advantages of the tiered network hierarchies used to implement complex local networks.

- As you study this lesson, answer the following questions:
- What is a corporate datacenter?
- What does the core switch do?
- How do the distribution and access layer switches work?
- What is the access layer switch?


Hybrid Topology

The network topology establishes how nodes are physically and logically connected. Recall that the basic topologies are as follows:

- Point to point is a one-to-one link between two hosts only.
- Star is many-to-many links where each node relies on a central forwarding device to establish paths to other nodes.
- Mesh is many-to-many links that depend on each node having multiple interfaces and cable paths to other nodes.

These basic topologies do not always support network requirements. Often, a more complex hybrid topology is required. A **hybrid topology** is anything that uses a mixture of point-to-point, star, and mesh physical and/or logical topologies. On modern networks, hybrid topologies are often used to implement redundancy and fault tolerance or to connect sites in WANs and in enterprise campus networks:

- **Hierarchical star**—Corporate networks are often designed in a hierarchy, also known as a tree topology. This can be combined with a star topology to implement each node in the overall tree. The links between nodes in the tree are referred to as backbones or trunks because they aggregate and distribute traffic from multiple different areas of the network.
- **Hierarchical Star-mesh**—Alternatively, nodes at the top of the hierarchy can be configured in a partial or full mesh for redundancy. Switches or routers lower in the hierarchy establish star topologies that connect end systems to the network.




Hierarchical hybrid topology example where a mesh network distributes traffic between multiple star networks.

- **Star of stars**—A WAN might be configured as a hub and spoke between a central office and branch offices, with each site implementing a star topology to connect end systems. This is also referred to as a snowflake topology.

Three-Tiered Network Hierarchy

A hierarchical model breaks down a large and complex network design into smaller sections based on the functions performed. Each function can be assessed by network designers to identify the most efficient hardware and software to use to implement it.

As a practical example of this type of model, many corporate office networks follow Cisco's design principles for a **three-tiered hierarchy**: access, distribution, and core.



Core, distribution, and access layers in three-tiered network architecture. (Images © 123RF.com.)


Access/Edge Layer

The **access or edge layer** allows end user devices, such as computers, printers, and smartphones to connect to the network. The access layer is implemented for each site using structured cabling and wall ports for wired access and access points for wireless access. Both are ultimately connected to workgroup switches. Switches deployed to serve the access layer might also be referred to as LAN switches or data switches. End systems connect to switches in the access/edge layer in a star topology. There are no direct links between the access switches.

Distribution/Aggregation Layer

The **distribution or aggregation layer** provides fault-tolerant interconnections between different access blocks and either the core or other distribution blocks. Each access switch has full or partial mesh links to each router or layer 3 switch in its distribution layer block. The distribution layer is often used to implement traffic policies, such as routing boundaries, filtering, or quality of service (QoS).

The **layer 3 switches** used to implement the distribution/aggregation layer have different capabilities to the layer 2 workgroup switches used in the access tier. Rather than 1 Gbps access port and 10 Gbps uplink ports, as would be typical of a workgroup switch, basic interfaces on an aggregation switch would be 10 Gbps and uplink/backbone ports would be 40 Gbps (or possibly 40 Gbps/100 Gbps). Layer 3 switches work on the principle of “route once, switch many,” which means that once a route is discovered, it is cached with the destination MAC address, and subsequent communications are switched without invoking the routing lookup. Layer 3 switches can be far faster, but they are not always as flexible. Layer 3 switches cannot usually perform WAN routing and work with interior routing protocols only. Often layer 3 switches support Ethernet only.



An example of a core/distribution switch. (Image © 123RF.com.)

Core Layer

The **core layer** provides a highly available network backbone. Devices such as client and server computers should not be attached directly to the core. Its purpose should be kept simple: provide redundant traffic paths for data to continue to flow around the access and distribution layers of the network. Routers or layer 3 switches in the core layer establish a full mesh topology with switches in distribution layer blocks.

Collapsed Core

Medium-sized networks might not need separate core and distribution layers. In a two-tier or **collapsed core** model, a monolithic core layer is implemented as a full mesh. This is impractical if there are large numbers of core switches, making the design less scalable.

Lesson 5.6

Virtual LANs



Exam Objectives Covered

- 2.1 Explain characteristics of routing technologies.
- 2.2 Given a scenario, configure switching technologies and features.
- 5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

Most networks make use of virtual LANs (VLANs), both to improve network security and network performance, so they are an important concept for you to understand. In this topic, you will identify the benefits of network segmentation and the characteristics and functions of VLANs.

As you study this lesson, answer the following questions:

- What is the purpose of configuring VLANs?
- How is information about VLANs communicated between switches and routers?
- How do voice VLANs help VoIP work effectively?
- How does routing between VLANs work?


Virtual LANs and Subnets

Modern Ethernet networks are built using switches. In its default configuration, every port on a switch will be in the same local segment or, put another way, in the same broadcast domain. At the Network layer, this group is identified as either an IP network or as a subnet within an IP network. Any host within a broadcast domain can contact any other host using the same logical addressing scheme (IP subnet) and by hardware/MAC addressing.

If too many hosts are attached to the same switch, broadcast traffic can become excessive and reduce performance. At layer 2, virtual LANs (VLANs) are a means of addressing this issue. Each interface on a managed switch can be assigned a VLAN ID. Using VLANs means that different groups of computers on the same cabling and attached to the same switch(es) can appear to be in separate LAN segments. Each VLAN is a separate broadcast domain.

At layer 3, subnetting is the process of logically dividing an IP network into smaller subnetworks (subnets), with each subnet having a unique address. A subnetting scheme can be applied to represent the VLAN design in the layer 3 topology.

For example, the following subnet design allocates separate subnets (10 . 0 . 1 . 0 and 10 . 0 . 2 . 0) for the two VLANs configured on Switch A and for the serial WAN links configured between Router A and Routers B and C (10 . 0 . 3 . 0 and 10 . 0 . 4 . 0). Routers B and C also have a subnet each for their local networks (10 . 0 . 5 . 0 and 10 . 0 . 6 . 0).



Subnet design. (Images © 123RF.com.)

Apart from breaking up broadcast domains, VLANs and subnets can be used to achieve other network design goals:

- Many organizations have more than one site with WAN links between them. The WAN link normally forms a separate subnet.
- It is useful to divide a network into logically distinct zones for security and administrative control. VLANs isolate a group of hosts, allowing incoming and outgoing traffic for the group to easily be filtered at the router.
- Some hosts and appliances have similar performance characteristics. VLANs and subnets can be used to group these hosts and optimize their network performance more easily.

Virtual LAN IDs and Membership

Implementing VLANs can reduce broadcast traffic when a network has expanded beyond a certain number of hosts or users. As well as reducing the impact of broadcast traffic, from a security point of view, each VLAN can represent a separate zone. VLANs are also used to separate nodes based on traffic type and the need for quality of service. For example, it is commonplace to put all VoIP handsets on a voice VLAN to minimize interference coming from nodes that are sending email or downloading large files on the same network. The switches and routers can then be configured to give the VoIP VLAN priority over VLANs that contain ordinary client PCs.

VLANs are configured on switches. The VLAN with ID 1 is referred to as the default VLAN. Unless configured differently, all ports on a switch default to being in VLAN 1.

The simplest means of assigning a node to a VLAN is by configuring the port interface on the switch with a VLAN ID in the range 2 to 4094. For example, from the switch management interface, ports 1 through 10 could be configured as a VLAN with the ID 10 and ports 11 through 20 could be assigned to VLAN 20. Host A connected to port 2 would be in VLAN 10, and Host B connected to port 12 would be in VLAN 20. Host A and Host B would not be able to communicate directly, even though they are connected to the same switch. Each VLAN is typically configured with its own subnet address and IP address range. Communications between VLANs must go through an IP router or layer 3 capable switch.

```
VB1F1#show vlan
```

VLAN	Name	Status	Ports
1	default	active	
111	B1F1-WiFi	active	Fa0/1
112	B1F1-Wired	active	Fa0/2, Fa0/3, Fa0/4, Fa0/5 Fa0/6, Fa0/7, Fa0/8, Fa0/9 Fa0/10, Fa0/11
113	B1F1-VoIP	active	Fa0/12, Fa0/13, Fa0/14, Fa0/15 Fa0/16, Fa0/17, Fa0/18, Fa0/19 Fa0/20
114	B1F1-Embedded	active	Fa0/21, Fa0/22, Fa0/23, Fa0/24
1002	fddi-default	active	
1003	token-ring-default	active	
1004	fddinet-default	active	
1005	trnet-default	active	

VLAN	Type	SAID	MTU	Parent	RingNo	BridgeNo	Stp	BrdgMode	Trans1	Trans2
1	enet	100001	1500	-	-	-	-	-	0	0
111	enet	100111	1500	-	-	-	-	-	0	0
112	enet	100112	1500	-	-	-	-	-	0	0
113	enet	100113	1500	-	-	-	-	-	0	0
114	enet	100114	1500	-	-	-	-	-	0	0

Cisco switch showing port 1 assigned to VLAN 111, ports 2–11 in VLAN 112, and so on.
(Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)



VLANs 2–1001 are referred to as the normal range. VLANs 1002–1005 are reserved. Extended VLAN IDs (1006–4094) are not available on older switch OS versions.

A VLAN database is the list of VLANs configured on a switch. A VLAN is created using the following commands from global configuration mode:

```
vlan 16
```

```
name VLAN16
```

The following commands put port 1–13 into VLAN 16:

```
interface range gigabitEthernet 0/1-13
```

```
switchport access vlan 16
```




Do note that on Cisco switches, there is an explicit `vlan database` command. This invokes a legacy configuration mode that is now deprecated. VLAN database mode only supports the normal range 2–1001.

The `show vlan` command reports the VLAN IDs configured on the switch, plus the ports assigned to them. You can use `no vlan` to delete a VLAN.


Trunking and IEEE 802.1Q

On a large network, a single switch will not provide enough ports for all the hosts that need to be connected to the network. This means that multiple switches must be interconnected to build the network fabric. Multiple switches may also be deployed to provide redundant links. The interconnections between switches are referred to as **trunks**. One of the ports on each switch would be configured as a trunk port for this purpose.



VLAN trunk link. (Images © 123RF.com.)

When frames designated for different VLANs are transported across a trunk, the VLAN ID (VID) of each frame must be preserved for the receiving switch to forward it correctly. VIDs are normally defined by the IEEE **802.1Q** standard. Under 802.1Q, per-VLAN traffic is identified by a tag inserted in the Ethernet frame between the Source Address and EtherType fields. The tag contains information about the VID (from 1 to 4,094) and priority (used for QoS functions). The EtherType value is set to identify the frame as 802.1Q.



Construction of an 802.1Q (VLAN-tagged) Ethernet frame.

Tagged and Untagged Ports

If a switch port will only ever participate in a single VLAN, that port can be configured as untagged. This is also referred to as an access port or host port. An untagged/access port uses the following **port tagging** logic:

- If a frame is addressed to a port in the same VLAN on the same switch, no tag needs to be added to the frame.
- If the frame needs to be transported over a trunk link, the switch adds the relevant 802.1Q tag to identify the VLAN, and then forwards the frame over the trunk port.
- If the switch receives an 802.1Q tagged frame on an access port, it strips the tag before forwarding it to the host.

Conversely, a tagged port will normally be one that is operating as a trunk; that is, capable of transporting traffic addressed to multiple VLANs using the 802.1Q frame format. A trunk might be used to connect switches or to connect a switch to a router. In some circumstances, a host attached to a port might need to be configured to use multiple VLANs and would need to be attached to a trunk port, rather than an access port. One example of this is a virtualization host with multiple guest operating systems. The virtual servers might need to be configured to use different VLANs.

Default VLAN and Native VLAN

The VLAN with ID 1 is referred to as the **default VLAN**. This cannot be changed. Unless configured differently, all ports on a switch default to being in VLAN 1. When you are implementing VLANs, you should avoid sending user data traffic over the default VLAN. It should remain unused or used only for inter-switch protocol traffic, where necessary. For example, spanning tree traffic would be permitted to run over the default VLAN. Make sure that unused ports are not assigned to VLAN 1.

A **native VLAN** is one into which any untagged traffic is put when receiving frames over a trunk port. When a switch receives an untagged frame over a trunk, it assigns the frame to the native VLAN. Untagged traffic might derive from legacy devices such as hubs or older switches that do not support 802.1Q encapsulated frames. The native VLAN is initially set with the same VLAN ID (VID) as the default VLAN (VID 1). You can and should change this, however, to make the native VID any suitable ID. This should not be the same as any VLAN used for any other data traffic. The same native VID should be configured for the trunk port on both switches.

Voice VLANs

Voice over IP (VoIP) transmits voice traffic as data packets, rather than over circuit-based transmission lines. The bandwidth and latency requirements of voice traffic mean that it is often necessary to prioritize it over other types of data packets. This can be accomplished using a dedicated VLAN for voice traffic. However, in many cases, VoIP has been implemented into network infrastructures that were originally designed for just desktop and laptop computers, with limited numbers of physical network wall ports.

To accommodate the lack of dedicated wall ports for handsets, most VoIP endpoints incorporate an embedded switch with just two external ports. The handset is connected via its uplink port to the wall port and via the structured cabling to an access switch. The PC or laptop is connected to the handset via the other port. The handset forwards data traffic from the PC to the access switch as untagged frames. The handset sends voice traffic over the same physical link but uses 802.1Q tagged frames.

Normally, for a switch interface to process tagged frames, it would have to be configured as a trunk port. This adds a lot of configuration complexity, so most switches now support the concept of a **voice or auxiliary VLAN** to distinguish the PC and VoIP traffic without having to configure a trunk. In the following example, the interface configuration assigns traffic from the PC to VLAN 100 and the voice traffic to VLAN 101:

```
interface GigabitEthernet0/0
  switchport mode access
  switchport access vlan 100
  switchport voice vlan 101
```



*Sharing a single physical wall port between a PC and VoIP handset. The handset and switch interface configuration allow VoIP traffic to be assigned to a different VLAN than the PC's data traffic.
(Images © 123RF.com.)*


The switch will only accept tagged frames that match the configured voice VLAN ID. To avoid having to configure this manually, the voice VLAN ID and other configuration parameters can be communicated to the handset using a protocol such as Cisco Discovery Protocol (CDP).

VLAN Routing

Many networks are segmented using the VLAN feature of managed switches. Traffic between VLANs must be routed. There are various ways of accomplishing this.

Subinterfaces

One method is to deploy a router with a single interface (a one-armed router or router on a stick) connected to a trunk port on the switch. The trunk port carries all the VLAN-to-VLAN traffic that must be routed. The router's physical interface is configured with multiple **subinterfaces**. Each subinterface is configured with a specific VLAN ID and IP address. The subinterface acts as the default gateway for its VLAN/subnet. The router forwards inter-VLAN traffic between the subinterfaces.



Router-on-a-Stick topology with subinterfaces serving each VLAN/subnet.
(Images © 123RF.com.)

The following commands would configure a subinterface for VLAN16 on G0, using the last available host address as the default gateway:

```
interface G0.16
encapsulation dot1Q 16
ip address 198.51.100.30 255.255.255.240
```



Remember that the choice of default gateway is by convention. Many organizations use the first available host address. In these examples, we're using the last available address (just to focus attention on how many addresses are available in any given subnet). The key is to apply the same convention consistently across the network.



Be aware that it's not necessary for the router to have a single physical interface. This is just a conventional example. It could have multiple physical interfaces each with subinterfaces connected to different switches. It could also have a WAN interface.

Switch Virtual Interfaces

Passing traffic between a router appliance and the switch over a trunk link is relatively inefficient and does not scale well to tens of VLANs. Consequently, enterprise networks usually deploy layer 3 switches in the core and distribution layers of their networks. A layer 3 switch is one that is optimized for routing between VLANs. It can use static and dynamic routing to identify which VLAN a packet with a given destination IP address should be forwarded to.

With layer 3 switches, each VLAN can be assigned a **Switch Virtual Interface (SVI)** to act as the default gateway. For example, if the topology shown in the previous figure were implemented using a single layer 3 switch rather than a router plus layer 2 switch, the SVI for VLAN 16 would be configured as follows:

```
interface VLAN16  
  
ip address 198.51.100.30 255.255.255.240
```

The hosts in that VLAN would be configured with 198.51.100.30 as the default gateway.



Do be aware that a layer 3 switch could also be configured with subinterfaces. Any port on a layer 3 switch can be designated as routed rather than switched using the `no switchport` command. To distinguish the concepts independently of device types, remember that an SVI is bound to a VLAN and doesn't require a physical interface (it's like configuring a virtual router for the virtual LAN); a subinterface is bound to a physical interface and then allocated a VLAN ID.

Lesson 5.7

Routing and VLAN Troubleshooting



Exam Objectives Covered

5.3 Given a scenario, troubleshoot common issues with network services.

Troubleshooting issues with routing and VLANs can be a daunting challenge, so it is important to understand the basic principles of how hosts and routers apply route selection logic to forward packets. You should also be able to use diagnostic tests methodically to establish the scope and likely cause of an issue.

As you study this lesson, answer the following questions:

- How do I apply route selection logic with information in routing tables to diagnose the cause of forwarding problems?
- How can VLAN assignment cause network communication problems?

Routing Table Issues

If you can ping a host's default gateway, but you cannot ping some or all hosts on remote networks, then you should suspect a routing issue. In many cases, this will be because a router has gone offline, and there is no alternative path to the network. Verify that the router is powered on, that cabling to interface ports is correct, and that each interface is up and configured with the correct IP address. If the router's interfaces are working and correctly configured, next investigate the routing topology.

If you suspect a problem with router configuration and the network topology, use `tracert` to try to identify where the network path is failing and the `route` or `show route` commands to investigate the routing tables of intermediate systems at that point in the path.

When inspecting a routing table, you can use `show ip route w.x.y.z` to check for the presence of a route to a specific IP network. A **missing route** may arise because a required static routing entry has not been entered or has been entered incorrectly. Missing routes may also arise because a router fails to communicate with its neighbors and so does not receive routing protocol updates. You might start troubleshooting this by pinging the router nodes that are neighbors of the system with the issue to check basic connectivity. If there is a network path and the neighbors are up, you would investigate the protocol configuration (perhaps there is an authentication issue or incorrect parameter).

If all expected routes are present, there may be a priority problem. Remember that route selection uses the following factors:

1. The most specific path is preferred. For example, if there are paths to 198.50.100.0/28 via 198.50.100.30 and 198.59.100.0/24 via 198.50.100.254, the path to 198.50.100.0/28 will be preferred for destination addresses from .1 to .30 because /28 is a longer prefix than /24.

2. If there are paths with equal prefixes, the path with the lowest administrative distance will be selected. Administrative distance is a measure of how trustworthy the source of the route is. Directly connected and static routes have lower AD values than routing protocols.
3. If there are identical paths with equal AD, the path with the lowest metric value is preferred.


Investigate any paths with overlapping ranges, such as the /24 and /28 example quoted. These are likely to indicate an error, especially if they come from different sources. For example, a misconfigured static route might be in conflict with a learned route.

Default Route and Routing Loop Issues

Each end host must be configured with a default gateway so that it can communicate with remote hosts. The host's routing table should contain an entry for 0.0.0.0/0 that uses the router's IP address as the path. If there is no entry for 0.0.0.0/0, check that the default gateway is configured. If the host cannot contact remote networks, verify that you can ping the default gateway address. If the default gateway router is up, verify its configuration. If the router configuration is faulty or the router or one of its interfaces is down, it's likely that multiple hosts will have connectivity problems.

A default route can also be configured on a router. This might be used as the path to a service provider's routers to contact external networks. Misconfigured default routes or propagating default routes to other routers can be the cause of routing loops.

A **routing loop** occurs when two routers use one another as the path to a network. Packets caught in a routing loop circle around until the TTL expires. One symptom of a potential routing loop is for routers to generate ICMP Time Exceeded error messages.



A routing loop created between Routers B and C. (Images © 123RF.com.)

Routing protocols use various mechanisms to prevent loops. For example, distance vector protocols use the following mechanisms:

- **Maximum hop count**—If the cost exceeds a certain value (16 in RIP), the network is deemed unreachable. A poison route is one advertised with a hop count of 16. This can provide an explicit failure notice to other routers.
- **Holddown timer**—If a node declares a network unreachable, its neighbors start a holddown timer. Any updates about that route received from other nodes are discarded for the duration of the timer. This is designed to ensure that all nodes have converged information about an unreachable network.
- **Split horizon**—Prevents a routing update from being copied back to the source. In the example above, this would prevent Router C from sending an update about a route to Router A via Router B to Router B.

Link state protocols try to ensure that each node has a consistent view of the network through continual, timely updates flooded to all nodes in the routing domain. A loop in a link state routing domain typically indicates that updates are not being propagated correctly.

You can use `tracert` to diagnose a routing loop by looking for IP addresses that appear multiple times in the output.

VLAN Assignment Issues

When you configure a virtual LAN (VLAN), the switch ports assigned to that VLAN are in a segmented network. The VLAN is likely to be assigned its own subnet address. Any device connecting to a port in the VLAN must have an appropriate IP configuration for that subnet in terms of IP address, subnet mask, default gateway, and DNS servers. Hosts in the VLAN must use the router to contact other VLANs or other remote networks.

The VLAN's router (default gateway) is likely to be configured as a type of virtual interface. If the router has a single physical trunk link to the switch hosting the VLAN, then that virtual interface will be configured as a subinterface of the physical interface, such as `interface G0.16`. If a layer 3 switch is used to implement the VLAN, the default gateway address will usually be configured as a Switched Virtual Interface (SVI), such as `interface VLAN16`. Use `show` commands on the router or switch to verify that these parameters have been set correctly and that ports have been assigned to the correct VLANs.

A host must be physically connected to the correct switch port configured with the appropriate VLAN ID. If the switch port is configured with the correct VLAN ID, and the host is still assigned to the wrong VLAN, you should suspect a physical cabling problem. Verify that the cable from the patch panel is connected to the correct switch port. If there are continuing problems, verify that the patch panel labeling correctly identifies the wall jack that the computer is connected to.

If you cannot diagnose the issue by looking at the configuration, use `ping` to test connectivity. You can use exactly the same process as you would a physical LAN:

1. Ping the loopback address, then the host's own IP address to verify that TCP/IP is working and that the host's IP is correctly configured.
2. Ping the default gateway to verify that it is contactable. Optionally, also ping another host on the same VLAN/subnet to verify that local communications can be established. If you cannot ping the default gateway, look for a problem with either the host or switch/router configuration. If you can rule out these and there is no connectivity with any local hosts, check that the patch cable for the workstation's wall port is connected to the correct switch port and for any other physical cable issues.
3. Ping a remote host or server to verify that routing is available. If you cannot ping any remote hosts, verify the switch and router configuration. Determine if the problem is isolated to a single VLAN or is more widespread across the network. If you can ping some remote hosts but not others, use `tracert` to determine the cause.



Hosts in each VLAN must be able to reach Dynamic Host Configuration Protocol (DHCP) servers and Domain Name System (DNS) servers. One option is to place these as hosts within the VLAN, but that could mean provisioning many servers. More typically, they will be in a server VLAN, and hosts in the client VLAN must use the default gateway to contact them. In this setup, a DHCP relay must also be configured. Network service configuration is discussed in more detail in the next module.

Module 5

Summary

You should be able to compare and contrast routing technologies, troubleshoot common general routing issues, and explain how topologies such as three-tiered hierarchy and VLANs can make corporate networks more manageable.

Guidelines for Supporting Routing and Campus Network Design

Follow these guidelines to deploy and support routing technologies and campus network design principles:

- Work out the topology of switches, virtual LANs (VLANs), and routers to create broadcast domain network segments that meet requirements for performance, security, and Physical/Data Link network technologies:
 - Determine whether to implement a core and distribution layer or a single collapsed core layer, based on network size and projected requirements for future expansion.
 - Determine bandwidth requirements within the core/distribution layer (typically 10 Gbps+) and provision appropriate switch modules, transceivers, and cabling (typically fiber optic).
 - Provision redundant trunk links within the core and between the core and distribution layer.
 - Determine bandwidth requirements for the access layer (typically 1 Gbps) and provision appropriate workgroup/LAN switches based on media type.
 - Provision redundant trunk links between distribution layer switch blocks and access layer switches.
 - Connect client devices (PCs, VoIP endpoints, and printers) and non-datacenter servers to access layer switches.
- Determine the organizational principles that will guide VLAN assignment:
 - Design IP subnets for each VLAN and create a VLAN numbering system.
 - Map the logical topology to the physical switch topology and identify trunk links. Configure the interfaces that will participate in trunk links with the VLANs they are permitted to carry.
 - Configure other interfaces as untagged/access ports within the appropriate VLAN.
- Assess how static, default, and dynamic routing can best meet network design requirements.

- Develop understanding of routing tables and forwarding decisions to assist with network design and troubleshooting:
 - How the destination prefix length and route source administrative distance affects forwarding.
 - How routing protocols use metrics to determine the least-cost route.
 - Next hop and time to live attributes of network paths.
- Use the command line tools to investigate host and router routing tables.
- Use traceroute/tracert to test routing.

Module 6

Implementing Network Services

Module Introduction

Layers 1 through 3 of the OSI model are concerned with addressing and packet forwarding and delivery. This basic connectivity is established for the purpose of transporting application data. In this module, you will describe how protocols at layer 4 provision the transport services that network applications depend upon. Also, this module identifies application protocols that perform low-level network operations tasks, such as providing dynamic address or name resolution services.

Module Objectives

In this module, you will do the following:

- Compare and contrast transport protocols.
- Use command line tools to scan network ports.
- Explain the use of network addressing services.
- Explain the use of name resolution services.
- Configure and troubleshoot DHCP and DNS services.

Lesson 6.1

Transport and Application Layer Protocols



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

You have seen how IP provides addressing and delivery at layer 3 of the OSI model. At layer 4, the TCP/IP protocol suite also defines how different applications on separate hosts establish connections and track communications. Understanding how application protocols use ports to establish connections is critical to being able to configure and support network services.

As you study this section, answer the following questions:

- What is the function of a Transport layer port?
- What is the difference between connectionless and connection-oriented services?
- How does TCP differ from UDP?
- What tools can be used to investigate Transport layer connections?

Transport Layer Ports and Connections

The first three layers of the OSI model are primarily concerned with basic connectivity, address, and forwarding. Protocols at the Transport layer (layer 4) are concerned with delivery of multiplexed application data. Transport layer headers don't describe how to forward a packet across a network. They instruct a host what to do with the data in a packet, and optionally, how to verify that it is complete.

A TCP/IP host may be running multiple services or communicating with multiple servers, clients, or peers in parallel. This means that incoming packets must be directed to the appropriate service or application. To facilitate this, each application protocol is assigned a unique identification number called a **port**. A host can operate multiple ports simultaneously.


Port numbers 0 through 1,023 are preassigned by the Internet Assigned Numbers Authority (IANA) to "well-known" server applications. These port assignments are documented at iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml. Other server applications have been registered in the port range 1,024 through 49,151.

The remaining ports (up to 65,535) are designated for private or dynamic use. As well as the server application needing a port, each client application must assign its own port number to track its requests. Client ports are also referred to as ephemeral ports or source ports.



OS implementations of TCP/IP have not always conformed to these recommendations. For example, early versions of Windows and UNIX/Linux used 1,024–5,000 for client ports. Modern Linux kernels often use 32,768–60,999.

The port number is used in conjunction with the source IP address to form a **socket**. Each socket is bound to a software process. Only one process can operate a socket at any one time. A connection is formed when a client socket requests a service from the server socket. A connection is uniquely identified by the combination of server port plus IP address and client port plus IP address. A server socket can therefore support multiple connections from a number of client sockets.



Multiplexing application ports as sockets at the Transport layer. (Images © 123RF.com.)

Transmission Control Protocol

The **Transmission Control Protocol (TCP)** works at the Transport layer to provide connection-oriented, guaranteed communication using acknowledgments to ensure that delivery has occurred. If packets are missing, they can be retransmitted. TCP can be used for unicast transmission only.

TCP takes data from the Application layer as a stream of bytes and divides it up into segments, each of which is given a header. The TCP segments become the payload of the underlying IP datagrams. The use of sequencing, acknowledgments, and retransmissions means that TCP requires numerous header fields to maintain state information. The main fields in the header of a TCP segment are the following:

Field	Explanation
Source port	TCP port of sending host.
Destination port	TCP port of destination host.
Sequence number	The ID number of the current segment (the sequence number of the last byte in the segment). This allows the receiver to rebuild the message correctly and deal with out-of-order packets.
Ack number	The sequence number of the next segment expected from the other host (that is, the sequence number of the last segment received +1). Packets might be out of order because they are delayed, but they could also be lost completely or arrive in a damaged state. In the first case, the lack of acknowledgment results in the retransmission of data and, in the second case, a Negative Acknowledgment (NAK or NACK) forces retransmission.
Data length	Length of the TCP segment.
Flags	Type of content in the segment (ACK, SYN, FIN, and so on).
Window	The amount of data the host is willing to receive before sending another acknowledgment. TCP's flow control mechanism means that if it is getting overwhelmed with traffic, one side can tell the other to slow the sending rate.
Checksum	Ensures validity of the segment. The checksum is calculated on the value of not only the TCP header and payload but also part of the IP header, notably the source and destination addresses. Consequently, the mechanism for calculating the checksum is different for IPv6 (128-bit addresses) than for IPv4 (32-bit addresses).
Urgent Pointer	If urgent data is being sent, this specifies the end of that data in the segment.
Options	Allows further connection parameters to be configured. The most important of these is the maximum segment size. This allows the host to specify how large the segments it receives should be, minimizing fragmentation as they are transported over data link frames.

TCP Handshake and Teardown

A TCP connection is typically established to transfer a single file, so a client session for something like a webpage (HTTP) might involve multiple TCP connections being opened with the server. These connections are managed using handshake transactions, which make use of a number of **TCP flags**.

TCP Three-Way Handshake

A connection is established using a three-way handshake:

No.	Time	Source	Destination	Protocol	Length	Info
1	0.0...	10.1.0.101	10.1.0.2	TCP	66	1624 → 80 [SYN] Seq=0 Win=65535 Len=0 MSS=1460 WS=256
2	0.0...	10.1.0.2	10.1.0.101	TCP	66	80 → 1624 [SYN, ACK] Seq=0 Ack=1 Win=8192 Len=0 MSS=1460
3	0.0...	10.1.0.101	10.1.0.2	TCP	54	1624 → 80 [ACK] Seq=1 Ack=1 Win=262144 Len=0
4	0.0...	10.1.0.101	10.1.0.2	HTTP	433	GET / HTTP/1.1
5	0.0...	10.1.0.2	10.1.0.101	TCP	54	80 → 1624 [ACK] Seq=1 Ack=380 Win=2102272 Len=0
6	0.2...	10.1.0.2	10.1.0.101	TCP	1514	80 → 1624 [ACK] Seq=1 Ack=380 Win=2102272 Len=1460

Internet Protocol Version 4, Src: 10.1.0.101, Dst: 10.1.0.2

Transmission Control Protocol, Src Port: 1624, Dst Port: 80, Seq: 1, Ack: 1, Len: 0

Source Port: 1624

Destination Port: 80

[Stream index: 0]

[TCP Segment Len: 0]

Sequence number: 1 (relative sequence number)

[Next sequence number: 1 (relative sequence number)]

Acknowledgment number: 1 (relative ack number)

0101 = Header Length: 20 bytes (5)

Flags: 0x010 (ACK)

Window size value: 1024

[Calculated window size: 262144]

[Window size scaling factor: 256]

0000 00 15 5d 01 ca 76 00 15 5d 01 ca 77 08 00 45 00 ...]...v...]...w...E...
 0010 00 28 16 49 40 00 80 06 00 00 0a 01 00 65 0a 01 ...(-I@... ..e...
 0020 00 02 06 58 00 50 61 ff 66 ce 99 26 04 92 50 10 ...X...a...f...&...P...
 0030 04 00 14 83 00 00

*Observing the three-way handshake with the Wireshark protocol analyzer.
(Screenshot courtesy of Wireshark.)*

1. The client sends a segment with the TCP flag SYN set to the server with a randomly generated sequence number. The client enters the SYN-SENT state.
2. The server, currently in the LISTEN state (assuming it is online), responds with a SYN/ACK segment, containing its own randomly generated sequence number. The server enters the SYN-RECEIVED state.
3. The client responds with an ACK segment. The client assumes the connection is ESTABLISHED.
4. The server opens a connection with the client and enters the ESTABLISHED state.



Servers can (usually) support thousands or even millions of TCP connections simultaneously.

The sending machine expects regular acknowledgments for segments it sends and, if a period elapses without an acknowledgment, it assumes the information did not arrive and automatically resends it. This overhead makes the system relatively slow. Connection-oriented transmission is suitable when reliability and data integrity are important.

TCP Connection Teardown

There are also functions for resetting a connection and (in some implementations) keeping a connection alive if no actual data is being transmitted (hosts are configured to time out unused connections). To close a connection, also referred to as teardown, the following basic steps are performed:

1. The client sends a FIN segment to the server and enters the FIN-WAIT1 state.
2. The server responds with an ACK segment and enters the CLOSE-WAIT state.
3. The client receives the ACK segment and enters the FIN-WAIT2 state. The server sends its own FIN segment to the client and goes to the LAST-ACK state.
4. The client responds with an ACK and enters the TIME-WAIT state. After a defined period, the client closes its connection.
5. The server closes the connection when it receives the ACK from the client.

Some implementations may use one less step by combining the FIN and ACK responses into a single segment operation.

```
Microsoft Windows [Version 10.0.14393]
(c) 2016 Microsoft Corporation. All rights reserved.

C:\Users\administrator>netstat -ano

Active Connections

    Proto Local Address           Foreign Address         State           PID
    TCP    0.0.0.0:135             0.0.0.0:0              LISTENING      652
    TCP    0.0.0.0:445             0.0.0.0:0              LISTENING      4
    TCP    0.0.0.0:5985            0.0.0.0:0              LISTENING      4
    TCP    0.0.0.0:47001           0.0.0.0:0              LISTENING      4
    TCP    0.0.0.0:49664           0.0.0.0:0              LISTENING      428
    TCP    0.0.0.0:49665           0.0.0.0:0              LISTENING      912
    TCP    0.0.0.0:49666           0.0.0.0:0              LISTENING      864
    TCP    0.0.0.0:49669           0.0.0.0:0              LISTENING      1996
    TCP    0.0.0.0:49670           0.0.0.0:0              LISTENING      524
    TCP    0.0.0.0:49703           0.0.0.0:0              LISTENING      516
    TCP    0.0.0.0:49706           0.0.0.0:0              LISTENING      524
    TCP    10.1.0.100:139          0.0.0.0:0              LISTENING      4
    TCP    10.1.0.100:49764        10.1.0.192:3000        ESTABLISHED    4280
    TCP    [::]:135                [::]:0                 LISTENING      652
    TCP    [::]:445                [::]:0                 LISTENING      4
    TCP    [::]:5985                [::]:0                 LISTENING      4
    TCP    [::]:47001              [::]:0                 LISTENING      4
```

Observing TCP connections with the netstat tool. (Screenshot used with permission from Microsoft.)

A host can also end a session abruptly using a reset (RST) segment. This would not be typical behavior and might need to be investigated. A server or security appliance might refuse connections using RST, a client or server application might be faulty, or there could be some sort of suspicious scanning activity ongoing.

User Datagram Protocol

The **User Datagram Protocol (UDP)** also works at the Transport layer, but unlike TCP, it is a connectionless, nonguaranteed method of communication with no acknowledgments or flow control. There is no guarantee regarding the delivery of messages or mechanism for retransmitting lost or damaged packets. If an application uses UDP and needs reliability mechanisms, these must be implemented in the Application layer or software logic.

UDP is suitable for applications that send small amounts of data in each packet and do not require acknowledgment of receipt. It is used by Application layer protocols that need to send multicast or broadcast traffic. It may also be used for applications that transfer time-sensitive data but do not require complete reliability, such as voice or video. Using small packets means that if a few are lost or arrive out of order, they only manifest as minor glitches in playback quality. The reduced overhead means that overall delivery is faster.

This table shows the structure of a UDP datagram.

Field	Explanation
Source port	UDP port of sending host.
Destination port	UDP port of destination host.
Message length	Size of the UDP packet.
Checksum	Ensures validity of the packet.

The header size is 8 bytes, compared to 20 bytes (or more) for TCP.

Netstat

The **netstat** command allows you to check the state of ports on the local host. You can use `netstat` to check for service misconfigurations, such as a host running a web or FTP server that a user installed without authorization. You may also be able to identify suspicious remote connections to services on the local host or from the host to remote IP addresses.

On Windows, used without switches, the command outputs active TCP connections, showing the local and foreign addresses and ports. Using the `-a` switch displays all open ports, including both active TCP and UDP connections and ports in the listening state.

On Linux, running `netstat` without switches shows active connections of any type. If you want to show different connection types, you can use the switches for Internet connections for TCP (`-t`) and UDP (`-u`), raw connections (`-w`), and UNIX sockets/local server ports (`-x`). Using the `-a` switch includes ports in the listening state in the output. `-l` shows only ports in the listening state, omitting established connections.

For example, the following command shows listening and established Internet connections (TCP and UDP) only: `netstat -tua`.


```
lamp@lamp:~$ netstat -tua
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address           Foreign Address         State
tcp        0      0 0.0.0.0:5000            0.0.0.0:*               LISTEN
tcp        0      0 localhost:mysql         0.0.0.0:*               LISTEN
tcp        0      0 localhost:domain        0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:ssh             0.0.0.0:*               LISTEN
tcp        0      0 localhost:33060         0.0.0.0:*               LISTEN
tcp        0      1 172.16.0.201:52492      172.16.0.254:domain     SYN_SENT
tcp6       0      0 [::]:http               [::]:*                  LISTEN
tcp6       0      0 [::]:ssh                 [::]:*                  LISTEN
udp        0      0 172.16.0.201:43367     172.16.0.254:domain     ESTABLISHED
udp        0      0 172.16.0.201:42410     172.16.0.254:domain     ESTABLISHED
udp        0      0 172.16.0.201:47084     172.16.0.254:domain     ESTABLISHED
udp        0      0 localhost:domain        0.0.0.0:*               ESTABLISHED
udp        0      0 172.16.0.201:bootpc    0.0.0.0:*               ESTABLISHED
```

Linux netstat output showing active and listening TCP and UDP connections.

On both Windows and Linux, `-n` displays ports and addresses in numerical format. Skipping name resolution speeds up each query. On Linux, using `-4` or `-6` filters sockets by IPv4 or IPv6 addresses respectively. In Windows, use the `-p` switch with the protocol type (TCP, TCPv6, UDP, or UDPv6).

Another common task is to identify which software process is bound to a socket. On Windows, `-o` shows the process ID (PID) number that has opened the port, while `-b` shows the process name. In Linux, use `-p` to show the PID and process name.

`netstat -s` reports per protocol statistics, such as packets received, errors, discards, unknown requests, port requests, failed connections, and so on. The tool will report Ethernet statistics using `-e` (Windows) or `-I` (Linux). `netstat -r` displays the routing table.

```
lamp@lamp:~$ netstat -i
Kernel Interface table
Iface    MTU     RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0     1500    4069   0      0 0      8134   0      0 0 BMRU
lo       65536   5322   0      0 0      5322   0      0 0 LRU
```

Linux netstat interface statistics showing receive and transmit packets numbers plus errors and dropped packets.

`netstat` can also be set to run continuously. In Windows, run `netstat -nn`, where `nn` is the refresh interval in seconds (press `Ctrl+C` to stop); in Linux, run `netstat -c`.



The Linux `netstat` command is part of the deprecated `net-tools` package. The preferred package `iproute2` contains a number of different commands to replace `netstat` functionality. Most of the port scanning functions are performed by `ss`, while interface statistics are reported by `nstat`.

Common TCP and UDP Ports

The following table lists some of the well-known and registered port numbers.

Port Number	Transport Protocol	Service or Application	Description
20	TCP	ftp-data	File Transfer Protocol—Data
21	TCP	ftp	File Transfer Protocol—Control
22	TCP	ssh/sftp	Secure Shell/FTP over SSH
23	TCP	telnet	Telnet
25	TCP	smtp	Simple Mail Transfer Protocol
53	TCP/UDP	domain	Domain Name System
67	UDP	bootps	BOOTP/DHCP Server
68	UDP	bootpc	BOOTP/DHCP Client
69	UDP	tftp	Trivial File Transfer Protocol
80	TCP	http	HTTP
110	TCP	pop	Post Office Protocol
123	UDP	ntp/sntp	Network Time Protocol/Simple NTP
143	TCP	imap	Internet Message Access Protocol
161	UDP	snmp	Simple Network Management Protocol
162	UDP	snmp-trap	Simple Network Management Protocol Trap
389	TCP/UDP	ldap	Lightweight Directory Access Protocol
443	TCP	https	HTTP-Secure (Secure Sockets Layer (SSL)/Transport Layer Security (TLS))

Port Number	Transport Protocol	Service or Application	Description
445	TCP	smb	Server Message Block over TCP/IP
514	UDP	syslog	Syslog
546	UDP	dhcpcv6-client	DHCPv6 Client
547	TCP	dhcpcv6-server	DHCPv6 Server
587	TCP	smtps	SMTP-Secure
636	TCP	ldaps	LDAP-Secure
993	TCP	imaps	IMAP-Secure
995	TCP	pop3s	POP3-Secure
1433	TCP	sql-server	MS Structured Query Language (SQL) Server
1521	TCP	sqlnet	Oracle SQL*Net
3306	TCP	mysql	MySQL/MariaDB
3389	TCP	rdp	Remote Desktop Protocol
5004	UDP	rtp	Real-Time Protocol
5005	UDP	rtcp	Real-Time Control Protocol
5060	TCP/UDP	sip	Session Initiation Protocol
5061	TCP/UDP	sips	SIP-Secure

Lesson 6.2

Dynamic Host Configuration Protocol



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.4 Given a scenario, implement IPv4 and IPv6 network services.

Every host interface needs an IP configuration to communicate on a TCP/IP network. An administrator can manually assign these IP addresses, or the assignment can be done automatically. By understanding the different methods available to you for assigning IP addresses, you can choose the method that best suits different networks and hosts.


As you study this lesson, answer the following questions:

- What is a DHCP scope?
- What type of configuration parameters can be delivered using DHCP?
- What type of devices can be used as a DHCP server?
- What are the advantages and disadvantages of static and reserved IP address assignments?

DHCP Process

The **Dynamic Host Configuration Protocol (DHCP)** provides an automatic method for allocating an IP address, subnet mask, and optional parameters, such as the default gateway and DNS server addresses, when a host joins the network.

A host is configured to use DHCP by specifying in the TCP/IP configuration that it should automatically obtain an IP address.



DHCP Discover, Offer, Request, Ack process. (Images © 123RF.com.)

1. When a DHCP client initializes, it broadcasts a DHCPDISCOVER packet to find a DHCP server. All communications are sent using UDP, with the server listening on port 67 and the client on port 68.
2. Presuming it has an IP address available, the DHCP server responds to the client with a DHCPOFFER packet, containing the address and other configuration information.
3. The client may choose to accept the offer using a DHCPREQUEST packet—also broadcast onto the network.
4. Assuming the offer is still available, the server will respond with a DHCPACK packet. The client broadcasts an ARP message to check that the address is unused. If so, it will start to use the address and options; if not, it declines the address and requests a new one.

The IP address is leased by the server for a limited period only. A client can attempt to renew or rebind the lease before it expires. If the lease cannot be renewed, the client must release the IP address and start the discovery process again.



Sometimes, the DHCP lease process is called the DORA process: Discover, Offer, Request, and Ack(nowledge).

DHCP Server Configuration

A DHCP server must be allocated a static IP address and configured with a range (or pool) of IP addresses and subnet masks plus option values to allocate.

DHCP Settings		
DHCP Server:	☐ Auto ☒ On ☐ Off	
IP Address Pool:	192.168.1.100	- 192.168.1.199
Address Lease Time:	720	(1-2880) minutes
Default Gateway:	192.168.1.254	(Optional)
Primary DNS:	192.168.1.254	(Optional)
Secondary DNS:	0.0.0.0	(Optional)

Firmware Version: 1.0.0 Build 20160909 Rel. 53890 Hardware Version: AP300 v1.0 [Support](#)

*Configuring DHCP on a TP-LINK wireless access point.
(Screenshot courtesy of TP-Link Technologies Co., Ltd.)*

A range of addresses and options configured for a single subnet is referred to as a **scope**. To define a scope, you must provide a start and end IP address along with a subnet mask. The server maintains a one-to-one mapping of scopes to subnets. That is, no scope can cover more than one subnet, and no subnet can contain more than one scope.

The multifunction device shown only supports a single scope. The DHCP server must be placed in the same subnet as its clients. More advanced DHCP servers might be configured to manage multiple scopes. Where a server provides IP configuration for multiple subnets/scopes, it must choose the pool to service each request based on the subnet from which the request originated.



There is no mechanism for a client to choose between multiple servers. Therefore, if multiple DHCP servers are deployed—for fault tolerance, for instance—they must either be configured with non-overlapping split scopes or use a failover mechanism. DHCP for multiple subnets is usually handled by configuring relay agents to forward requests to a central DHCP server.

DHCP Options

Along with an address scope, you also need to define other parameters, such as lease time and options.

DHCP Lease Time and Available Leases

The client can renew the lease when at least half the lease's period has elapsed (T1 timer) so that it keeps the same IP addressing information. If the original DHCP server does not respond to the request to renew the lease, the client attempts to rebind the same lease configuration with any available DHCP server. By default, this happens after 87.5% of the lease duration is up (T2 timer). If this fails, the client releases the IP address and continues to broadcast to discover a server.

A long lease time means the client does not have to renew the lease often, but the DHCP server's available pool of IP addresses is not replenished frequently. Where IP addresses are in short supply, a short lease period enables the DHCP server to allocate addresses previously assigned to hosts that are now not active on the network.



A Windows client can be forced to release a lease by issuing a command such as `ipconfig`. In Linux, the utility `dhclient` is often used for this task, though modern distributions might use `NetworkManager` or `systemd-networkd`.

DHCP Options

When the DHCP server offers a configuration to a client, at a minimum it must supply an IP address and subnet mask. Typically, it will also supply other IP-related settings, known as **DHCP options**. Each option is identified by a tag byte or decimal value between 0 and 255 (though neither 0 nor 255 can be used as option values). Some widely used options include the following:

- The default gateway (IP address of the router).
- The IP address(es) of DNS servers that can act as resolvers for name queries.
- The DNS suffix (domain name) to be used by the client.
- Other useful server options, such as time synchronization (NTP), file transfer (TFTP), or VoIP proxy.

A set of default (global) options can be configured on a server-wide basis. Default options can be overridden by setting scope-specific options.

DHCP Reservations and Exclusions

One disadvantage of the standard dynamic assignment method is that it does not guarantee that any given client will retain the same IP address over time. There are some cases where it would be advantageous for certain hosts, such as network printers or wireless access points, to retain their IP addresses.

One solution is to configure static assignments, using IP addresses outside the DHCP scope. Alternatively, statically assigned addresses can be assigned from a specially configured exclusion range if this is supported by the server. While these solutions are functional, they lose the advantages of centralized configuration management.

An alternative approach is to create a **reservation**. A reservation is a mapping of a MAC address or interface ID to a specific IP address within the DHCP server's address pool. When the DHCP server receives a request from the given interface, it always provides the same IP address. This is also referred to as static or fixed address assignment. An automatically allocated reservation refers to an address that is leased permanently to a client. This is distinct from static allocation as the administrator does not predetermine which specific IP address will be leased.

Lesson 6.3

APIPA and SLAAC



Exam Objectives Covered

- 1.7 Given a scenario, use appropriate IPv4 network addressing.
- 3.4 Given a scenario, implement IPv4 and IPv6 network services.

Not all networks have DHCP servers. Also, the DHCP process can sometimes fail to work properly. There are various methods for hosts to perform autoconfiguration when DHCP is not available. You should also understand how IPv6 networks use autoconfiguration and appreciate differences in the way DHCPv6 functions, compared to DHCPv4.

As you study this lesson, answer the following questions:

- What is the purpose of Automatic Private IP Addressing (APIPA)?
- How do you know if a host is using an APIPA address?
- Which IP configuration parameters are set when APIPA is used? Which parameters are not set?
- How does address autoconfiguration work on IPv6 networks?
- What configuration parameters can be supplied using DHCPv6?

Automatic Private IP Addressing

A host's IP configuration can either be applied statically, or it can use an autoconfiguration method. Autoconfiguration on an IPv4 network usually means using a Dynamic Host Configuration Protocol (DHCP) server.

Automatic Private IP Addressing (APIPA) was developed by Microsoft as a means for clients that could not contact a DHCP server to communicate on the local network anyway. If a Windows host does not receive a response from a DHCP server within a given time frame, it selects an address at random from the range 169.254.1.1 to 169.254.254.254.



These addresses are from one of the address ranges reserved for private addressing (169.254.0.0/16). The first and last subnets are supposed to be unused.

This type of addressing is referred to as link local in standards documentation (RFC 3927).

APIPA has no mechanism for assigning default gateway or DNS server addresses. Hosts using APIPA are restricted to communicating on the local network.

IPv6 Interface Autoconfiguration and Testing

In IPv6, an interface must always be configured with a link local address. One or more routable addresses can be assigned to the interface in addition to the link local address. As with IPv4, you can either assign a routable IPv6 address statically or use an automatic addressing scheme. Static address configuration would generally be reserved to routers and possibly some types of servers.

Neighbor Discovery Protocol and Router Advertisements

The Neighbor Discovery (ND) Protocol performs some of the functions on an IPv6 network that ARP and ICMP perform under IPv4. The main functions of ND are as follows:

- **Address autoconfiguration**—Enables a host to configure IPv6 addresses for its interfaces automatically and detect whether an address is already in use on the local network, by using neighbor solicitation (NS) and neighbor advertisement (NA) messages.
- **Prefix discovery**—Enables a host to discover the known network prefixes that have been allocated to the local segment. This facilitates next-hop determination (whether a packet should be addressed to a local host or a router). Prefix discovery uses router solicitation (RS) and **router advertisement (RA)** messages. An RA contains information about the network prefix(es) served by the router, information about autoconfiguration options, plus information about link parameters, such as the MTU and hop limit. Routers send RAs periodically and in response to a router solicitation initiated by the host.
- **Local address resolution**—Allows a host to discover other nodes and routers on the local network (neighbors). This process also uses neighbor solicitation (NS) and neighbor advertisement (NA) messages.
- **Redirection**—Enables a router to inform a host of a better route to a particular destination.

Stateless Address Autoconfiguration

IPv4 has a system for generating link local addresses, but these are not routable outside the local network. Consequently, IPv4 depends heavily on the Dynamic Host Configuration Protocol (DHCP) for address autoconfiguration. IPv6 uses a more flexible system of address autoconfiguration called **stateless address autoconfiguration (SLAAC)**:

- The host generates a link local address and uses Neighbor Discovery (ND) messages to test that it is unique.
- The host listens for a router advertisement (RA) or transmits a router solicitation (RS) using ND protocol messaging. The router can either provide a network prefix, direct the host to a DHCPv6 server to perform stateful autoconfiguration, or perform some combination of stateless and stateful configuration.

ICMPv6


IPv6 uses an updated version of ICMP. The key new features are the following:

- **Error messaging**—ICMPv6 supports the same sort of destination unreachable and time exceeded messaging as ICMPv4. One change is the introduction of a Packet Too Big class of error. Under IPv6, routers are no longer responsible for packet fragmentation and reassembly, so the host must ensure that they fit in the MTUs of the various links used.
- **Informational messaging**—ICMPv6 supports ICMPv4 functions, such as echo and redirect, plus a whole new class of messages designed to support ND and MLD, such as router and neighbor advertisements and solicitations.

DHCPv6 Server Configuration


IPv6's Stateless Address Autoconfiguration (SLAAC) process can locate the default gateway and generate a host address with a suitable network prefix automatically. In this context, the role of a DHCP server in IPv6 is different. DHCPv6 is often just used to provide additional option settings, rather than leases for host IP addresses. The format of messages is different, but the process of DHCP server discovery and address leasing (if offered) is fundamentally the same. As IPv6 does not support broadcast, clients use the multicast address `ff02::1:2` to discover a DHCP server. DHCPv6 uses ports 546 (clients) and 547 (servers), rather than ports 68 and 67 as in DHCPv4.

In stateless mode, a client obtains a network prefix from a router advertisement and uses it with the appropriate interface ID. The router can also set a combination of flags to tell the client that a DHCP server is available. If so configured, the client solicits a DHCPv6 server using the multicast address `ff02::1:2` and requests additional configuration information.



DHCPv6 stateless mode. (Images © 123RF.com.)

By contrast, stateful mode means that a host can also obtain a routable IP address from a DHCPv6 scope. In either mode, a DHCPv6 server can be used to supply options information, such as DNS server addresses, DNS suffix/domain lists, time servers, and so on.



DHCPv6 stateful mode. (Images © 123RF.com.)

Configuring the scope requires you to define the network prefix and then any IP addresses that are to be excluded from being offered. All other addresses that are not explicitly excluded can be offered. The host must still listen for a router advertisement to obtain the network prefix and configure a default gateway. There is no mechanism in DHCPv6 for setting the default route.



When using stateful DHCPv6, it is possible to configure a static reservation. However, a DHCPv6 reservation doesn't use a MAC address. Instead, each system generates a host DHCP Unique Identifier (DUID) plus an Identity Association Identifier (IAID) for each interface.

Lesson 6.4

DHCP Relay and Troubleshooting



Exam Objectives Covered

3.4 Given a scenario, implement IPv4 and IPv6 network services.

5.3 Given a scenario, troubleshoot common issues with network services.

Automatic addressing is a critical service for most types of networks. Where a SOHO network might have a single DHCP server for a single subnet, enterprises must support multiple subnets. This lesson will help you to implement and troubleshoot DHCP services in complex network environments.


As you study this section, answer the following questions:

- What is the purpose of a DHCP relay agent?
- What is the purpose of an IP helper? How does it differ from a DHCP relay agent?
- What issues arise when using DHCP services, and how can they be resolved?

DHCP Relay and IP Helper

Normally, routers do not forward broadcast traffic. This means that each broadcast domain must be served by its own DHCP server. On a large network with multiple subnets, this would mean provisioning and configuring many DHCP servers. To avoid this scenario, a **DHCP relay** agent can be configured to provide forwarding of DHCP traffic between subnets. Routers that can provide this type of forwarding are described as RFC 1542 compliant.

The DHCP relay intercepts broadcast DHCP frames, applies a unicast address for the appropriate DHCP server, and forwards them over the interface for the subnet containing the server. The DHCP server can identify the original IP subnet from the packet and offer a lease from the appropriate scope. The DHCP relay also performs the reverse process of directing responses from the server to the appropriate client subnet.



Configuring a DHCP relay agent. (Images © 123RF.com.)

This **IP helper** functionality can be configured on routers to allow set types of broadcast traffic (including DHCP) to be forwarded to an interface. The IP helper function supports the function of the DHCP relay agent. For example, in the diagram, hosts in the 10.1.20.0/24 and 10.1.10.0/24 subnets need to use a DHCP server for autoconfiguration, but the DHCP server is located in a different subnet. The router is configured as a DHCP relay agent, using the following commands to enable forwarding of DHCP broadcasts on the interfaces serving the client subnets:

```

interface eth1
ip helper-address 10.1.0.200

interface eth2
ip helper-address 10.1.0.200
  
```

UDP forwarding is a more general application of the same principle. As well as DHCP, it is used for the Network Time Protocol (NTP) and other broadcast-based applications.

DHCP Issues

A Windows host that is configured to use dynamic addressing but that fails to obtain a lease will revert to an automatic IP address (APIPA) configuration and select an address in the 169.254.0.0/16 range. Linux might use link local addressing, set the address to unknown (0.0.0.0), or leave the interface unconfigured.

Possible reasons for a client to fail to obtain a lease include the following:

- The DHCP server is offline. If your DHCP servers go offline, users will continue to connect to the network for a period and thereafter start to lose contact with network services and servers as they come to try to renew a lease.
- No more addresses available (DHCP scope exhaustion). Create a new scope with enough addresses or reduce the lease period. A shorter lease period can mitigate exhaustion issues in networks with high client turnover, such as guest Wi-Fi. IP Address Management (IPAM) software suites can be used to track address usage across a complex DHCP infrastructure.

- The router between the client and DHCP server doesn't support BOOTP forwarding. Either install RFC 1542-compliant routers or add another type of DHCP relay agent to each subnet or VLAN.

If you reconfigure your DHCP servers and their scopes, you will need to plan for the fact that not all clients' IP configurations will be updated when the server scopes are edited and could be left with an expired IP, default gateway, or DNS server address. You can mitigate this by lowering the lease duration in advance of changes, forcing all clients to renew, or running parallel settings for a period.



Also be aware that address pool exhaustion might be a symptom of a malicious attack.

Lesson 6.5

Domain Name System



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.4 Given a scenario, implement IPv4 and IPv6 network services.

Each host that has an IP address assigned to it can also have a descriptive name. This makes it easier for human users to identify and access it on the network and for application services to be configured with an addressing scheme that allows for changes in the underlying network. Almost all networks depend on this name resolution functionality to operate smoothly and securely, so it is important to understand how it works. In this topic, you will identify methods for host name resolution for TCP/IP networks.

As you study this lesson, answer the following questions:

- How do Domain Name Service (DNS) queries resolve host names and domains to IP addresses?
- What is the role of the root servers in the DNS hierarchy?
- In DNS, what is the difference between a zone and a domain?
- What is the difference between a forward lookup zone and a reverse lookup zone?
- What functions do different record types have, such as A records and PTR records?
- What methods can be used to ensure the security of name queries?

Host Names and Domain Names

The Internet Protocol uses a binary IP address to locate a host on an internetwork. The dotted decimal (IPv4) or hex (IPv6) representation of this IP address is used for configuration purposes, but it is not easy for people to remember. For this reason, a “friendly” name is also typically assigned to each host. There are two types of names: host names and fully qualified domain names (FQDNs).

A **host name** is assigned to a computer by the administrator, usually when the OS is installed. The host name needs to be unique on the local network.

To avoid the possibility of duplicate host names on the Internet, a **fully qualified domain name (FQDN)** is used to provide a unique identity for the host belonging to a particular network. An example of an FQDN might be `nut.widget.example`. An FQDN is made up of the host name and a domain suffix. In the example, the host name is `nut`, and the domain suffix is `widget.example`. This domain suffix consists of the domain name `widget` within the top-level domain (TLD) `.example`. A domain suffix could also contain subdomains between the host and domain name. The trailing dot or period represents the root of the hierarchy.



When you are configuring name records, an FQDN must include the trailing period to represent the root, but this can be omitted in most other use cases.

A domain name must be registered with a registrar to ensure that it is unique within a top-level domain. Once a domain name has been registered, it cannot be used by another organization. The same domain name may be registered within different top-level domains, however—`widget.example.` and `widget.example.uk.` are distinct domains, for instance.

Numerous hosts may exist within a single domain. For example: `nut`, `bolt`, and `washer` might all be hosts within the `widget.example.` domain. Given that, FQDNs must follow certain rules:

- The host name must be unique within the domain.
- The total length of an FQDN cannot exceed 253 characters, with each label (part of the name defined by a period) no more than 63 characters (excluding the periods).
- A DNS label should use letter, digit, and hyphen characters only. A label should not start with a hyphen. Punctuation characters such as the period (.) or forward slash (/) should not be used.
- DNS labels are not case-sensitive.


Additionally, Internet registries may have their own restrictions.

DNS Hierarchy

The **Domain Name System (DNS)** is a global hierarchy of distributed name server databases that contain information on domains and hosts within those domains. At the top of the DNS hierarchy is the root, which is represented by the null label, consisting of just a period (.). There are 13 root level servers (A to M).

Immediately below the root lie the top-level domains (TLDs). There are several types of top-level domains, but the most prevalent are generic (such as `.com`, `.org`, `.net`, `.info`, `.biz`), sponsored (such as `.gov`, `.edu`), and country code (such as `.uk`, `.ca`, `.de`). DNS is operated by ICANN ([icann.org](https://www.icann.org)), which also manages the generic TLDs. Country codes are generally managed by an organization appointed by the relevant government.

Information about a domain is found by tracing records from the root down through the hierarchy. The root DNS servers have complete information about the top-level domain servers. In turn, these servers have information relating to servers for the second level domains. No name server has complete information about all domains. Records within the DNS tell them where an authoritative name server for the missing information is found.




DNS hierarchy. (Images © 123RF.com.)

An FQDN reflects this hierarchy, from most specific on the left (the host's resource record with its name:IP address mapping) to least specific on the right (the TLD followed by the root). An example is `pc.corp.515support.com`.

Name Resolution Using DNS

The signal for the name resolution process to commence occurs when a user presents an FQDN (often within a web address) to an application program, such as a web browser. The client application, referred to as a stub resolver, checks its local cache for the mapping. If no mapping is cached, it forwards the query to its local name server. The IP addresses of one or more name servers that can act as resolvers are usually set in the TCP/IP configuration. The resolution process then takes place as follows:



DNS name resolution process. (Images © 123RF.com.)

Most queries between name servers are performed as **iterative lookups**. This means that a name server responds to a query with either the requested record or the address of a name server at a lower level in the hierarchy that is authoritative for the namespace. It makes no effort to try to make additional queries to locate information that it does not have. In the figure, at steps 4 and 5, the root server and .net name server simply pass the querying server the address of an authoritative name server. They do not take on the task of resolving the original query for `www.515web.net`.

A **recursive lookup** means that if the queried server is not authoritative, it does take on the task of querying other name servers until it finds the requested record or times out. The name servers listed in a client's TCP/IP configuration accept recursive queries. This is the type of querying performed by the `corp.515support.com` name server.




A DNS server may be configured to only perform recursive querying (a resolver), or it may perform recursive querying and maintain zone records, or it may only maintain zone records. Usually the roles are split, especially if the servers are open to the Internet. Most Internet-accessible DNS servers disable recursive queries. Recursive resolvers are typically only accessible by authorized clients—subscribers within an ISP's network or clients on a private LAN, for instance.

Resource Record Types

DNS name servers maintain the DNS namespace in zones. A single zone namespace might host records for multiple domains. Conversely, subdomains within a domain might be managed as multiple zones, possibly hosted on multiple servers.

A DNS zone will contain numerous **resource records**. These records allow a DNS name server to resolve queries for names and services hosted in the domain into IP addresses. Resource records can be created and updated manually (statically), or they can be generated dynamically from information received from client and server computers on the network.

The Start of Authority (SOA) record identifies the primary **authoritative name server** that maintains complete resource records for the zone. The primary name server can be used to modify resource records. The SOA also includes contact information for the zone and a serial number for version control.



Configuring a Start of Authority record in Windows DNS.
(Screenshot courtesy of Microsoft.)

Name server (NS) records identify authoritative DNS name servers for the zone. As well as the primary name server, most zones are configured with secondary name servers for redundancy and load balancing. Secondary name servers hold read-only copies of resource records but can still be authoritative for the zone.

```


; BIND data file for 515support.com
$TTL 604800
@ IN SOA ns1.515support.com. hostmaster.515support.com. (
    2012080402 ; Serial
    3600 ; Refresh
    1800 ; Retry
    604800 ; Expire
    86400 ; Negative Cache TTL
)
IN NS ns1.515support.com.
IN A 203.0.113.33
IN MX 10 mail.515support.com.
ns1 IN A 203.0.113.33
www IN A 198.51.100.29
mail IN A 198.51.100.29
dvwa IN CNAME www

```

Resource records configured on a BIND DNS server.

Host Address and Canonical Name Records

An address (A) record is used to resolve a host name to an IPv4 address. An AAAA record resolves a host name to an IPv6 address.



Name	Type	Data	Timestamp
BACKUP	Host (A)	10.1.16.13	4/10/2021 3:00:00 AM
BACKUP	IPv6 Host (AAAA)	fdf0:2413:6d1c:0020:7b88:f308:ff44:7991	4/10/2021 3:00:00 AM
ca	Alias (CNAME)	DC10.corp.515support.com.	static
CS01	Host (A)	10.1.16.17	4/10/2021 3:00:00 AM
CS01	IPv6 Host (AAAA)	fdf0:2413:6d1c:0020:bf79:db95:594e:018e	4/10/2021 2:00:00 AM
dc10	Host (A)	10.1.16.1	static
dc10	IPv6 Host (AAAA)	fdf0:2413:6d1c:0020:0000:0000:0000:0001	static
dc10	IPv6 Host (AAAA)	fdf0:2413:6d1c:0020:41f3:89dd:42ad:92ae	static
git	Host (A)	10.1.16.9	8/4/2021 2:00:00 AM
mail	Host (A)	10.1.16.2	static

Both types of host records (A and AAAA) plus a CNAME record in Windows Server DNS. (Screenshot courtesy of Microsoft.)



DNS uses the UDP transport protocol over port 53 by default, and UDP has a maximum packet size of 512 bytes. Due to the larger address sizes of IPv6, AAAA records can exceed this limit. This can result in UDP packets being fragmented into several smaller packets. This can result in these packets being blocked by firewalls if they are not configured to expect them. Network administrators should check that their DNS servers can accept these transmissions and that intermediary components are not blocking them.

A canonical name (CNAME) (or alias) record is used to configure an alias for an existing address record (A or AAAA). For example, the IP address of a web server with the host record `lamp` could also be resolved by the alias `www`. CNAME records are also often used to make DNS administration easier. For example, an alias can be redirected to a completely different host temporarily during system maintenance.

Multiple different-named resource records can refer to the same IP address (and vice versa in the case of load balancing).



There is nothing to stop an administrator configuring multiple address records to point different host names to the same IP address. Using CNAME records is usually considered better practice, however. It is also possible to configure multiple A or AAAA records with the same host name but different IP addresses. This is usually done as a basic load balancing technique, referred to as round robin DNS.




A name server can be configured to allow automatic creation, updating, and deletion of host records using Dynamic DNS (DDNS). DDNS allows a client or DHCP server to configure records, rather than requiring the DNS server administrator to create and update them manually. In Windows, running `ipconfig /registerdns` causes a client to attempt to use DDNS.

Mail Exchange, Service, and Text Records

A mail exchange (MX) record is used to identify an email server for the domain. In a typical network, multiple servers are installed to provide redundancy, and each one will be represented with an MX record. Each server record is given a preference value with the lowest numbered entry preferred. The host identified in an MX record must have an associated A or AAAA record. An MX record must not point to a CNAME record.

While most DNS records are used to resolve a name into an IP address, a Service (SRV) record contains the service name and port on which a particular application is hosted. SRV records are often used to locate VoIP or media servers. SRV records are also an essential part of the infrastructure supporting Microsoft's Active Directory; they are used by clients to locate domain controllers, for instance. As with MX, SRV records can be configured with a priority value.



The screenshot shows the Windows DNS Manager console. On the left, the tree view is expanded to 'Forward Lookup Zones' > 'corp.515support.com'. The right pane displays a list of SRV records for this zone. The records are as follows:


Name	Type	Data	Timestamp
_gc	Service Location (SRV)	[0][100][3268] dc10.corp.515support.com.	8/3/2021 11:00:00 PM
_kerberos	Service Location (SRV)	[0][100][88] dc10.corp.515support.com.	8/3/2021 11:00:00 PM
_kpasswd	Service Location (SRV)	[0][100][464] dc10.corp.515support.com.	8/3/2021 11:00:00 PM
_ldap	Service Location (SRV)	[0][100][389] dc10.corp.515support.com.	8/3/2021 11:00:00 PM

SRV records in Windows Server DNS. (Screenshot courtesy of Microsoft.)

A TXT record is used to store any free-form text that may be needed to support other network services. A single domain name may have many TXT records, but most commonly they are used as part of Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM). An SPF record is used to list the IP addresses or names of servers that are permitted to send email from a particular domain and is used to combat the sending of spam. DKIM records are used to decide whether you should allow received email from a given source, preventing spam and mail spoofing. DKIM can use encrypted signatures to prove that a message really originated from the domain it claims.

Pointer Records

A DNS server may have two types of zones: forward lookup and reverse lookup. Forward lookup zones contain the resource records listed previously. For example, given a name record, a forward lookup returns an IP address; an MX record returns a host record associated with the domain's mail services. Conversely, a **reverse DNS** query returns the host name associated with a given IP address. This information is stored in a reverse lookup zone as a pointer (PTR) record.



Reverse lookup zone and pointer records in Windows Server DNS.
(Screenshot courtesy of Microsoft.)

Reverse DNS querying uses a special domain named by the first three octets of IP addresses in the zone in reverse order and appended with `in-addr.arpa`. The name server is configured with a reverse lookup zone. This zone contains PTR records consisting of the final octet of each host record. For example, the reverse lookup for a host record containing the IP address 198.51.100.1 is:

`1.100.51.198.in-addr.arpa`

IPv6 uses the `ip6.arpa` domain; each of the 32 hex characters in the IPv6 address is expressed in reverse order as a subdomain. For example, the IPv6 address:

`2001:0db8:0000:0000:0bcd:abcd:ef12:1234`

is represented by the following pointer record:

`4.3.2.1.2.1.f.e.d.c.b.a.d.c.b.0.0.0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa`

DNS Server Configuration

DNS is essential to the function of the Internet. Windows Active Directory and most Linux networks also require a DNS service to be running and correctly configured. It is important to realize that there are different kinds of DNS servers however, fulfilling different roles in network architecture.

DNS Server Types

A DNS server is usually configured to listen for queries on UDP port 53. Some DNS servers are also configured to allow connections over TCP port 53, as this allows larger record transfers (over 512 bytes). Larger transfers might be required if IPv6 is deployed on the network or if the DNS servers are using a security protocol (DNSSEC).

A name server can maintain primary and/or secondary zones:

- Primary means that the zone records held on the server are editable. A zone can be hosted by multiple primary servers for redundancy. As the zone records are editable on all primaries, changes must be carefully replicated and synchronized. It is critically important to update the serial number for each change.

- Secondary means that the server holds a read-only copy of the zone. This is maintained through a process of replication known as a **zone transfer** from a primary name server. A secondary zone would typically be provided on two or more separate servers to provide fault tolerance and load balancing. Again, the serial number is a critical part of the zone transfer process.

A name server that holds complete records for a domain can be defined as authoritative. This means that a record in the zone identifies the server as a name server for that namespace. Both primary and secondary name servers are authoritative.

Servers that don't maintain a zone (primary or secondary) are referred to as cache-only servers. A non-authoritative answer from a server is one that derives from a cached record, rather than directly from the zone records.

DNS Caching

Each resource record can be configured with a default **time to live (TTL)** value, measured in seconds. This value instructs resolvers how long a query result can be kept in cache. Setting a low TTL allows records to be updated more quickly but increases load on the server and latency on client connections to services. Some common TTL values include 300 (five minutes), 3,600 (one hour), 86,400 (one day), and 604,800 (one week).

DNS caching is performed by both servers and client computers. In fact, each application on a client computer might be configured to manage its own DNS cache. For example, separate web browser applications typically maintain their own caches rather than relying on a shared OS cache.

If there is a change to a resource record, server and client caching means that the updated record can be relatively slow to propagate around the Internet. These changes need to be managed carefully to avoid causing outages. Planning for a record change involves reducing the TTL in the period before the change, waiting for this change to propagate before updating the record, and then reverting to the original TTL value when the update has safely propagated.

Internal Versus External DNS

As well as making sure that resource records for the managed domain(s) are accurate, administrators should ensure that DNS services are highly available and secure, to prevent DNS spoofing, where an attacker is able to supply false name resolutions to clients.

A company will use primary and secondary name servers to maintain authoritative zone records for the domains that it manages. Internal DNS zones refer to the domains used on the private network only. These name records should only be available to internal clients. For example, a company might run a Windows Active Directory network using the domain name `corp.515support.com`. The zone records for the subdomain `corp.515support.com` would be served from internal name servers. This would allow a client PC (`pc1.corp.515support.com`) to contact a local application server (`crm.corp.515support.com`). The name servers hosting these internal subdomain records must not be accessible from the Internet.

External DNS zones refer to records that Internet clients must be able to access. For example, the company might run web and email services on the domain `515support.com`. In order for Internet hosts to use a web server at `www.515support.com` or send email to an `@515support.com` address, the zone records for `515support.com` must be hosted on a name server that is accessible over the Internet.

Companies must also provide name resolution services to support their internal clients contacting other domains. The function of a resolver is to perform recursive queries in response to requests from client systems (stub resolvers). If a name server is not authoritative for the requested domain, it can either perform a recursive query to locate an authoritative name server, or it can forward the request to another name server. A recursive resolver must be configured with a root hints file so that it can query the whole DNS hierarchy from the root servers down. DNS servers should allow recursive queries only from authorized internal clients. It is also a good idea to separate the DNS servers used to host zone records from ones used to service client requests for non-authoritative domains.



It is possible for the same DNS server instance to perform in both name server and resolver roles, but more typically these functions are separated to different servers for security reasons.

As an alternative to recursion (or to supplement it), name servers can be configured to resolve queries via forwarding. A forwarder transmits a client query to another DNS server and routes the replies it gets back to the client. A conditional forwarder performs this task for certain domains only. For example, you might configure a DNS server that is authoritative for the local private network (internal DNS), but that forwards any requests for Internet domains to an external DNS resolver run by your ISP.

DNS Security

DNS is a critical service that should be configured to resist spoofing and poisoning attacks. These attacks mean that a threat actor changes the record returned by a DNS query to point to a different IP address, potentially redirecting the victim machine to connect to a malicious host.

DNS Security Extensions

DNS Security Extensions (DNSSEC) help to mitigate against spoofing and poisoning attacks on DNS servers by providing a validation process for DNS responses. With DNSSEC enabled, the authoritative server for the zone creates a “package” of resource records (called an RRset) signed with a private key (the Zone Signing Key). When another server requests a secure record exchange, the authoritative server returns the package along with its public key, which can be used to verify the signature.

The public Zone Signing Key is itself signed with a separate Key Signing Key. Separate keys are used so that if there is some sort of compromise of the Zone Signing Key, the domain can continue to operate securely by revoking the compromised key and issuing a new one.

Name	Type	Data	Timestamp
o7i67bgdrlr78a8qs1de2n30bkvc3v	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][...	static
ofajigu1o4cvl6kuevutfc6jqt9idvson	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
ofajigu1o4cvl6kuevutfc6jqt9idvson	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][...	static
q24qcu9n11ei5qb07reuf57887rtaal	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
q24qcu9n11ei5qb07reuf57887rtaal	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][r...	static
rpbrskks6g83jrmfficimqig04a4f64c	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
rpbrskks6g83jrmfficimqig04a4f64c	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][t...	static
ts2r500ntarmmtkmkfqtp8jnodsks1r	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
ts2r500ntarmmtkmkfqtp8jnodsks1r	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][...	static
u7k7b9gskb4c0pot5u83te7k7vrmgs6	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
u7k7b9gskb4c0pot5u83te7k7vrmgs6	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][v...	static
updates	Alias (CNAME)	WIN2016-MS.classroom.local	static
updates	RR Signature (RRSIG)	[CNAME][Inception(UTC): 9/10/2017 2:51:21 PM][...	static
v7gcrlejbjopg1vbei3rgv873lc61hb8	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
v7gcrlejbjopg1vbei3rgv873lc61hb8	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][v...	static
vp0393g6eupsgcdfp6cg5n4h51e9ks3	RR Signature (RRSIG)	[NSEC3][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
vp0393g6eupsgcdfp6cg5n4h51e9ks3	Next Secure 3 (NSEC3)	[SHA-1][NO Opt-Out][50][4E6B4B560BC1EA8C][0...	static
WIN07-WS	Host (A)	10.1.0.101	9/8/2017 6:00:00 AM
WIN07-WS	RR Signature (RRSIG)	[A][Inception(UTC): 9/10/2017 2:51:21 PM][Expira...	static
WIN10-WS	Host (A)	10.1.0.100	9/8/2017 6:00:00 AM
WIN10-WS	RR Signature (RRSIG)	[A][Inception(UTC): 9/10/2017 2:51:21 PM][Expira...	static
win2016-dc	Host (A)	10.1.0.1	static
win2016-dc	IPv6 Host (AAAA)	fdab:cdef:0000:0001:0000:0000:0000:0001	static
win2016-dc	RR Signature (RRSIG)	[A][Inception(UTC): 9/10/2017 2:51:21 PM][Expira...	static
win2016-dc	RR Signature (RRSIG)	[AAAA][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
WIN2016-MS	Host (A)	10.1.0.2	9/8/2017 5:00:00 AM
WIN2016-MS	IPv6 Host (AAAA)	fdab:cdef:0000:0001:0000:0000:0000:0002	9/8/2017 5:00:00 AM
WIN2016-MS	RR Signature (RRSIG)	[A][Inception(UTC): 9/10/2017 2:51:21 PM][Expira...	static
WIN2016-MS	RR Signature (RRSIG)	[AAAA][Inception(UTC): 9/10/2017 2:51:21 PM][E...	static
www	Host (A)	10.1.0.10	static
www	RR Signature (RRSIG)	[A][Inception(UTC): 9/10/2017 2:51:21 PM][Expira...	static

*Windows Server DNS services with DNSSEC enabled.
(Screenshot used with permission from Microsoft.)*

The Key Signing Key for a particular domain is validated by the parent domain or host ISP. The top-level domain trusts are validated by the Regional Internet Registries, and the DNS root servers are self-validated, using a type of M-of-N control group key signing. This establishes a chain of trust from the root servers down to any particular subdomain.

DNS Client Security

Where DNSSEC validates the records held by a name server, DNS client security uses transport encryption to prevent an on-path threat actor tampering with responses to DNS queries. There are two main protocols for securing DNS queries:

- **DNS over transport layer security (DoT)**—This uses Transport Layer Security (TLS) to validate the resolver name server's digital certificate. This mitigates the risk of a threat actor using a rogue DNS server to spoof the legitimate one. If the client trusts the certificate, the subsequent DNS traffic will be encrypted. DoT works over TCP port 853.
- **DNS over hypertext transfer protocol secure (DoH)**—This also validates the resolver certificate and encrypts the DNS traffic but does so by encapsulating it within HTTP Secure packets. This uses the HTTPS standard port TCP/443, which completely disguises the fact that the client is making DNS queries. The downside is that the additional HTTP headers add overhead to each query and response.



As well as protecting against malicious attacks, DoH and DoT provide better privacy. Plain text queries can be read by anyone operating a network appliance in the path between the client and resolver. Encrypting the queries and responses prevents this type of snooping. Conversely, administrators of a corporate network need to ensure that clients use authorized resolvers and will often prefer to monitor DNS traffic.

Lesson 6.6

DNS Troubleshooting



Exam Objectives Covered

3.4 Given a scenario, implement IPv4 and IPv6 network services.

5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

The Domain Name System (DNS) is critical for locating services and hosts on the Internet and on corporate networks. SOHO and enterprise Linux and Windows systems usually rely on DNS server infrastructure for name resolution and service discovery. In the absence of DNS servers, network client machines will be unable to log on or connect to services or servers. DNS problems can also affect external websites and services. As a network technician, you will often be called upon to troubleshoot issues with name resolution.

As you study this lesson, answer the following questions:

- What are the symptoms of name resolution problems?
- What is the role of the HOSTS file in the name resolution process?
- What is the difference between nslookup and dig?

Client DNS Issues

A DNS issue is typically indicated when a host can ping a server by its IP address, but not by its name.

When a host receives a client request to access a name and it does not have the IP mapping cached, it asks a name server configured as a resolver to perform the lookup and return the IP address. As name resolution is a critical service, most hosts are configured with primary and secondary name server resolvers for redundancy. The server addresses are entered as IPv4 and IPv6 addresses. In a majority of cases, these addresses are likely to be autoconfigured via DHCP.

If a single client is unable to resolve names, the issue is likely to lie with that client's configuration. In Windows, you can view the name servers configured as resolvers using `ipconfig /all`. In Linux, the DNS server addresses are recorded in `/etc/resolv.conf`. Typically, a package such as NetworkManager or systemd-networkd would add the entries. Entries added directly will be overwritten at reboot.

If a host cannot resolve names, check that the correct name server addresses have been configured and that you can ping them. If there are configuration errors, either correct them (if the interface is statically configured) or investigate the automatic addressing server. If there are connectivity errors, check the network path between the host and its name servers.

If multiple clients are affected, the issue is likely to lie with the server service (or the way a subnet accesses the server service). Check that the server configured as a DNS resolver is online and available (that you can ping the server from the client). Bear in mind that DHCP might be configuring DNS server settings incorrectly. Check the server options or scope options configuration on the DHCP server as well.



While we are focusing on name resolution via DNS here, note that a host can use multiple methods, especially on Windows workgroup networks. Link Local Multicast Name Resolution (LLMNR) and Multicast DNS (mDNS) are modified forms of DNS that allow clients to perform name resolution on a local link without needing a server.



Hosts have a system DNS configuration, but apps such as browsers might use separately configured name servers.

Name Resolution Issues

If some DNS queries work from the client, but others don't, then the problem is more complex. When you perform a successful connectivity test by IP address, and have verified that the host's DNS resolvers are working, and the target host or service still does not respond to pinging it by name, you need to check for a fault within the name resolution process.

Name Resolution Methods

To troubleshoot name resolution, you should establish exactly how the process works on that specific host. A host can use a variety of methods to resolve a name to an IP address. In very general terms, these will be as follows:

1. Check local name caches. One complication here is that there are different types of cache and separate caches for individual applications, such as web browsers. On Windows, you can use `ipconfig/displaydns` and `ipconfig/flushdns` to monitor and clear the system's DNS cache.
2. Check HOSTS. The **HOSTS** file is a static list of host name to IP address mappings. The local resolver is likely to try to use any HOSTS file mappings first (or the mappings might be cached automatically). The default location under Windows is `%SystemRoot%\system32\drivers\etc\`, while under Linux it is usually placed in the `/etc` directory. In most cases, HOSTS should not contain any entries (other than the loopback address). Any static entries in HOSTS could be the cause of a name resolution issue. The file can also be used for troubleshooting.



Any text preceded by the `#` symbol in a HOSTS file is a comment and will not be processed. To verify a name resolution problem, edit the HOSTS file and place the correct name and IP address record in the file for the test host. When you ping that name, if that is successful, it suggests a name resolution service problem.

3. Verify DNS records using the `nslookup` or `dig` tools. There might be some discrepancy between the records returned by the resolver compared to the records configured on the authoritative DNS server that maintains the zone.

Use the `nslookup` or `dig` utilities to check what records are returned by the resolver. If trying to connect to an Internet resource, compare these records to those returned by public resolvers (such as Google's servers at 8.8.8.8). Consider whether clients have cached a record that has been changed recently. Reconfiguration of DNS records should be planned and implemented carefully to avoid caching problems.

nslookup

Name resolution troubleshooting typically involves testing multiple clients and servers. The use of caching and the distributed nature of the system means that configuration errors can occur in several different places.

You might start investigating a name resolution issue by verifying the name configured on a host. In Windows, you can use the command `ipconfig /all` to display the FQDN of the local host. In Linux, you can use the command `hostname --fqdn`.



On a local network, each host is normally configured with a DNS suffix. For example, PC1 might be configured as part of a Windows network with the suffix `ad.example.local`. If this suffix is not set correctly, some name queries could fail.

You can troubleshoot DNS name resolution with the **nslookup command**:

```
nslookup Option Host DNSServer
```

Host can be either a host name, domain name, FQDN, or IP address. *DNSServer* is the IP address of a server used to resolve the query; the default DNS server is used if this argument is omitted. *Option* specifies an nslookup subcommand. For example, the following command queries Google's public DNS server (8.8.8.8) for information about 515support.com's mail records:

```
nslookup -type=mx 515support.com 8.8.8.8
```

If nslookup is run without any arguments (or by specifying the server only with `nslookup – DNSServer`), the tool is started in interactive mode. You can perform specific query types and output the result to a text file for analysis.

```
C:\Users\Admin>nslookup -type=mx comptia.org 8.8.8.8
Server: dns.google
Address: 8.8.8.8

Non-authoritative answer:
comptia.org      MX preference = 10, mail exchanger = comptia-org.mail.protection.outlook.com

C:\Users\Admin>nslookup -type=ns comptia.org 8.8.8.8
Server: dns.google
Address: 8.8.8.8

Non-authoritative answer:
comptia.org      nameserver = ns2.comptia.org
comptia.org      nameserver = ns1.comptia.org

C:\Users\Admin>nslookup -type=mx comptia.org ns1.comptia.org
Server: UnKnown
Address: 209.117.62.56

comptia.org      MX preference = 10, mail exchanger = comptia-org.mail.protection.outlook.com
C:\Users\Admin>
```

The first two nslookup commands identify comptia.org's MX and primary name server records using Google's public DNS resolver (8.8.8.8). Note that the answers are non-authoritative. The third command queries CompTIA's name server for the MX record. This answer is authoritative. (Screenshot courtesy of Microsoft.)



The Windows PowerShell environment provides a more sophisticated scripted environment that you can use to issue cmdlets to test DNS name resolution (and change DNS settings as well, if required). PowerShell provides a cmdlet called `Resolve-DnsName`, which allows a more flexible method of testing name resolution than nslookup, as it allows testing of the different methods of name resolution (HOSTS file, DNS cache, and DNS server).

dig

Domain Information Groper (dig) is a command line tool for querying DNS servers that ships with the BIND DNS server software published by the Internet Systems Consortium (ISC) (isc.org/downloads/bind).

`dig` can be run pointing at a specific DNS server; otherwise, it will use the default resolver. Without any specific settings, it queries the DNS root zone. A simple query uses the syntax: `dig host`. This will search for the address record for the host, domain, or FQDN or PTR record for an IP address.

The following command example directs the resolve request to the specific DNS server identified after the `@` symbol. This can be an FQDN or IP address.

```
dig @ns1.isp.example host
```

Other examples of `dig` are to display all the resource records about a domain or just specific ones such as Mail Exchange:

```
dig @ns1.isp.example host all
```

```
dig @ns1.isp.example host MX
```

`dig` often generates a lot of information, so it is possible to add parameters to the end of the command like `+nocmd` or `+nostats`, which will reduce the output.

```
toor@LX20D:~$ dig +nocmd +noedns +noquestion 515support.com MX
;; Got answer:
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 12449
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; ANSWER SECTION:
515support.com.      6968      IN        MX       10 mail.515support.com.

;; Query time: 0 msec
;; SERVER: 127.0.0.53#53(127.0.0.53)
;; WHEN: Wed Aug 04 10:08:20 BST 2021
;; MSG SIZE rcvd: 53

toor@LX20D:~$ dig +nocmd +noedns +noquestion mail.515support.com
;; Got answer:
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 57840
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; ANSWER SECTION:
mail.515support.com. 7029      IN        A        198.51.100.29

;; Query time: 0 msec
;; SERVER: 127.0.0.53#53(127.0.0.53)
;; WHEN: Wed Aug 04 10:08:32 BST 2021
;; MSG SIZE rcvd: 53
```

Using dig to locate MX and A records.



You can use `dig` on Windows by downloading the BIND DNS server package and installing it using the tools-only option.

Module 6

Summary

You should be able to compare and contrast appropriate uses of TCP and UDP and select appropriate tools to support and troubleshoot Transport layer issues.

You should be able to explain the uses and purposes of the network services protocols DHCP and DNS.

Guidelines for Supporting Network Services

Follow these guidelines to make effective use of network addressing services:

- Identify required ports for each host and record this information in configuration management documentation. Ensure that only legitimate applications and services can bind to a server port.
- Understand the use of handshakes and acknowledgments to support and troubleshoot reliable transport mechanisms using TCP.
- Understand that applications may use UDP for unreliable unicast, multicast, or broadcast transmissions to minimize protocol overheads.
- Configure secure DHCP and DNS servers and ensure that all network hosts can contact them, using DHCP relay where appropriate.
- Ensure DHCP servers are configured with accurate IP, default gateway, and DNS server parameters for the scopes/subnets that they serve.
- Configure reservations or static assignments for hosts that need to be allocated a consistent IP address.
- If the address pool is limited, use short lease times to prevent address exhaustion.
- Understand the role of Neighbor Discovery and router advertisements in IPv6 address autoconfiguration.
- Set up primary and secondary name servers to host records for your LAN. These name services should be accessible only by authorized clients.
- Configure the appropriate host, MX, and service records for the forward lookup zone on the primary server.
- Optionally, configure a reverse lookup zone to allow clients to resolve IP addresses to host names.
- Configure the secondary server to obtain up-to-date records periodically through a zone transfer with the primary server.
- To resolve client Internet queries, set up a forwarder to pass queries to trusted resolvers on the Internet, such as your ISP's DNS server or trusted public services such as those from Google or Quad9.

- For external DNS, consider using a third-party provider, ideally with a cloud service, to ensure high availability. Without public DNS, your customers will not be able to browse your websites or send you email.
- Set up a process for checking that your external DNS records are accurate and working correctly.

Module 7

Explaining Application Services

Module Introduction

Where DHCP and DNS support basic network operations, other Application layer protocols provide platforms for user-level services, such as websites, databases, file/printer sharing, email, and voice/video calling.

You must be able to identify the ports used by these services and their performance and security requirements so that you can assist with product deployments and upgrades and perform basic troubleshooting.

Module Objectives

In this module, you will do the following:

- Explain the importance of time synchronization and the role of NTP.
- Explain the use of web, file/print, and database services.
- Explain the use of email and voice services.
- Explain how high availability services are provisioned using redundancy and load balancing.

Lesson 7.1

Application Security and Time Synchronization



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.4 Given a scenario, implement IPv4 and IPv6 network services.

Many TCP/IP application protocols have been developed to meet the different purposes of transferring files, exchanging messages, and publishing pages. However, when TCP/IP protocols were being developed, it was assumed that access to the network was sufficient security, so there were no mechanisms to authenticate services and protect data as it crossed the network. On modern networks, we have to assume that no network channel is entirely safe, so robust application security is critical. Along with security, most application protocols also require hosts to be synchronized to the same time. This is particularly important for authentication and auditing functions. This lesson will help you to explain features of the protocols that fulfill these security and synchronization functions.

As you study this lesson, answer the following questions:

- How do clients identify and authenticate application services, and how do servers protect data exchanged with clients from snooping?
- What methods are available to synchronize network hosts to the same time?

Transport Layer Security

One of the critical problems for the provision of services is that TCP/IP application protocols were originally devised without any security mechanisms. Without security, there is no authentication of the servers running the applications (or of the clients accessing them), and all data is sent in plaintext. This makes these services highly vulnerable to spoofing, eavesdropping, and unauthorized modification.

Transport Layer Security (TLS) was developed as an IETF standard to solve this issue.

TLS works as a layer between the Application and Transport layers of the TCP/IP stack, or, in OSI terms, at the Session layer. It's normally used to authenticate and encrypt TCP connections. When it is used with the HTTP application, it is referred to as HTTP Secure (HTTPS). TLS can also be used to secure other TCP application protocols, such as DNS, NTP, FTP, POP3/IMAP, SMTP, and LDAP. The secure form of the protocol typically uses a different port than the insecure version.



TLS can also be used with UDP, referred to as Datagram Transport Layer Security (DTLS), most often in virtual private networking (VPN) solutions.

To implement TLS, the server is installed with a **digital certificate** issued by some trusted **certificate authority (CA)**. When a client connects to a secure service, a TLS handshake is performed. During the handshake, the server provides its certificate to the client. The cryptographic data in the certificate proves the identity of the server, assuming that the client also trusts the CA. The certificate contains the **public key** part of a public/private encryption key pair. The **private key** is kept a secret known only to the server.

If authentication is successful, the server and client use the key pair in the digital certificate and a chosen cryptographic **cipher suite** within the TLS protocol to set up an encrypted tunnel. Even though someone else might know the public key and be in a position to record traffic passing between the server and client, they cannot decrypt the contents of the tunnel without obtaining the server's private key. This means that the communications cannot be read or changed by a third party.



The latest versions of TLS can use a mechanism called Perfect Forward Secrecy (PFS). When this is configured, not even obtaining the server's private key allows decryption of captured packets.

TLS has been developed through a number of versions, with TLSv1.3 being current at the time of writing. A server and client must be able to agree on a compatible version. As older versions can contain serious weaknesses, many servers are configured to allow only TLSv1.3 or TLSv1.2. Additionally, the client and server must be able to agree on a mutually supported cipher suite.




TLS itself was developed from an older protocol called Secure Sockets Layer (SSL). SSL is now completely obsolete.

Network Time Protocol

Many applications on networks require hosts to be synchronized to the same time. These include authentication and auditing/logging mechanisms, scheduling applications, and backup software. The **Network Time Protocol (NTP)** enables the synchronization of these time-dependent applications. NTP works over UDP on port 123.

Top-level NTP servers (stratum 1) obtain the Coordinated Universal Time (UTC) via a direct physical link to an accurate clock source, such as an atomic clock accessed over the **Global Positioning System (GPS)**. An NTP server that synchronizes its time with a stratum 1 server over a network is operating at stratum 2. Each stratum level represents a step away from the accurate clock source over a network link. These lower stratum servers act as clients of the stratum 1 servers and as servers or time sources to lower stratum NTP servers or client hosts. Most switches and routers can be configured to act as time servers to local client hosts and this function is also typically performed by network directory servers. It is best to configure each of these devices with multiple reference time sources (at least three) and to establish them as peers to allow the NTP algorithm to detect drifting or obviously incorrect time values.



Stratum 1 NTP servers are directly connected to an accurate clock source. Each stratum level below one represents a network hop away from that accurate time source. (Images © 123RF.com.)

Client hosts (application servers and workstations) usually obtain the time by using a modified form of the protocol called Simple NTP (SNTP). SNTP works over the same port as NTP. A host that supports only SNTP cannot act as a time source for other hosts. In Windows, the Time Service can be configured by using the `w32tm` command. In Linux, the `ntp` package can be configured via `/etc/ntp.conf`.

Time drift is when a system's clock begins to deviate from the source clock. NTP can use two methods to deal with time drift:

- **Slew method**—If the time is off by only a few seconds, NTP adjusts the time a few milliseconds at a time to get it back on track. Slewing is a slower, methodical method of correcting the time, but the risk of problems occurring is much less.
- **Slam method**—If the time is off by more than a few seconds and slewing will take too long, NTP will hard reset the time. While this is a quick and immediate fix, slamming can cause some programs to not function properly.

If a server or host is configured with the incorrect time, it may not be able to access network services. Authentication and other security mechanisms will often fail if the time is not synchronized on both communicating devices. In this situation, errors are likely to be generic failure or invalid token type messages. Always try to rule out time synchronization as an issue early in the troubleshooting process.



If a local stratum 1 server cannot be implemented on the local network, the time source can be configured using one or more public NTP server pools, such as time.google.com, time.windows.com, time.apple.com, time.nist.gov, or pool.ntp.org.

To mitigate risks from unauthorized time sources or manipulation of synchronization data, NTP can be protected using Transport Layer Security (TLS). **Network Time Security (NTS)** works over TCP port 4460. NTS servers may also support ports 3443 and 4443, which were specified in early draft versions of the protocol.

Precision Time Protocol

NTP is accurate enough for many network services, but not for the most timing critical application requirements. Networks supporting industrial processes, 5G cellular data, medical devices, market trading and financial services, or broadcasting use the **Precision Time Protocol (PTP)**. Where NTP can produce millisecond precision, PTP is capable of nanosecond precision. PTP can also be seen as a general replacement for NTP. It is defined in the IEEE 1588 standard.

PTP can use layer 2 messaging plus hardware clocks in compatible network adapters and switches to ensure greater levels of accuracy than NTP can support. It uses mechanisms to measure and account for delay.

PTP uses the following clock types:

- Grandmaster clock is the authoritative time source within a PTP domain.
- Boundary clock is one with interfaces in multiple PTP segments.
- Ordinary clock is one with a single PTP interface.

When two clocks are connected, one interface has a timeTransmitter role, and the other has a timeReceiver role. The grandmaster clock's interfaces are always timeTransmitter. A boundary clock would have the timeReceiver role on its interface with the grandmaster and the timeTransmitter role on other interfaces. Ordinary clock interfaces are usually timeReceiver.

Additionally, transparent clocks can be deployed. These can measure path delay and adjust P2P messages to compensate.



PTP can also be deployed as a layer 3 protocol over IP, but it will not work as accurately as a layer 2 implementation with PTP-compatible hardware-timestamping adapters and switches.

Lesson 7.2

Web, File/Print, and Database Services



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 1.4 Explain common networking ports, protocols, services, and traffic types.

So far, you have studied lower-layer services and application protocols that enable basic connectivity between nodes. Above these are the services that provide useful functions to users, such as web browsing, file/print sharing, and databases. The services that form part of the TCP/IP protocol suite are mostly client-server protocols and applications. Client-server applications are based around a centralized server that stores information and waits for requests from clients. You need a good understanding of how these protocols are used so that you can support them on your networks.

As you study this lesson, answer the following questions:


- How are secure service protocols distinguished from insecure ones?
- What specific protocols support publishing, file transfer, file/printer sharing, and database access?

Hyper Text Transfer Protocol

Websites and web applications are perhaps the most useful and ubiquitous of network services. Web technology can be deployed for a huge range of functions and applications, in no way limited to the static pages of information that characterized the first websites. The foundation of web technology is the **HyperText Transfer Protocol (HTTP)**. HTTP enables clients (typically web browsers) to request resources from an HTTP server. A client connects to the HTTP server using an appropriate TCP port (TCP/80, by default) and submits a request for a resource, using a uniform resource locator (URL). The server acknowledges the request and either responds with the data or with an error message.

HTTP Headers and Payload

The response and request formats are defined in the HTTP header. The HTTP payload is usually used to serve HyperText Markup Language (HTML) webpages, which are plain text files with coded tags describing how the page should be formatted. A web browser can interpret the tags and display the text and other resources associated with the page, such as binary picture or sound files linked to the HTML page.



Using Firefox's web developer tools to inspect the HTTP requests and response headers involved in serving a typical modern webpage. (Screenshot courtesy of Mozilla Foundation.)

HTTP also features a forms mechanism (POST) that enables a user to submit data from the client to the server. HTTP is nominally a stateless protocol; this means that the server is not required to preserve information about the client during a session. However, the basic functionality of HTTP servers is also often extended by support for scripting and programmable features (web applications). Servers can also set text file cookies to preserve session information. These coding features, plus integration with databases, increase flexibility and interactivity, but also increase the attack surface and expose more vulnerabilities.



Many argue that HTTP is a stateful protocol. Version 2 of HTTP adds more state-preserving features (blog.zamicol.com/2017/05/is-http2-stateful-protocol-application.html).

Web Servers

Most organizations have an online presence, represented by a website. In order to run a website, it must be hosted on an HTTP server connected to the Internet. Larger organizations or SMEs with the relevant expertise may host websites themselves, but more typically, an organization will lease a server or space on a server from an ISP. The following types of hosting packages are common:


- **Dedicated server**—The ISP allocates your own private server computer. This type of service is usually unmanaged (or management comes at additional cost).
- **Virtual private server (VPS)**—The ISP allocates you a virtual machine (VM) on a physical server. This is isolated from other customer instances by the hypervisor.
- **Cloud hosting**—Your website is run on a cloud over several hardware computers, allowing more scalability if demand patterns change.

- **Shared hosting**—Your website is hosted within a private directory on a shared server. Performance can be severely affected by other sites hosted on the server, because all the sites are competing for the same resources.

The main web server platforms are Apache, Microsoft Internet Information Services (IIS), and NGINX

HTTP Secure

Plaintext HTTP is highly vulnerable. A modern browser will warn users before initiating an unencrypted connection (or may refuse such connections altogether). HTTP protected by Transport Layer Security (TLS) is referred to as HTTP Secure (HTTPS). HTTPS encrypted traffic is sent over TCP port 443 (by default), rather than the open and unencrypted port 80. A web browser will open a secure session to a server providing this service by using a URL starting with `https://`, and it will also show a padlock icon in the address bar to indicate that the connection is secure. The padlock icon allows inspection of the site's security data, including the certificate authority (CA) that issued the certificate.




HTTPS padlock icon. (Screenshot courtesy of Microsoft.)

File Transfer Protocol

It is often necessary to transfer files to and from appliances or servers from a remote host. Many methods of remote file access use some form of the **File Transfer Protocol (FTP)**. While HTTPS-based web services and web applications can now offer file downloads to end users, FTP is still often used to perform the administrative upload/download of files to and from servers and appliances. For these uses, it is important to secure the FTP session.


Active Versus Passive FTP

An FTP client connects to TCP port 21 on an FTP server and opens a chosen dynamic client port number (n). The TCP port 21 control port is used to transfer commands and status information, but not for data transfer. Data transfer can operate in one of two modes: active or passive. In active mode, the client sends a `PORT` command specifying its chosen data connection port number (typically $n+1$), and the server opens the data connection between the chosen client port and TCP port 20 on the server.



FTP in active mode. (Images © 123RF.com.)

In passive mode, the client opens a data port (again, typically $n+1$) and sends the PASV command to the server's control port. The server then opens a random high port number and sends it to the client using the PORT command. The client then initiates the connection between the two ports.



FTP in passive mode. (Images © 123RF.com.)

Active FTP poses a configuration problem for some firewalls, as the server is initiating the inbound connection, but there is no way of predicting which port number will be utilized. However, not all FTP servers and clients can operate in passive mode. If this is the case, check that firewalls installed between the client and server can support active FTP (stateful inspection firewalls).




Another problem is that the control connection can remain idle when the data connection is in use, meaning that the connection can be "timed out" by the firewall (or other routing device).

Trivial File Transfer Protocol

The **Trivial File Transfer Protocol (TFTP)** is a connectionless protocol running over UDP port 69. Consequently, TFTP does not provide the guaranteed delivery offered by FTP and is only suitable for transferring small files. Also, it only supports reading (GET) and writing (PUT) files, not directory browsing, file deletion, or any of the other features of FTP. A TFTP server is most commonly used by legacy network appliances (switches, routers, diskless workstations, and printers) to download configuration files. It can also be used as a backup and restore method for configuration files. However, TFTP has no security mechanisms, and appliances are no longer as resource constrained as they were in the early days of networking. Consequently, secure protocols are now preferred for these functions.

Secure File Transfer Protocol

Secure FTP (SFTP) addresses the privacy and integrity issues of FTP by encrypting the authentication and data transfer between client and server. In SFTP, a secure link is created between the client and server using Secure Shell (SSH) over TCP port 22. Ordinary FTP commands and data transfer can then be sent over the secure link without risk of eavesdropping. This solution requires an SSH server that supports SFTP plus SFTP client software.



WinSCP SFTP client. (Screenshot courtesy of WinSCP.)

Another means of securing FTP is to use the connection security protocol SSL/TLS. There are two means of configuring FTP over TLS:

- **Explicit TLS (FTPES)**—Use the `AUTH TLS` command to upgrade an insecure connection established over TCP port 21 to a secure one. This protects authentication credentials. The data connection for the actual file transfers can also be encrypted (using the `PROT` command).
- **Implicit TLS (FTPS)**—Negotiate an SSL/TLS tunnel before the exchange of any FTP commands. This mode uses TCP port 990 for the control connection.

FTPS is tricky to configure when there are firewalls between the client and server. Consequently, FTPES is usually the preferred method.

Server Message Block

File and print services allow network clients to share access to disk and printer resources.

On a Windows network, the File/Print Sharing Service is provided by the **Server Message Block (SMB)** protocol. SMB allows a host to share its directories/files and printers to make them available for other machines to use. Support for SMB in UNIX- or Linux-based machines and network attached storage (NAS) appliances is provided by using the Samba software suite (samba.org/samba/what_is_samba.html), which allows a Windows client to access a Linux host as though it were a Windows file or print server. Samba also allows Linux hosts to access file/printer shares hosted on Windows.

On legacy networks, SMB ran as part of an older network services protocol called NetBIOS on TCP port 139. If no legacy client support is required, however, SMB is more typically run directly over TCP port 445. SMB should be restricted to use only on local networks. It is important that any traffic on the NetBIOS port ranges (137–139) and port 445 be blocked by a perimeter firewall.

SMB version 3 supports message encryption, which can be enabled on a file server or on a per-share basis. An encrypted share can only be accessed by an SMB 3.0 or higher client.




SMB has gone through several updates, with SMB3 as the current version. SMB1 has very serious security vulnerabilities and is now disabled by default on current Windows versions (docs.microsoft.com/en-us/windows-server/storage/file-server/troubleshoot/detect-enable-and-disable-smbv1-v2-v3).

Network Attached Storage

A **network attached storage (NAS)** appliance is a device dedicated to performing a file server role. A NAS uses a bespoke operating system, typically based on Linux, and will usually be operated via a web app. A NAS appliance is accessed via an IP address or domain name. Copy or backup operations are performed at file level using an application protocol, such as Server Message Block (SMB) or File Transfer Protocol (FTP).

Most NAS devices support some level of Redundant Array of Independent Disks (RAID). In a RAID array, information is spread between disks so that if one fails, the logical volume and its data will remain available.

The main drawback of NAS is that it shares bandwidth with other network applications. Adding a NAS to an already overwhelmed network increases network traffic and may result in unacceptable delays for users and applications to access data. On networks with adequate bandwidth, however, a NAS is a quick and easy way of adding shared storage.



Network attached storage. (Images © 123RF.com.)

Database Services

A database provisions information in a format that can be read and updated through some type of query language. There are two main types of databases. Relational databases store information in tables with rows (records) and columns (fields). Relationships between data fields in different tables is created using key fields that uniquely identify each record. Relational databases are operated using **Structured Query Language (SQL)**. SQL defines commands such as SELECT to retrieve information or UPDATE to change it.

SQL has been implemented in relational database management system (RDBMS) platforms by several different vendors. As well as providing an implementation of SQL, an RDBMS provides management tools and often a GUI to use to operate the database. A remote access protocol allows a client to connect to the database server over the network and allows replication traffic to move between database servers. Replication is a means of synchronizing the data held on each server. Each RDBMS uses a different TCP port to distinguish it as an application service:

- Oracle's remote data access protocol SQL*Net uses TCP/1521.
- Microsoft SQL Server uses TCP/1433.
- The open source MySQL platform uses TCP/3306. The MariaDB platform forked from MySQL uses the same port.
- The open source PostgreSQL platform uses TCP/5432.



These are the principal ports. An RDBMS is likely to use other TCP or UDP ports for additional functions.

By default, these ports are insecure. However, the RDBMS server can be installed with a certificate and configured to enable TLS transport encryption. The connection is still made over the same port. Either the server or the client can be configured to require encryption and drop the connection if a valid security profile is not available. Optionally, the client can also be installed with a certificate and the server configured to refuse connections from clients without a valid certificate.

The other type of database is referred to as NoSQL or “not only SQL.” Rather than highly structured relational tables, NoSQL data can use a variety of formats, such as key-value pairs or wide columns (where rows do not have to have the same set of fields). NoSQL databases are typically accessed using an **application programming interface (API)** over HTTPS.



All the RDBMS platforms also provide support for NoSQL datastores. There are also dedicated NoSQL platforms, such as MongoDB, Amazon DynamoDB, and CouchDB.

Lesson 7.3

Email and Voice Services



Exam Objectives Covered

1.4 Explain common networking ports, protocols, services, and traffic types.


Communication is a key function of computer networks, and there are various messaging, voice, and video services available to facilitate it. Voice and video protocols that transfer real-time data bring their own challenges for network architecture, and you need to understand these performance demands to build networks that can support them.

As you study this lesson, answer the following questions:

- What is the difference between mail transfer and mailbox access protocols?
- How does VoIP differ from traditional phone service?
- What is the difference between a hard VoIP phone and a soft VoIP phone?
- Why is quality of service (QoS) important for VoIP?

Simple Mail Transfer Protocol

Electronic mail enables a person to compose a message and send it to another user on their own network (intranet) or anywhere in the world via the Internet. Email uses separate mail transfer and mailbox access protocols:



Operation of delivery and mailbox email protocols. (Images © 123RF.com.)

The **Simple Mail Transfer Protocol (SMTP)** specifies how email is delivered from one system to another. The SMTP server of the sender discovers the IP address of the recipient SMTP server by using the domain name part of the recipient's email address. The SMTP servers for the domain are registered in DNS using mail exchange (MX) and host (A/AAAA) records.

SMTP does not queue messages indefinitely. If there is a communication problem, the SMTP server retries at regular intervals before timing out and returning a non-delivery report (NDR) to the sender. The NDR will contain an error code indicating the reason the item could not be delivered. SMTP provides no mechanism for the persistent storage of messages.

SMTP communications can be secured using TLS. This works much like HTTPS with a certificate on the SMTP server and a negotiation between client and server about which cipher suites to use. There are two ways for SMTP to use TLS:

- **STARTTLS**—This is a command that upgrades an existing insecure connection to use TLS. This is also referred to as explicit TLS or opportunistic TLS. This method is now deprecated but does remain in widespread use.
- **Implicit TLS**—This establishes the secure connection before any SMTP commands (HELO, for instance) are exchanged. Implicit TLS is now considered the preferred method.

Typical SMTP configurations use the following ports and secure services:

- **Port 25**—Used for message relay between SMTP servers, or message transfer agents (MTAs). If security is required and supported by both servers, the STARTTLS command can be used to set up the secure connection.
- **Port 465**—Used for SMTP Submission with implicit TLS. SMTP Submission is a subset of SMTP that allows the message submission agent (MSA) part of a mail client to transfer messages for delivery by a server.
- **Port 587**—Used for SMTP Submission with explicit TLS. Servers configured to support port 587 should use STARTTLS and require authentication before message submission.



Mail clients can use port 25 to submit messages to the server for delivery, but this is not best practice. Use of port 25 is typically reserved for relay between servers.

Internet Message Access Protocol

SMTP is useful only to deliver mail to hosts that are permanently available. When a message is received by an SMTP server, it delivers the message to a mailbox server. This could be a separate machine or a separate process running on the same server. A mailbox access protocol allows the user's client email software to operate the mailbox.

```

GNU nano 2.2.2      File: /etc/dovecot/dovecot.conf      Modified
protocols = imap imap
#protocols = none

# A space separated list of IP or host addresses where to listen in for
# connections. "*" listens in all IPv4 interfaces. "[::]" listens in all IPv6
# interfaces. Use "*", [::]" for listening both IPv4 and IPv6.
#
# If you want to specify ports for each service, you will need to configure
# these settings inside the protocol imap/pop3/managesieve { ... } section,
# so you can specify different ports for IMAP/POP3/MANAGESIEVE. For example:
#
# protocol imap {
#     listen = *:143
#     ssl_listen = *:943
# }
#
# protocol pop3 {
#     listen = *:10100
#     ..
# }
#
# protocol managesieve {
#     listen = *:12000
#     ..
# }
#listen = *

# Disable LOGIN command and all other plaintext authentications unless
# [ Read 1280 lines ]
^G Get Help  ^O WriteOut  ^R Read File  ^V Prev Page  ^K Cut Text   ^C Cur Pos
^X Exit      ^J Justify   ^W Where Is   ^U Next Page  ^U UnCut Text ^T To Spell

```

Configuring mailbox access protocols on a server. (Screenshot courtesy of Mozilla Foundation.)

Internet Message Access Protocol (IMAP) is the most widely used mail retrieval protocol. IMAP supports permanent connections to a server and connecting multiple clients to the same mailbox simultaneously. It also allows a client to manage the mailbox on the server (to organize messages in folders and to control when they are deleted, for instance) and to create multiple mailboxes.

A client connects to an IMAP server over TCP port 143, but this port is insecure. Connection security can be established using a TLS. The default port for IMAPs is TCP/993.



In a Windows environment, the proprietary Messaging Application Programming Interface (MAPI) protocol is typically used to access Microsoft Exchange mailboxes. MAPI uses HTTPS as a secure transport protocol.

Voice and Video Services

Voice over IP (VoIP), web conferencing, and video teleconferencing (VTC) solutions have become standard methods for the provision of business and social communications. Many networks are upgrading from legacy voice services to IP-based protocols and products.

Private Branch Exchange

Legacy voice services use the public switched telephone network (PSTN). A residential telephone installation would be serviced by a simple box providing a one- or two-line analog interface to the local exchange. This analog interface is also referred to as the plain old telephone service (POTS). Each line provides a single channel for an incoming or outgoing call. A typical business requires tens or hundreds of lines for voice communications, let alone capacity for data communications. Historically, this requirement would have been facilitated by a digital trunk line, also referred to as a time division multiplexing (TDM) circuit. A TDM can multiplex separate voice and data channels for transmission over a single cable.


A **private branch exchange (PBX)** is an automated switchboard providing a single connection point for an organization's voice lines. A TDM-based PBX connects to the telecommunications carrier over a digital trunk line, which will support multiple channels (inward and outward calls). The PBX allows for the configuration of the internal phone system to direct and route calls to local extensions and provides other telephony features such as call waiting, music on hold, and voicemail.

VoIP-Enabled PBX

TDM-based PBXes are being replaced by hybrid and fully IP/VoIP PBXes. For internal calls and conferences, a VoIP PBX establishes connections between local VoIP endpoints with data transmitted over the local Ethernet network. A VoIP PBX can also route incoming and outgoing calls from and to external networks. This might involve calls between internal and external VoIP endpoints, or with voice telephone network callers and receivers. A VoIP PBX will also support features such as music on hold and voicemail.

A TDM PBX is supplied as vendor-specific hardware. A VoIP PBX can be implemented as software running on a Windows or Linux server. Examples of software-based solutions include 3CX ([3cx.com](https://www.3cx.com)) and Asterisk ([asterisk.org](https://www.asterisk.org)). There are also hardware solutions, where the VoIP PBX runs on a router, such as Cisco Unified Communications Manager ([cisco.com/c/en/us/products/unified-communications/unified-communications-manager-callmanager/index.html](https://www.cisco.com/c/en/us/products/unified-communications/unified-communications-manager-callmanager/index.html)).

A VoIP PBX would normally be placed at the network edge and be protected by a firewall. Internal clients connect to the PBX over Ethernet data cabling and switching infrastructure, using Internet Protocol (IP) at the Network layer for addressing. The VoIP PBX uses the organization's Internet link to connect to a VoIP service provider, which facilitates inward and outward dialing to voice-based telephone networks.



A VoIP PBX facilitates internal IP calls and calls to and from external VoIP networks and the landline and cellular telephone networks. (Images © 123RF.com.)

VoIP Protocols

Voice and video services can be challenging to support because they require response times measured in milliseconds (ms). Delayed responses will result in poor call or video quality. This type of data can be one-way, as is the case with media streaming, or two-way, as is the case with VoIP and VTC.

The protocols designed to support real-time services cover one or more of the following functions:

- **Session control**—Used to establish, manage, and disestablish communications sessions. They handle tasks such as user discovery (locating a user on the network), availability advertising (whether a user is prepared to receive calls), negotiating session parameters (such as use of audio/video), and session management and termination.
- **Data transport**—Handles the delivery of the actual video or voice information.
- **Quality of service (QoS)**—Provides information about the connection to a QoS system, which in turn ensures that voice or video communications are free from problems, such as dropped packets, delay, or jitter.

Session Initiation Protocol

The **Session Initiation Protocol (SIP)** is one of the most widely used session control protocols. SIP endpoints are the end user devices (also known as user agents), such as IP-enabled handsets or client and server web conference software. Each device, conference, or telephony user is assigned a unique SIP address known as a SIP Uniform Resource Identifier (URI). Examples of SIP URIs include:

```
sip:jaime@515support.com
```

```
sip:2622136227@515support.com
```

```
sip:jaime@2622136227
```

```
meet:sip:organizer@515support.com;ms-app=conf;ms-conf-id=subg42
```



There is also a tel: URI scheme allowing SIP endpoints to dial a landline or cell phone. A tel: URI can either use the global (E.164) format (such as tel:+1-866-8358020) or a local format (for internal extensions).

SIP typically runs over UDP or TCP ports 5060 (insecured) and 5061 (SIP-TLS). SIP has its own reliability and retransmission mechanisms and can thus be seen to benefit most from the lower overhead and reduced latency and jitter of UDP. Some enterprise SIP products use TCP anyway.

Real-Time Transport Protocol and RTP Control Protocol

While SIP provides session management, the actual delivery of real-time data uses different protocols. The principal one is Real-time Transport Protocol (RTP). RTP enables the delivery of a stream of media data via UDP, while implementing some of the reliability features usually associated with TCP communications. RTP works closely with the RTP Control Protocol (RTCP). Each RTP stream uses a corresponding RTCP session to monitor the quality of the connection and to provide reports to the endpoints. These reports can then be used by the applications to modify codec parameters or by the network stacks to tune quality of service (QoS) parameters.

VoIP Phones

A VoIP/SIP endpoint can be implemented as software running on a computer or smartphone or as a dedicated hardware handset. **VoIP phones** use VLAN tagging to ensure that the SIP control and RTP media protocols can be segregated from normal data traffic. In a typical voice VLAN configuration, the LAN port on the handset is connected to the wall port, while the PC is connected to the PC port on the handset. The two devices share the same physical link, but data traffic is distinguished from voice traffic by configuring separate VLAN IDs.

Handsets can use Power over Ethernet (PoE), if available, to avoid the need for separate power cabling or batteries. There are also wireless handsets that work over 802.11 Wi-Fi networks.

Connection security for VoIP works in a similar manner to HTTPS. To initiate the call, the secure version of SIP (SIPS) uses digital certificates to authenticate the endpoints and establish a TLS tunnel. The secure connection established by SIPS can also be used to generate a master key to use with the secure versions of the transport and control protocols.

When you are installing a new handset, you should also test that the connection works and that the link provides sufficient call quality. Most service providers have test numbers to verify basic connectivity and perform an echo test call, which replays a message you record so that you can confirm voice quality.

Lesson 7.4

Disaster Recovery and High Availability



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 2.1 Explain characteristics of routing technologies.
- 3.3 Explain disaster recovery (DR) concepts.

The services offered over a network use physical media and processing assets installed at a single site. The site premises might be an office or datacenter. When deploying services, it is critical to anticipate and mitigate issues that might arise from disasters that prevent a site from working normally. The plans used to minimize the risk of site-wide problems are referred to as business continuity, while the plans used to mitigate these issues if they do occur are called disaster recovery. At this stage in your career, you must understand the concepts and technologies underpinning these plans so that you can assist with disaster recovery and high availability planning and provisioning.

As you study this lesson, answer the following questions:

- What are the relationships among business continuity, disaster recovery, high availability, fault tolerance, load balancing, and redundancy?
- What is the meaning of common metrics used in disaster recovery planning?
- What are the roles of hot, warm, and cold sites within disaster recovery?
- What is the difference between active/active and active/passive high availability?

Disaster Recovery Concepts

A **disaster recovery plan (DRP)** addresses large-scale network outage incidents. These will typically be incidents that threaten the performance or security of a whole site. A DRP should accomplish the following:

- Identify scenarios for natural and non-natural disasters and options for protecting systems.
- Identify tasks, resources, and responsibilities for responding to a disaster. Disaster recovery focuses on tasks such as switching services to failover systems or sites and restoring systems and data from backups.
- Train staff in the disaster planning procedures and how to react well to adverse events.

Testing system resilience and incident response effectiveness are crucial for organizations to recover from disruptions and maintain business continuity. By conducting various tests, organizations can identify potential vulnerabilities, evaluate the efficiency of their recovery strategies, and improve their overall preparedness for real-life incidents.

- **Tabletop exercises** involve teams discussing and working through hypothetical scenarios to assess their response plans and decision-making processes. These exercises help identify knowledge, communication, and coordination gaps, ultimately strengthening the organization's incident response capabilities. For example, a tabletop exercise might be an earthquake that destroys processing ability at a primary site, testing failover to an alternate processing location.
- **Validation tests** involve performing simulations of failovers. This tests that services can be restored using backup configurations and data. Validation tests can also test metrics for recovery time. They can also reveal any unexpected problems, such as dependencies between services not being met during the failover process.

Where disaster recovery focuses on plans for specific scenarios, a **business continuity plan (BCP)** or continuity of operations plan (COOP) is a collection of processes and resources that enable an organization to maintain normal business operations in the face of some adverse event. Continuity planning activity focuses on the functions performed by a business or other organization:

- Business impact analysis (BIA) identifies mission essential and primary business functions and the risks that would arise if the organization cannot fulfill them.
- IT contingency planning (ITCP) or IT service continuity planning (ITSCP) ensures that these functions are supported by resilient IT systems, working to identify and mitigate all single points of failure from a process or function.

Disaster Recovery Metrics

Disaster recovery planning is governed by a variety of metrics that express how reliable services are and how long it takes to recover from critical events.

One of the key properties of a resilient system is availability. Availability is the percentage of time that the system is online, measured over a certain period, typically one year. The corollary of availability is downtime; that is, the percentage or amount of time during which the system is unavailable.

High availability is a characteristic of a system that can guarantee a certain level of availability. The **Maximum Tolerable Downtime (MTD)** metric states the requirement for a business function. Downtime is calculated from the sum of scheduled service intervals (Agreed Service Time) plus unplanned outages over the period. High availability might be implemented as 24x7 (24 hours per day, seven days per week) or 24x365 (24 hours per day, 365 days per year). For a critical system, availability will be described as two-nines (99%) up to five- or six-nines (99.9999%).

Availability	Annual MTD (hh:mm:ss)
99.9999%	00:00:32
99.999%	00:05:15
99.99%	00:52:34
99.9%	08:45:36
99%	87:36:00


A system where there is almost no scheduled downtime and outages are extremely rare is also referred to as continuous availability. This sort of availability is required when there is not just a commercial imperative, but a danger of injury or loss of life associated with systems failure. Examples include networks supporting medical devices, air traffic control systems, communications satellites, networked autonomous vehicles, and smart traffic signaling systems.

The MTD metric sets the upper limit on the amount of recovery time that system and asset owners have to resume operations. Additional metrics can be used to govern recovery operations:

- **Recovery time objective (RTO).** This is the period following a disaster that an individual IT system may remain offline. This represents the maximum amount of time allowed to identify that there is a problem and then perform recovery (restore from backup or switch in an alternative system, for instance).
- **Work recovery time (WRT).** Following systems recovery, there may be additional work to reintegrate different systems, restore data from backups, test overall functionality, and brief system users on any changes or different working practices so that the business function is again fully supported.

RTO+WRT must not exceed MTD!

Recovery point objective (RPO). This is the amount of data loss that a system can sustain, measured in time units. That is, if a database is destroyed by a virus, an RPO of 24 hours means that the data can be recovered from a backup copy to a point not more than 24 hours before the database was infected.



Metrics governing mission essential functions. (Images © 123RF.com.)



Any data that has been lost between the RPO and the present needs to either be accepted as a loss or reconstructed.

Disaster Recovery Sites

Within the scope of business continuity planning, disaster recovery plans (DRPs) describe the specific procedures to follow to recover a system or site to a working state. A disaster could be anything from a loss of power or failure of a minor component to human-made or natural disasters, such as fires, earthquakes, or acts of terrorism.

Providing redundant devices and spares or network links allows the spare devices to be swapped in if existing systems fail. Enterprise-level networks often also provide for spare sites. A spare site is another location that can provide the same (or similar) level of service. A disaster or systems failure at one site will cause services to failover to the alternate processing site. Disaster recovery planning must demonstrate how this will happen, what checks need to be made to ensure that failover has occurred successfully (without loss of transactional data or service availability), and how to revert to the primary site once functionality is restored there.

Site resiliency is described as hot, warm, or cold:

- A **hot site** can failover almost immediately. It generally means that the site is already within the organization's ownership and is ready to deploy. For example, a hot site could consist of a building with operational computer equipment that is kept updated with a live dataset.
- A **warm site** could be similar but with the requirement that the latest dataset will need to be loaded.
- A **cold site** takes longer to set up. A cold site may be an empty building with a lease agreement in place to install whatever equipment is required when necessary.

Clearly, providing redundancy on this scale can be very expensive. Sites are often leased from service providers. However, in the event of a nationwide emergency, demand for the services is likely to exceed supply! Another option is for businesses to enter into reciprocal arrangements to provide mutual support. This is cost-effective but complex to plan and set up.

For many companies, the most cost-effective solution is to move processing and data storage to a cloud site. A cloud operator should be able to maintain hot site redundancy so that a disaster in one geographic area will not disrupt service, because the cloud will be supported by a datacenter in a different region.

Fault Tolerance and Redundancy

Switching services over to a disaster recovery site should be a rare occurrence. A resilient network should be able to deal with routine faults by provisioning highly available systems. A fault is usually defined as an event that causes a service to become unavailable. Each IT system will be supported by assets, such as servers, disk arrays, switches, routers, and so on. Each asset is susceptible to faults. Key performance indicators (KPIs) can be used to determine the reliability of each asset and assess whether goals for MTD, RTO, and RPO can be met. Some of the main KPIs relating to component reliability are as follows:

- **Mean time between failures (MTBF)** represents the expected lifetime of a product. The calculation for MTBF is the total operational time divided by the number of failures. For example, if you have 10 appliances that run for 50 hours and two of them fail, the MTBF is 250 hours/failure ($10 \times 50 / 2$).

- **Mean time to failure (MTTF)** expresses a similar metric for non-repairable components. For example, a hard drive may be described with an MTTF, while a server, which could be repaired by replacing the hard drive, would be described with an MTBF. The calculation for MTTF is the total operational time divided by the number of devices. For example, say two drives were installed in the server in a RAID array. One had failed after 10 years, but had never been replaced, and the second failed after 14 years, bringing down the array and the server. The MTTF of the drives is $(10+14)/2 = 12$ years.

MTTF/MTBF can be used to determine the amount of asset redundancy a system should have. A redundant system can failover to another asset if there is a fault and continue to operate normally. It can also be used to work out how likely failures are to occur.


- **Mean time to repair (MTTR)** is a measure of the time taken to correct a fault so that the system is restored to full operation. This can also be described as mean time to replace or recover. MTTR is calculated as the total number of hours of unplanned maintenance divided by the number of failure incidents. This average value can be used to estimate whether a recovery time objective (RTO) is achievable.

A system that can experience failures in individual components and subsystems and continue to provide the same (or nearly the same) level of service is said to be fault tolerant. Fault tolerance is often achieved by provisioning redundancy for critical components to eliminate single points of failure. A redundant or failover component is one that is not essential to the normal function of a system but that allows the system to recover from the failure of another component. Examples of devices and solutions that provide fault tolerance include the following:

- **Redundant spares**—Components such as power supplies, network cards, drives (RAID), and cooling fans provide protection against hardware failures. A fully redundant server configuration is configured with multiple components for each function (power, networking, and storage). A faulty component will then automatically failover to the working one.
- **Network links**—If there are multiple paths between switches and routers, these devices can automatically failover to a working path if a cable or network port is damaged.
- **Uninterruptible power supplies (UPSs) and standby power supplies**—Provide power protection in the event of complete power failure (blackout) and other types of building power issues.
- **Backup strategies**—Provide protection for data.
- **Cluster services**—A means of ensuring that the total failure of a server does not disrupt services generally.

Load Balancers

Where NIC teaming allows load balancing at the component level, a **load balancer** can be deployed as a hardware appliance or software instance to distribute client requests across server nodes in a farm or pool. You can use a load balancer in any situation where you have multiple servers providing the same function. Examples include web servers, front-end email servers, and web conferencing, video conferencing, or streaming media servers. The load balancer is placed in front of the server network and distributes requests from the client network or Internet to the application servers. The service address is advertised to clients as a virtual server. This is used to provision services that can scale from light to heavy loads, provision fault tolerant services, and to provide mitigation against distributed denial of service (DDoS) attacks.



Topology of basic load balancing architecture. (Images © 123RF.com.)

There are two main types of load balancers:

- **Layer 4 switch**—Basic load balancers make forwarding decisions on IP address and TCP/UDP header values, working at the Transport layer of the OSI model.
- **Layer 7 switch (content switch)**—As web applications have become more complex, modern load balancers need to be able to make forwarding decisions based on application-level data, such as a request for a particular URL or data types like video or audio streaming. This requires more complex logic, but the processing power of modern appliances is sufficient to deal with this.




We are used to associating switches with layer 2 (Ethernet), but appliances can perform switch-like forwarding at layer 3, layer 4, and layer 7. These are collectively referred to as multilayer switches.

High Availability Clusters

Where a load balancer distributes traffic between independent processing nodes, **clustering** allows multiple redundant processing nodes that share data with one another to accept connections. If one of the nodes in the cluster stops working, connections can failover to a working node. To clients, the cluster appears to be a single server.

Virtual IP

For example, you might want to provision two load balancer appliances so that if one fails, the other can still handle client connections. Unlike load balancing with a single appliance, the public IP used to access the service is shared between the two instances in the cluster. This is referred to as a **virtual IP** or shared or floating address. The instances are configured with a private connection, on which each is identified by its “real” IP address. This connection runs some type of redundancy protocol, such as Common Address Redundancy Protocol (CARP), that enables the active node to “own” the virtual IP and respond to connections. The redundancy protocol also implements a heartbeat mechanism to allow failover to the passive node if the active one should suffer a fault.



Topology of clustered load balancing architecture. (Images © 123RF.com.)

The same sort of topology can be used to deploy routers and firewalls for high availability and load sharing.

Active-Passive and Active-Active Clustering

In the previous example, if one node is active, the other is passive. This is referred to as **active-passive** clustering. The major advantage of active/passive configurations is that performance is not adversely affected during failover. However, the hardware and operating system costs might be higher because of the unused capacity.

An **active-active** cluster means that both nodes are processing connections concurrently. This allows the administrator to use the maximum capacity from the available hardware while all nodes are functional. In the event of a failover, the workload of the failed node is immediately and transparently shifted onto the remaining node. At this time, the workload on the remaining nodes is higher, and performance is degraded.



In a standard active-passive configuration, each active node must be matched by a passive node. There are N+1 and N+M configurations that provision fewer passive nodes than active nodes to reduce costs.


First Hop Redundancy

In a full or partial mesh network topology, alternate routes can be found to bypass failed routers or faulty connections. However, end systems are typically served by a single router configured as the default gateway. While it is possible to configure hosts with multiple default gateways for fault tolerance, this does not work well in practice, as it requires a greater degree of complexity in the hosts' routing algorithms than is typically implemented on an end system host.

To address this problem, various types of **first hop redundancy protocols (FHRP)** have been developed.

Hot Standby Router Protocol

The proprietary Hot Standby Router Protocol (HSRP) developed by Cisco allows multiple physical routers to serve as a single default gateway for a subnet. To do this, each router must have an interface connected to the subnet, with its own unique MAC address and IP address. In addition, they also need to be configured to share a common virtual IP address and a common MAC address. The group of routers configured in this way is known as a standby group. They communicate among themselves using IP multicasts and choose an active router based on priorities configured by an administrator. The active router responds to any traffic sent to the virtual IP address. Of the remaining routers in the standby group, the router with the next highest priority is chosen as the standby router. The standby router monitors the status of the active router and takes over the role if the active router becomes unavailable, also triggering the selection of a new standby router from the remaining routers in the group.



Hot Standby Router Protocol (HSRP) topology. (Image © 123RF.com.)



Cisco also has the Gateway Load Balancing Protocol (GLBP) which allows for an active/active load-balanced configuration.

Virtual Router Redundancy Protocol

The open standard protocol Virtual Router Redundancy Protocol (VRRP) is similar to HSRP, with the differences mainly being in terminology and packet formats. In VRRP, the active router is known as the master, and all other routers in the group are known as backup routers. There is no specific standby router; instead, all backup routers monitor the status of the master, and in the event of a failure, a new master router is selected from the available backup routers based on priority.

One advantage of VRRP over HSRP is that it does not require each router interface to be assigned a unique IP address. It is possible to configure VRRP routers to use only the virtual IP address. This can be useful on subnets where address space utilization is high.

Module 7

Summary

You should be able to explain the characteristics of common application ports and protocols, especially in terms of security/encryption requirements.

Guidelines for Supporting Network Application Services

Follow these guidelines to support network applications effectively:

- Deploy web servers to intranets as required. Public websites can be deployed to a perimeter network, but most organizations use some type of hosted or cloud service.
- Configure web servers with a valid certificate issued by a locally trusted or public certificate authority (CA) to enable HTTPS over TCP/443.
- Enable secure FTP on web servers, file servers, and appliances as a means of transferring files securely. FTP can be secured using SSH (SFTP) or TLS (FTPES or FTPS).
- Ensure that unencrypted local file and printer sharing services such as SMB are used only on trusted local networks. Block the SMB ports (TCP/UDP/137–139 and TCP/445) at the network perimeter. Ensure that legacy versions of the protocol are disabled.
- Deploy database services for access by application servers, rather than being directly accessible to client workstations and devices. Use access control lists to block access to RDBMS ports TCP/1521 (Oracle SQL*Net), TCP/1433 (MS SQL), TCP/3306 (MySQL/MariaDB), or TCP/5432 (PostgreSQL).
- Deploy SMTP servers to the network edge to transfer email messages to and from external recipients over TCP/25. Use TCP/587 and TLS to allow mail clients to submit messages for delivery securely. IMAP mailbox servers should be deployed as secure version (TCP/993).
- Deploy VoIP/hybrid PBX with voice gateways to local and perimeter networks to support legacy and packetized telephony devices. Configure VoIP endpoints to use secure SIP (TCP/5061) for session control and RTP/RTCP for data transfer.
- Develop disaster recovery and high availability plans and provision supporting resources:
 - Identify requirements for cold, warm, or hot alternate processing sites.
 - Test DR plans using tabletop exercises and validation tests, using the latter to develop key metrics, such as MTD, RPO, RTO, MTTR, and MTBF.

- Consider using active/passive or active/active load balancers to distribute requests and compensate for server failures.
- Use FHRP to provision multiple default gateways to compensate for router failures.

Module 8

Supporting Network Management

Module Introduction

So far in this course, you have learned about all the different network media and topologies plus the application protocols that go toward building network connectivity and services. In this module, you will demonstrate use of tools and management methods that will help you document network assets, determine baselines, and optimize your network's performance.

Module Objectives

In this module, you will do the following:

- Explain the use of configuration and change management documentation.
- Use discovery and monitoring tools to identify network assets.
- Use event management to ensure network availability.
- Use packet analysis and traffic metrics to troubleshoot performance issues.

Lesson 8.1

Organizational Policies and Documentation



Exam Objectives Covered

3.1 Explain the purpose of organizational processes and procedures.

In a well-organized network, administrators keep close control of the ways systems are configured using change management procedures and documentation. You need to understand the importance of these organizational processes and procedures, as most service requests and troubleshooting you will perform will be in the context of ticket systems.

As you study this lesson, answer the following questions:

- What is the difference between configuration and change management?
- What information should be captured in inventory documentation?
- What information is needed to create different types of network diagrams?

Configuration Management

Running an efficient network is not just about installing cabling and network devices. The administration of the network in terms of configuration documentation, change management, and monitoring is a critical task.

Configuration management means identifying and documenting all the infrastructure and devices installed at a site. It is a systematic approach to ensuring that the desired state of an IT system is maintained throughout its lifecycle. Configuration management is implemented using the following elements:

- Service assets are things, processes, or people that contribute to the delivery of an IT service. Each asset must be identified by some sort of label.
- A configuration item (CI) is an asset that requires specific management procedures for it to be used to deliver the service. CIs are defined by their attributes.
- A configuration management system (CMS) is the tools and databases that collect, store, manage, update, and present information about CIs. A small network might capture this information in spreadsheets and diagrams; there are dedicated applications for enterprise CMSs.

When discussing configuration management concepts, you need to distinguish between various configuration states:

- A baseline documents the approved or authorized state of a CI. This allows auditing processes to detect unexpected or unauthorized change. A baseline can be a **configuration baseline** (the ACL applied to a firewall, for instance) or a performance baseline (such as the throughput achieved by the firewall). A baseline configuration is sometimes referred to as a golden configuration.

- The **production configuration** is the state of a CI as used within a working network. This might deviate temporarily or persistently from the baseline. This deviation is often referred to as configuration drift.
- A **backup configuration** is a copy of the production configuration made at a particular time. As the production configuration might have drifted, a given backup might also not match the golden configuration.

Monitoring configurations requires production and backup states to be compared to the baseline. Where a configuration has drifted, it might be appropriate either to revert to the golden configuration or to update the baseline template. You will need to perform testing to determine which is the better approach for a given scenario. Preventing unexpected configuration drift requires effective change management procedures.

Network Device Backup Management

All business continuity and disaster recovery planning procedures make use of backups. The execution and frequency of backups must be carefully planned and guided by policies. In network management, backup policies are less focused on the data stored on servers and more on swiftly restoring faulty switches, routers, firewalls, and load balancers.

Each device should have a documented baseline configuration. The deployment process should be capable of applying this configuration or a backup configuration to a replacement device or when restoring a faulty device.

One complication here is that most network appliances have a startup or persistent configuration and a running configuration. In most cases, these should be the same. It is possible that a configuration oversight left a valid running configuration that was never saved as the startup configuration. Regular audits and other configuration management procedures should be used to detect and remediate running configs that differ from the saved config. It is also possible that deviations from the baseline cannot be reverted without causing disruption, and it is the baseline that must be updated.

Most devices will also support a version history of previous configurations, enabling a change to be rolled back if it causes problems.

An appliance may also support two backup modes:

- **State/bare metal**—A snapshot-type image of the whole system. This can be redeployed to any device of the same make and model as a system restore.
- **Configuration file**—A copy of the configuration data in a structured format, such as Extensible Markup Language (XML). This file can be used in a two-stage restore where the OS or firmware image is applied first (or a new appliance provisioned) and then the configuration is restored by importing the backup file.

A network appliance may also hold state information that has not been written to a log and that will not be captured by a backup of the configuration file only. State information includes data such as the MAC tables in switches or the NAT table in a firewall. Advanced firewalls may contain additional data such as malware/intrusion detection signatures. Some devices might log state data to an internal database that can be backed up periodically. In other cases, if this information needs to be preserved, the appliance should be configured to log state data to a remote server, using a protocol such as syslog.

Change Management

A documented **change management** process minimizes the risk of configuration drift and unscheduled downtime by implementing changes in a planned and controlled way. The need to change is often described either as reactive, where the change is forced on the organization, or as proactive, where the need for change is initiated internally. Changes can also be categorized according to their potential impact and level of risk (major, significant, minor, or normal, for instance).

In a formal change management process, the need for change and the procedure for implementing the change is captured in a Request for Change (RFC) document. The RFC will then be considered at the appropriate level, and affected stakeholders will be notified. Major or significant changes might be managed as a separate project and require approval through a Change Advisory Board (CAB).

Configuration changes should be made only when there is a service request ticket authorizing the change. This means that the activity of all network personnel, whether it be installing new devices or troubleshooting, is recorded in job logs. In a fully documented environment, each task will be governed by a **standard operating procedure (SOP)**. A SOP sets out the principal goals and considerations, such as budget, security, or customer contact standards, for performing a task and identifies lines of responsibility and authorization for performing it. A SOP may also contain detailed steps for completing a task in an approved way, or these steps may be presented as work instructions.


Managing changes using a ticket system facilitates request process tracking. The ticket documents request and approval, identifies stakeholders plus change and rollback plans, and monitors progress through implementation and testing of the change.

Asset Inventory Documentation

To support configuration and change management, it is crucial for an organization to have an inventory of its tangible and intangible assets. In terms of network management, hardware assets are network appliances (routers, switches, threat management devices, access points), servers, workstations, and passive network infrastructure (cabling and patch panels). Software assets are operating systems and applications.

Inventory Tools

There are many software suites and associated hardware solutions available to assist with managing inventory. An asset management database can be configured to store as much or as little information as is deemed necessary, though typical data would be type, model, serial number, asset ID, location, user(s), value, and service in



A product such as Lansweeper assists inventory management by scanning network hosts and compiling an asset information database automatically. (Screenshot used with permission from Lansweeper.)

Warranty Support and Licensing

For each asset record, there should also be a copy of or link to the appropriate vendor documentation. This includes both an invoice and warranty/support contract and support and troubleshooting guidance. The software inventory should also track license usage to help ensure that all installation complies with the vendor's licensing agreement.

Licensing for servers and network appliances can be complex, and it is easy to make configuration errors. When faced with an unexpected problem, it is often worth considering whether a licensing or feature activation issue could be the cause. On a switch or router, license failures could restrict the number of ports available, the number of routes allowed in the routing table, or the availability of routing protocols. Security and management features may have been configured under a trial or evaluation period and suddenly stop working when that grace period ends.

The starting point for troubleshooting license issues will be the log. This should show whether an evaluation/trial period has just expired or when a seat/instance count has been exceeded. Verify that the appliance has the correct licenses or activation keys installed. If relevant, ensure that the appliance can connect to its licensing or activation server.

Lifecycle Management

One of the functions supported by change management and inventory documentation is **system lifecycle management**. A system lifecycle refers to the managed acquisition, deployment, use, and decommissioning of assets. Up-to-date inventory documentation can identify assets that are no longer fully supported by the vendor or that otherwise no longer meet performance or security requirements.

When a manufacturer discontinues sales of a product, it enters an **end of life (EOL)** phase in which support and availability of spares and updates become more limited. An **end of support (EOS)** system is one that is no longer supported by its developer or vendor. EOS products no longer receive security updates and so represent a critical vulnerability if any remain in active use.



The exact terminology can vary between vendors. EOL is sometimes referred to as End of Sale (EOS). If that is the case, End of Support can be referred to as End of Service Life (EOSL).

Each type of firmware, OS, and applications software has vulnerabilities that present opportunities for would-be attackers. As soon as a vulnerability is identified in a supported product, the vendor will (or should) try to correct it. At the same time, threat actors will try to exploit it. There can never be a single comprehensive list of vulnerabilities for each bit of firmware or software, so you must stay up to date with the system security advisories posted on vendor websites and in other security reference sources. Patch management refers to the procedures put in place to manage the installation of updates for hardware (firmware) and software.



*A **patch** is a publicly released collection of updates. These can include fixes and feature changes/improvements. A **hotfix** is a code change that addresses a specific issue that can be applied without incurring downtime; conversely, a **coldfix** is one that requires the software or host to be restarted. The term **bugfix** is usually reserved for issues that are caught during product development and testing.*

The firmware on a device such as a router/firewall may be a very sophisticated piece of software. It is quite common for such software to have known vulnerabilities, so it is vital to use a secure version. Updating firmware is known as flashing the chip. This is generally done via a vendor-supplied setup program. It is important to make a backup of the system configuration (especially for a firewall) before performing a firmware update or upgrade.

A host OS, such as Windows, can apply patches individually. An appliance OS, such as Cisco IOS, must be patched to a particular version number by applying a new software image. To address a particular vulnerability, you could use a tool such as the IOS Software Checker (tools.cisco.com/security/center/softwarechecker.x) to identify the “first fix” version of IOS for that security advisory. This does mean that other changes could be introduced, so careful testing and impact assessment is required.

Once you have completed environment and compatibility checks and backed up the existing configuration, the basic upgrade process is to copy the new system image to the appliance’s flash memory. This can be done over a network using Trivial File Transfer Protocol (TFTP) or remote file copy or by using a removable flash memory card. Once the image update is in place, you run a command sequence to replace the old image and load the new one at startup.

Most software and firmware version changes and updates are upward, toward newer versions. Downgrading (or rollback) refers to reverting to a previous version of the software or firmware. This might be necessary to fix a problem caused by a recently upgraded or updated device or software. In some circumstances, downgrading might not be possible. For instance, a network appliance might not support downgrading to an earlier firmware version, and an OS might have to be reinstalled completely. When applying a patch or upgrade, it is common practice to make a configuration backup, in case settings must be reapplied after the update. When downgrading, a configuration backup might not work because it may involve settings not included in the earlier version.

Decommissioning

Each system also has a **decommissioning** phase of its lifecycle. When a server or appliance is disposed of by resale, gift, or recycling, there is a risk that software licenses could be misused or that configuration information valuable to an attacker could be leaked. These risks can be mitigated by ensuring that the built-in factory reset routine is invoked to wipe any custom configuration settings or modifications when decommissioning a server, switch, router, firewall, or printer.

A factory reset may leave **data remnants**, however. Data remnant removal is critical because an organization's confidential data or personal/sensitive data held could be compromised.

Data remnant removal refers to ensuring that no data is recoverable from hard disk drives (HDDs), flash devices or solid state drives (SSDs), tape media, and CD and DVD ROMs before they are disposed of or put to a different use. Paper documents must also be disposed of securely. Data remnants can be dealt with either by destroying the media or by sanitizing it (removing the confidential information but leaving the media intact for reuse).

Methods of destroying media include incineration, pulverization, and degaussing (for magnetic media such as hard drives).

Media **sanitization** refers to erasing data from HDD, SSD, and tape media before they are disposed of or put to a different use. The standard method of sanitizing an HDD is called overwriting. This can be performed using the drive's firmware tools or a utility program. The basic type of overwriting is called zero filling, which just sets each bit to zero. Single-pass zero filling can leave patterns that can be read with specialist tools. A more secure method is to overwrite the content with one pass of all zeros, then a pass of all ones, and then one or more additional passes in a pseudorandom pattern.

Secure Erase

Since 2001, the SATA and Serial Attached SCSI (SAS) specifications have included a **Secure Erase (SE)** command. This command can be invoked using a drive/array utility or the `hdparm` Linux utility. On HDDs, this performs a single pass of zero-filling.

For SSDs and hybrid drives and some USB thumb drives and flash memory cards, overwriting methods are not reliable, because the device uses wear-leveling routines in the drive controller to communicate which locations are available for use to any software process accessing the device. On SSDs, the SE command marks all blocks as empty. A block is the smallest unit on flash media that can be given an erase command. The drive firmware's automatic garbage collectors then perform the actual erase of each block over time. If this process is not completed (and there is no progress indicator), there is a risk of remnant recovery, though this requires removing the chips from the device to analyze them in specialist hardware.

Instant Secure Erase

HDDs and SSDs that are self-encrypting drives (SEDs) support another option, invoking a SANITIZE command set in SATA and SAS standards from 2012 to perform a crypto erase. Drive vendors implement this as **Instant Secure Erase (ISE)**. With an SED, all data on the drive is encrypted using a media encryption key. When the erase command is issued, the MEK is erased, rendering the data unrecoverable.

Physical Network Diagrams

Asset and configuration item (CI) documentation makes significant use of diagrams. A diagram is the best way to capture the complex relationships between network elements. It is important not to try to include too much information as this tends to make the diagram too complex to be useful. Consequently, a large number of diagram types are used in network management. One basic distinction is between physical and logical network diagrams. Within the class of physical network diagrams, the following types are commonly used.

Cable Maps

A **cable map** or floor plan shows how wires are routed through conduit from telecommunications closets to work areas. For example, you might use floor plans to document wall port locations and cable runs in an office. Physically accurate floor plans are hard to design and are likely to require the help of an architect or graphics professional.

A port location diagram identifies how wall ports located in work areas are connected back to ports in a distribution frame or patch panel and then from the patch panel ports to the switch ports. Rack diagrams should also show how power outlets on the uninterruptible power supply (UPS) connect to appliance power supply units (PSUs).



In order for a physical diagram of cabling and assets to make any sense, there must be a system of labeling in place for identifying these assets. A typical type of port naming convention is for alphanumeric identifiers for the campus (for multicampus networks), building (for campus networks), telecommunications space, and port. For example, CB01-01A-D01 could refer to a cable terminating at Main Campus Building (CB01), telecommunications space A on floor 1 (01A), or data port 1 (D01). Structured cable and patch cords should be labeled at both ends to fully identify the circuit.



In addition to having a diagram, it can be very useful to take a photo of the current configuration by using a digital camera or smartphone. This provides an additional visual reference for troubleshooting and identifying unauthorized changes.

Wiring Diagram


A **wiring diagram** (or pin-out) shows detailed information about the termination of twisted pairs in an RJ45 jack or Insulation Displacement Connector (IDC). You might also use a wiring diagram to document how fiber optic strands are terminated.



You should document the wiring diagrams used to terminate twisted pairs. Ethernet is wired by T568A or T568B, and the same standard should be used consistently throughout the network.

Rack Diagrams

A **rack diagram** records the position of each appliance in the rack. You can obtain stencils that represent vendor equipment from their websites or a collection such as visiocalfe.com. You can record key configuration information for each item using labels. As well as service tags and port IDs and links, you should identify which power outlets on the uninterruptible power supply (UPS) connect to which appliance power supply units (PSUs).



Designing rack layout in Microsoft Visio. (Screenshot used with permission from Microsoft.)

Logical Network Diagrams

In contrast to floor plans drawn to an accurate scale, a schematic is a simplified or abstracted representation of a system. In terms of the physical network topology, a schematic diagram can show the general placement of equipment and telecommunications rooms, plus device and port IDs, without trying to capture the exact position or relative size of any one element. Schematics can also be used to represent the logical structure of the network in terms of zones and subnets.

When you make network schematics, do not try to represent too much information in a single diagram. For example, create separate diagrams for the PHY, Data Link, and Logical (IP) layers. Some of the information appropriate to show at each layer includes the following:

- **PHY (Physical layer)**—Asset IDs, cable links, and wall/patch panel/switch port IDs. You can use color-coding or line styles to represent the cable type (make sure the diagram has an accompanying legend to explain your scheme).
- **Data Link (layer 2)**—Shows interconnections between switches and routers, with asset IDs (or the management IP of the appliance), interface IDs, and link-layer protocol and bandwidth. You could use line thickness to represent bandwidth, but for clarity it is a good idea to use labels as well.

- **Logical (IP/layer 3)**—IP addresses of router interfaces (plus any other static IP assignments) and firewalls, plus links showing the IP network ID and netmask, VLAN ID (if used), and DHCP scopes.
- **Application**—Server instances and TCP/UDP ports in use. You might also include configuration information and performance baselines (CPU, memory, storage, and network utilization) at this level.

Schematics can either be drawn manually using a tool such as Microsoft Visio or compiled automatically from network mapping software.

Schematics can use either representative icons or pictures or drawings of actual product models. As far as icons go, the ones created by Cisco are recognized as standards. These are freely available (without alteration) from Cisco's website (cisco.com/c/en/us/about/brand-center/network-topology-icons.html). Some of the more commonly used devices are shown here:

Icon	Device	Icon	Device	Icon	Device
	Hub		Content Switch		IP Phone
	Access Point		Modem		Firewall
	Bridge		CSU/DSU		Router
	Basic Switch		PBX		Wireless Router
	Layer 3 Switch		Broadband Router		Router/Firewall
	Multilayer Switch		Cable Modem		Security Appliance

Common Cisco network icons. (Images © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)

IP Address Management

An enterprise will have to manage hundreds or even thousands of IPv4 and IPv6 networks and subnets across a wide range of physical infrastructure. Maintaining visibility into IP address assignments and name resolution across physical, virtualized, and cloud infrastructure and incorporating network appliances, servers and clients, plus mobile devices and “internet of things” devices is a challenging task. Historically, IT departments might have tracked IP usage in static files such as spreadsheets. **IP address management (IPAM)** software provides better automation and oversight than these manually compiled lists.

The core function of IPAM is to scan DHCP and DNS servers and log IP address usage to a database. Most suites can scan IP address ranges to detect use of statically assigned addresses. Some IPAM software may also be able to scan the hardware associated with an IP address (device fingerprinting) and save the information to an asset inventory. IPAM software can often be used to manage and reconfigure DHCP and DNS servers remotely.

The software also provides analysis tools to allow administrators to identify overloaded DHCP scopes or to make more valuable public IP addresses available.

Common Agreements

Agreements are used between a company and its employees and between companies to enforce performance and security objectives.

Service Level Agreement Requirements

A **service level agreement (SLA)** is a contractual agreement setting out the detailed terms under which an ongoing service is provided. This can be a legally binding formal contract between supplier and customer businesses or a less formal agreement, such as an SLA agreed on between internal departments. SLA requirements define aspects of the service, such as scope, performance characteristics, and responsibilities that are agreed upon between the service provider and the customer.

Depending on the nature of your organization's business, you may be responsible for maintaining SLA requirements agreed on with your customers, using SLAs to guarantee service standards from your suppliers, or both.

Nondisclosure Agreement

A **nondisclosure agreement (NDA)** is the legal basis for protecting information assets. It defines what uses of sensitive data are permitted, what storage and distribution restrictions must be enforced, and what penalties will be incurred by breaches of the agreement. A contract of employment is highly likely to contain NDA clauses. NDAs are also used between companies and contractors and between two companies.

Memorandum of Understanding

A **memorandum of understanding (MOU)** is a preliminary or exploratory agreement to express an intent to work together. MOUs are usually intended to be relatively informal and not to act as binding contracts. MOUs almost always have clauses stating that the parties shall respect confidentiality, however.

Lesson 8.2

Host Discovery and Monitoring



Exam Objectives Covered

3.2 Given a scenario, use network monitoring technologies.

5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

There are many types of network monitoring solutions. Discovery solutions enable administrators to identify which hosts are connected. Availability and performance monitoring verifies service status. Configuration monitoring protects against unauthorized changes. While the capabilities of individual monitoring suites vary widely, you should be able to explain the basic principles underlying these solutions so that you can support their use.


As you study this lesson, answer the following questions:

- What solutions and tools are available to perform different monitoring requirements?
- What is the difference between host and service discovery?
- What outputs do performance, availability, and configuration monitors have?
- What role does baselining play in performance monitoring?

Network Discovery

One of the management tasks facing a network administrator is to verify exactly what is connected to the network and what is being communicated over it. This is usually described as **network discovery** or visibility. Visibility is necessary to confirm that servers and clients are in the correct VLANs or subnets and to try to identify rogue or unauthorized machines. An **IP scanner** is a tool that performs host discovery and can establish the overall logical topology of the network in terms of subnets and routers.

IP scanning can be performed using lightweight standalone open source or commercial tools. Examples include Nmap, Angry IP, or PRTG. Enterprise network management suites will also perform IP scanning and combine that with asset or inventory information about each host. This functionality is often referred to as IP address management (IPAM). Suites that integrate with DHCP and DNS servers can be referred to as DHCP, DNS, and IPAM (DDI).



Angry IP Scanner.


Host discovery is a basic type of IP scanning that only attempts to determine whether an IP address is “up.” There are many different host discovery techniques. Some are best at discovering large numbers of legitimate hosts quickly; others are optimized for identifying rogue hosts that are attempting to remain hidden. The most basic techniques use the ping, arp, and traceroute tools. Some suites use **Simple Network Management Protocol (SNMP)** queries, which can also report more detailed information about interface statistics, while as noted above, enterprise suites can query local DHCP and DNS servers for information. Security-oriented scanners can use specially crafted probes to locate hosts that might be configured not to respond to pings.

A scan might be performed when an administrator determines it necessary (ad hoc), or there might be a regular schedule of scans. In security terms, if an adversary knows that scanning activity is only performed using a schedule, they could try to minimize malicious activity during those windows. Ad hoc scanning might be better at detecting unauthorized activity. However, some types of scans can potentially cause disruption to network performance or host reliability.

Nmap

The **Nmap Security Scanner** (nmap.org) is widely used for IP scanning, both as an auditing and as a penetration testing tool. The tool is open-source software with packages for most versions of Windows, Linux, and macOS. It can be operated with a command line or via a GUI (Zenmap).

The basic syntax of an Nmap command is to give the IP subnet (or IP address) to scan. When used without switches like this, the default behavior of Nmap is to ping and send a TCP ACK packet to ports 80 and 443 to determine whether a host is present. On a local network segment, Nmap will also perform ARP and Neighbor Discovery (ND) sweeps. If a host is detected, Nmap performs a port scan against that host to determine which services it is running. This OS fingerprinting can be time consuming on a large IP scope. If you want to perform only host discovery, you can use Nmap with the `-sn` switch to suppress the port scan. The tool can also work out hop counts by specifying the `--traceroute` switch.



Nmap discovery scan. (Screenshot used with permission from Nmap.)


A variety of options are available for custom scans to try to detect stealthy hosts (nmap.org/book/host-discovery-techniques.html).

Nmap Port Scanning

Many of the tools used for host discovery can also perform remote port scanning. A **port scanner** tries to identify which TCP and UDP ports are listening. By analyzing responses, a scanner can also try to perform service and host detection to identify software and OS versions. As with host discovery, there are many different techniques for performing port scans. Some techniques are designed for covert use (to try to avoid detection of the scanning activity by the target), and some are designed to probe beyond security barriers, such as firewalls.


As examples, the following represent some of the main types of scanning that Nmap can perform:

- **TCP SYN (-sS)**—This is a fast technique (also referred to as half-open scanning) as the scanning host requests a connection without acknowledging it. The target's response to the scan's SYN packet identifies the port state.
- **TCP connect (-sT)**—A half-open scan requires Nmap to have privileged access to the network driver so that it can craft packets. If privileged access is not available, Nmap must use the OS to attempt a full TCP connection. This type of scan is less stealthy.
- **UDP scans (-sU)**—Scan UDP ports. As these do not use ACKs, Nmap needs to wait for a response or timeout to determine the port state, so UDP scanning can take a long time. A UDP scan can be combined with a TCP scan.
- **Port range (-p)**—By default, Nmap scans 1,000 commonly used ports. Use the `-p` argument to specify a port range. You can also use `--top-ports n`, where *n* is the number of commonly used ports to scan. The frequency statistics for determining how commonly a port is used are stored in the `nmap-services` configuration file.



Half-open scanning with Nmap. (Screenshot used with permission from Nmap.)

When services are discovered, you can use Nmap with the `-sV` or `-A` switch to probe a host more intensively to discover the software or software version operating each port. The process of identifying an OS or software application from its responses to probes is called fingerprinting.



The responses to network probes can be used to identify the type and version of the host operating system. (Screenshot used with permission from Nmap.)

Discovery Protocols

Nmap is a relatively complex app that needs a Windows or Linux computer to run it. Switch, router, and access point appliances can run simpler discovery protocols to identify other devices on the same local (Data Link) network.

Cisco Discovery Protocol

Cisco Discovery Protocol (CDP) runs by default on all Cisco switch, router, and access point hardware. It uses Data Link layer multicast messaging to send status announcements over local interfaces every 60 seconds. CDP uses the multicast address 01:00:0c:cc:cc:cc. Each device keeps a cache table of the data compiled from announcements it has received.

The command `show cdp neighbors` reports information from the CDP cache. CDP can report device ID/host name, IOS version, interface addresses and statistics, VLAN information, and Power over Ethernet usage.

```
Switch#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
Device ID        Local Intrfce   Holdtme    Capability   Platform     Port ID
Switch           Fas 0/2         166        H P         7960
Router           Fas 0/24        166        R           C2800        Fas 0/0
Router           Fas 0/24        166        R           C2800        Fas 0/0.10
Router           Fas 0/24        166        R           C2800        Fas 0/0.20
Switch           Fas 0/1         166        H P         7960
Switch           Fas 0/3         166        H P         7960
```

Output from show cdp neighbors command listing router physical interface plus subinterfaces and a number of switches (these are actually VoIP handsets with embedded voice VLAN switches). (Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)

Link Layer Discovery Protocol

CDP performs a very useful function, but as a proprietary protocol is restricted to use on Cisco hardware or hardware from licensed vendors. The IEEE **Link Layer Discovery Protocol (LLDP)** is a standards-based approach that offers similar functionality. LLDP uses the multicast address 01:80:c2:00:00:0e. It sends announcements every 30 seconds, by default.

Performance Monitoring

Network monitoring tools fulfill a wide range of functions beyond host and service discovery. As input, they can capture and analyze traffic, monitor interface and device metrics, and consolidate log data. As output, they can alert you to events, help you define baselines, analyze traffic patterns and congestion, determine upgrade and forecast needs, and generate reports for management.

Performance Metrics

When you are monitoring a network host or appliance, several **performance metrics** can tell you whether the host is operating normally:

- **Bandwidth**—This is the rated speed of all the interfaces available to the device, measured in Mbps or Gbps. For wired Ethernet links, this will not usually vary, but the bandwidth of WAN and wireless links can change over time.

- **Utilization/throughput**—This is the actual amount of data transferred. Utilization expresses this as a percentage of the bandwidth, while throughput is the amount of data transferred per unit of time.
- **CPU and memory**—Devices such as switches and routers perform a lot of processing. If CPU and/or system memory utilization (measured as a percentage) is consistently very high, an upgrade might be required. High CPU utilization can also indicate a problem with network traffic.
- **Storage**—Some network devices require persistent storage (typically, one or more flash drives) to keep configuration information and logs. Storage is measured in MB or GB. If the device runs out of storage space, it could cause serious errors. Servers also depend on fast input/output (I/O) to run applications efficiently.

Baseline Metrics

Baseline metrics establish the level of resource utilization at a point in time, such as when the system was first installed. This provides a comparison to measure system responsiveness later. For example, if a company is expanding a remote office that is connected to the corporate office with an ISP's basic tier package, the baseline can help determine if there is enough reserve bandwidth to handle the extra user load, or if the basic package needs to be upgraded to support higher bandwidths.

Reviewing baselines is the process of evaluating whether a baseline is still fit for purpose or whether a new baseline should be established. Changes to the system usually require a new baseline to be taken.

Availability Monitoring

An **availability monitor** triggers an alert or alarm if a host or service experiences an outage or other unscheduled downtime. These tools are also referred to as heartbeat monitors or uptime monitors. Most work by sending a probe to the target service and checking for a non-error response. For example, an HTTP service should return a 200 status code when a resource is available. Some monitors may also be configured to check the expiry date of digital certificates.

When you are troubleshooting unresponsive service issues, they will usually manifest with multiple clients being unable to connect. There can be any number of underlying causes, but consider some of the following:

- The application or OS hosting the service has crashed (or there is a hardware or power problem).
- The server hosting the service is overloaded (high CPU/memory/disk I/O utilization/disk space utilization). Try throttling client connections until the server resources can be upgraded.
- There is congestion in the network, either at the client or server end (or both). Use `ping` or `tracert` to check the latency experienced over the link and compare to a network performance baseline. Again, throttling connections or bandwidth may help to ease the congestion until higher bandwidth links can be provisioned.

- A broadcast storm is causing loss of network bandwidth. Switching loops causes broadcast and unknown unicast frames to circulate the network perpetually, as each switch repeatedly floods each frame. A broadcast storm may quickly consume all link bandwidth and crash network appliances (check for excessive CPU utilization on switches and hosts). The Spanning Tree Protocol (STP) is supposed to prevent such loops, but this can fail if STP communications between switches do not work correctly, either because of a fault in cabling or a port/transceiver or because of a misconfiguration. Ports can also be configured with storm control. This will start to drop broadcasts and unknown unicasts if they reach a certain level.
- Network congestion or high host CPU/memory utilization may also be a sign that the service is being subject to a denial of service (DoS) attack. Look for unusual access patterns (for example, use GeoIP to graph source IP addresses by country and compare to baseline access patterns).



If users on a LAN cannot connect to an external service, such as a cloud application, use a site such as [isitdownrightnow.com](https://www.isitdownrightnow.com) to test whether the issue is local to your network or a problem with the service provider site.



Be proactive in monitoring service availability so that you can resolve problems before they affect large numbers of clients.

Configuration Monitoring

Configuration management processes ensure that all network appliances are in a known state. Recall that there are various configuration states:

- The baseline or golden configuration is a template for the state that a given device should be in.
- The production configuration is the state that the device is actually in. Also, a device could have a running configuration that is different from its startup configuration.
- A backup configuration is a point-in-time copy of a running or startup configuration.

A **configuration monitor** generates logs, alerts, or alarms when there is a change to a device's production configuration. Some tools may be capable of identifying line-by-line differences between production and baseline configurations.

Lesson 8.3

Simple Network Management Protocol



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.2 Given a scenario, use network monitoring technologies.

Performance and configuration monitoring depends on querying information and statistics from an appliance's local system monitor. While there are plenty of proprietary tools and agents available, the Simple Network Management Protocol (SNMP) provides a standardized way of retrieving this data.

As you study this lesson, answer the following questions:

- What is the role of the MIB when using SNMP?
- How is a trap used in network administration?
- Why doesn't the community string provide security for SNMP devices?

SNMP Agents and Monitors

Stand-alone devices may have a web console to use for performance and environmental monitoring. Local access is not scalable to managing tens or hundreds of devices, however.

The Simple Network Management Protocol (SNMP) is a widely used framework for remote management and monitoring of servers and network appliances. SNMP consists of agents and a monitoring system.

SNMP Agents

The agent is a process (software or firmware) running on a switch, router, server, or other SNMP-compatible network device. A device running an SNMP agent is referred to as a managed device. The agent maintains a data store called a **management information base (MIB)** that holds variables relating to the activity of the device, such as the number of frames per second handled by a switch. Each parameter stored in a MIB is referred to by a numeric Object Identifier (OID). OIDs are stored within a tree structure. Part of the tree is generic to SNMP, while part can be defined by the device vendor.

An agent is configured with the community string or community name of the computers allowed to manage the agent and the IP address or host name of the server running the management system. The community string acts as a rudimentary type of password. An agent can pass information only to management systems configured with the same community string. There are usually two community strings; one for read-only access and one for read-write access (or privileged mode).

Services: Net-SNMP

General **SNMPv3 Users** full help

Enable SNMP Service ☒

SNMP Community

SNMP Location

SNMP Contact

Add AgentX Support ☐

Layer 3 Visibility ☐

Display Version in OID ☐

Listen IPs Clear All Copy

Save


*Configuring an SNMP agent on an OPNsense security appliance.
(Screenshot courtesy of OPNsense.)*

SNMP Monitor

An SNMP monitor is management software that provides a location from which you can oversee network activity. The monitor polls agents at regular intervals for information from their MIBs and displays the information for review. It also displays any trap operations as alerts for the network administrator to assess and act upon as necessary. The monitor can retrieve information from a device in two main ways:

- **Get**—The software queries the agent for a single OID. This command is used by the monitor to perform regular polling (obtaining information from devices at defined intervals).
- **Trap**—The agent informs the monitor of a notable event, such as port failure. The threshold for triggering traps can be set for each value.

The monitor can be used to change certain variables using the Set command. It can also walk an MIB subtree by using multiple Get and Get Next commands. This is used to discover the complete layout of an MIB. Device queries take place over UDP port 161; traps are communicated over UDP port 162.



SNMP collects information from network devices for diagnostic purposes. (Images © 123RF.com.)

SNMP Security

Many networks run SNMP v2c. This protocol version has no support for robust authentication or encryption. When using SNMP v2c, apply the following guidelines:

- SNMP v2c **community strings** are sent in plaintext and should not be transmitted over the network if there is any risk of interception.
- Use difficult-to-guess community strings; never leave the community string blank or set it to the default.
- Use access control lists to restrict management operations to known hosts (that is, restrict to one or two host IP addresses).

SNMP v3 supports encryption and strong user-based authentication. Instead of community strings, the agent is configured with a list of usernames and access permissions. When authentication is required, the SNMP message is signed with a hash of the user's passphrase. The agent can verify the signature and authenticate the user using its own record of the passphrase.

If authNoPriv mode is used, packets are not encrypted. authPriv enables encryption using the credential as a key.

Lesson 8.4

Event Management



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.2 Given a scenario, use network monitoring technologies.

Gathering information from logs is essential for troubleshooting a wide range of issues. On a large network, administrators typically deploy “single pane of glass” solutions that collect and aggregate logs and display the results as dashboards. Understanding the processes that these products use to forward and manage event data will be crucial to your role in network support.

As you study this lesson, answer the following questions:

- What is log forwarding?
- What is the difference between log collection and log aggregation?
- How are events categorized by severity or priority?
- What is the role of security information and event management (SIEM) products?

Network Device Logs

Network device logs are one of the most valuable sources of performance, troubleshooting, and security auditing information. A single logged event consists of metadata, such as the date and time, category, and event ID, plus a description and contents of error or informational output. For example, you can use a system log to troubleshoot an IP conflict by looking for TCP/IP events or to determine when and why a system was shut down.

While the specifics of what and where events are logged can vary widely from platform to platform, it is possible to discern some general log types, including system, security, application, and performance or traffic.

System and Application Logs

A system log records startup events plus subsequent changes to the configuration at an OS level. This will certainly include kernel processes and drivers but could also include core services.

By contrast, an application log records data for a single specific service, such as DNS, HTTP, or a database. Note that a complex application could write to multiple log files. For example, the Apache web server logs errors to one file and access attempts to a

Audit Logs

An audit log records use of authentication and authorization privileges. It will generally record success/fail type events. An audit log might also be described as an access log or security log. Audit logging might be performed at an OS level and at a per-application level.

Firewall: Log Files: Live View

Interface	Time	Source	Destination	Proto	Label
wan	Aug 30 08:59:42	203.0.113.44:51964	198.51.100.29:25	tcp	Default deny rule
wan	Aug 30 08:59:09	203.0.113.44:51964	198.51.100.29:25	tcp	Default deny rule
lan	Aug 30 08:59:01	10.1.24.101:49884	172.16.0.201:80	tcp	Default allow LAN to any rule
lan	Aug 30 08:58:57	10.1.24.101:49881	172.16.0.201:80	tcp	Default allow LAN to any rule
wan	Aug 30 08:58:53	203.0.113.44:51964	198.51.100.29:25	tcp	Default deny rule
wan	Aug 30 08:58:46	203.0.113.44:49690	172.16.0.201:80	tcp	Allow web access (unencrypted)

Viewing audit logs on an OPNsense security appliance. (Screenshot courtesy of OPNsense.)



Audit logs typically associate an action with a particular user. This is one of the reasons that it is critical that users not share logon details. If a user account is compromised, there is no means of tying events in the log to the actual attacker.


Performance/Traffic Logs

Performance and traffic logs record metrics for compute, storage, and network resources over a defined period.

Log Collectors and Syslog

When an event is generated by a device, it might be logged to the console or to a local file. Accessing the console or log files on individual devices is impractical on a large network. A log collector receives event messages forwarded from numerous devices to a single storage location. As well as collecting logs, the system can be configured to run one or more status and alerting dashboards.

Syslog is an example of a protocol and supporting software that facilitates log collection. It has become a de facto standard for logging events from distributed systems. For example, syslog messages can be generated by Cisco routers and switches, as well as UNIX or Linux servers and workstations. A syslog collector usually listens on UDP port 514.



*Configuring an OPNsense security appliance to transmit logs to a remote syslog server.
The syslog server has the IP address 10.1.16.242 and listens on UDP port 514.
(Screenshot courtesy of OPNsense.)*

As well as a protocol for forwarding messages to a remote log collector, Syslog provides an open format for event data. A syslog message comprises a PRI code, a header containing a timestamp and host name, and a message part. The PRI code is calculated from the facility and a severity level. The message part contains a tag showing the source process plus content. The format of the content is application dependent. It might use space- or comma-delimited fields or name/value pairs, such as JavaScript Object Notation (JSON) data.

Event Prioritization and Alerting

Devices can generate thousands of events per hour. There must be a system for prioritizing those that require an immediate response over those that are nonurgent or informational only. Most logging systems categorize each event. For example, in Windows, system and application events are defined as Informational, Warning, or Critical, while audit events are categorized as Success or Fail.

Syslog severity levels are as follows:


Code	Level	Interpretation
0	Emergency	The system is unusable (kernel panic).
1	Alert	A fault requiring immediate remediation has occurred.
2	Critical	A fault that will require immediate remediation is likely to develop.
3	Error	A nonurgent fault has developed.
4	Warning	A nonurgent fault is likely to develop.
5	Notice	A state that could potentially lead to an error condition has developed.
6	Informational	A normal but reportable event has occurred.
7	Debug	Verbose status conditions used during development and testing.

The **logging level** configured on each host determines the maximum level at which events are recorded or forwarded. For example, if the logging level for remote forwarding is set to 4, events that are level 5, 6, or 7 are *not* forwarded.

An automated event management system can be configured to generate some sort of alert. An alert can indicate when certain event types of a given severity are encountered. Alerts can also be generated by setting thresholds for performance counters. Examples include packet loss, link bandwidth drops, number of sessions established, delay/jitter in real-time applications, and so on. Finally, an alert can reveal an anomaly, or patterns of behavior or usage that are not consistent with normal activity. Most network monitors also support heartbeat tests so that you can receive an alert if a device or server stops responding to probes.

Setting alerts is a matter of balance. On the one hand, you do not want performance to deteriorate to the point that it affects user activity; on the other hand, you do not want to be overwhelmed by alerts.

You can also make a distinction between alerts and notifications. An alert means that the system has matched some sort of pattern or filter that should be recorded and highlighted. A notification means that the system sends a message to advertise the occurrence of the alert. A low priority alert may simply be displayed in the system dashboard. A high priority alert might use some sort of active notification messaging, such as emailing a system administrator, sending a text message (SMS) to a phone, or triggering a physical alarm signal.




*Configuring alert and notification settings in EventSentry SIEM.
(Screenshot courtesy of NETIKUS.NET Ltd.)*

There should be some process for acknowledging and dismissing alerts as they are generated. A serious alert may need to be processed as an incident and assigned a job ticket for formal investigation. If an alert is a false positive, it can be dismissed. If the management system or dashboard is allowed to become cluttered with old alerts, it is much more difficult to identify new alerts and gauge the overall status of the network.

Security Information and Event Management

As distinct from collection, log aggregation refers to normalizing data from different sources so that it is consistent and searchable. Aggregation can also condense repetitive, individual events to a summary event that counts the number of instances. **Security Information and Event Management (SIEM)** is designed to integrate network and security monitoring through automated collection, aggregation, and analysis of log data. The core function of a SIEM tool is to aggregate logs from multiple sources. In addition to logs from Windows and Linux-based hosts, this could include devices such as switches, routers, firewalls, intrusion detection sensors, vulnerability scanners, malware scanners, and databases.

To perform log aggregation, SIEM software features connectors or plug-ins to interpret (or parse) data from distinct types of systems and to account for differences between vendor implementations. Usually parsing will be carried out using regular expressions tailored to each log file format to identify attributes and content that can be mapped to standard fields in the SIEM's reporting and analysis tools. Another important function is to normalize date/time zone differences to a single timeline.




Wazuh SIEM dashboard—Configurable dashboards provide a high-level status view of network security metrics. (Screenshot used with permission from Wazuh Inc.)

Different monitoring and logging products can also be integrated using an application programming interface (API). An API makes the functions of a particular product available to scripts. For example, a configuration monitor could have an API that allows a SIEM to initiate a scan and return the results directly to the SIEM.

Log Reviews

Monitoring involves viewing traffic, protocols, and events in real time. Network and log reviewing, or analysis, involves later inspection and interpretation of the captured data to determine what was happening on the network during the capture. Only referring to the logs after a major incident misses the opportunity to identify performance problems or security issues early and to respond proactively.

Not all performance incidents will be revealed by a single event. One of the features of log analysis and reporting software should be to identify trends. Examining each event in a log file makes it difficult to spot a trend. Plotting data as a graph is particularly helpful as it makes it easier to spot trends, spikes, or troughs in a visualization of events rather than the raw data. Most performance monitors, log collectors, and SIEMs can plot metrics in a graph.



Data retrieved from a performance log file. (Screenshot courtesy of Microsoft.)

Lesson 8.5

Packet Capture and Analysis



Exam Objectives Covered

3.2 Given a scenario, use network monitoring technologies.

5.5 Give a scenario, use the appropriate tool or protocol to solve networking issues.

One of the critical tasks for network administrators is to identify and analyze the traffic passing over network links. This information is used to troubleshoot network services and to verify the security of the network.

As you study this lesson, answer the following questions:

- What is the difference between a packet capture tool and a packet analyzer?
- What are the different methods of packet capture?
- What is the difference between packet and traffic analysis?

Packet Capture

A protocol analyzer is one of the most important tools used for network support. A protocol analyzer allows inspection of traffic received by a host or passing over a network link. A protocol analyzer depends on a **packet sniffer**. A sniffer captures frames moving over the network medium.



Often the terms “sniffer” and “protocol analyzer” are used interchangeably, but be aware that they might be implemented separately.

A basic software-based sniffer installed to a host will simply interrogate the frames received by the network adapter by installing a special driver. This allows the frames to be read from the network stack and saved to a file on disk. They also support filters to reduce the amount of data captured.

There are three main options for connecting a sniffer to the appropriate point in the network:

- **SPAN (switched port analyzer)/port mirroring**—This means that the sensor is attached to a specially configured port on the switch that receives copies of frames addressed to nominated access ports (or all the other ports). This method is not completely reliable. Frames with errors will not be mirrored, and frames may be dropped under heavy load.
- **Passive test access point (TAP)**—This is a box with ports for incoming and outgoing network cabling and an inductor or optical splitter that physically copies the signal from the cabling to a monitor port. There are types for copper and fiber optic cabling. Unlike a SPAN, no logic decisions are made so the monitor port receives every frame—corrupt or malformed or not—and the copying is unaffected by load.

- **Active TAP**—This is a powered device that performs signal regeneration (again, there are copper and fiber variants), which may be necessary in some circumstances. Gigabit signaling over copper wire is too complex for a passive tap to monitor, and some types of fiber links may be adversely affected by optical splitting. Because it performs an active function, the TAP becomes a point of failure for the links in the event of power loss.

A TAP will usually output two streams to monitor a full-duplex link (one channel for upstream and one for downstream). Alternatively, there are aggregation TAPs, which rebuild the streams into a single channel, but these can drop frames under very heavy load.

tcpdump

tcpdump is a command line packet capture utility for Linux, providing a user interface to the libpcap library. The basic syntax of the command is:

```
tcpdump -i eth0
```

...where `eth0` is the interface to listen on (you can substitute with the keyword `any` to listen on all interfaces of a multi-homed host). The utility will then display captured packets until halted manually (by pressing `Ctrl+C`). The operation of the basic command can be modified by switches. For example, the `-w` and `-r` switches write output to a file and read the contents of a capture file respectively. The `-v`, `-vv`, and `-vvv` can be used to increase the amount of detail shown about each frame while the `-e` switch shows the Ethernet header.

tcpdump is often used with some sort of filter expression:

- **Type**—Filter by `host`, `net`, `port`, or `portrange`.
- **Direction**—Filter by source (`src`) or destination (`dst`) parameters (`host`, `network`, or `port`).
- **Protocol**—Filter by a named protocol rather than port number (for example, `arp`, `icmp`, `ip`, `ip6`, `tcp`, `udp`, and so on).

Filter expressions can be combined by using Boolean operators:

- `and (&&)`
- `or (||)`
- `not (!)`

Filter syntax can be made even more detailed by using parentheses to group expressions. A complex filter expression should be enclosed by quotes. For example, the following command filters frames to those with the source IP 10.1.0.100 and destination port 53 or 80:

```
tcpdump -i eth0 "src host 10.1.0.100 and (dst port 53 or dst port 80)"
```




Refer to tcpdump.org for the full help and usage examples. [ngrep \(github.com/jpr5/ngrep\)](https://github.com/jpr5/ngrep) is another useful packet capture and analysis tool. As well as the standard filter syntax, it supports use of regular expressions (regexr.com) to search and filter capture output. You can also use the netcat tool (nmap.org/ncat) to copy network traffic from one host to another for analysis.

Protocol Analyzers

A **protocol analyzer** works in conjunction with a packet capture or sniffer tool. You can either analyze a live capture to analyze frames as they are read by a sniffer or open a saved capture (.pcap) file. Most protocol analyzer tools bundle a sniffer component with the analyzer in the same software package.

One function of a protocol analyzer is to parse each frame in a stream of traffic to reveal its header fields and payload contents in a readable format. Analyzing protocol data at the frame or packet level will help to identify protocol or service misconfigurations. **Wireshark** ([wireshark.org](https://www.wireshark.org)) is an open source graphical packet capture and analysis utility, with installer packages for most operating systems. Having chosen the interfaces to listen on, the output is displayed in a three-pane view, with the top pane showing each frame, the middle pane showing the fields from the currently selected frame, and the bottom pane showing the raw data from the frame in hex and ASCII.



Wireshark protocol analyzer. (Screenshot courtesy of Wireshark.)

As a live stream or capture file can contain hundreds or thousands of frames, you can use display filters to show only particular frames or sequences of frames. Another useful option is to use the **Follow TCP Stream** context command to reconstruct the packet contents for a TCP session.

Another function of a protocol analyzer is to perform traffic analysis. Rather than reading each frame individually, you use the tool to monitor statistics related to communications flows, such as bandwidth consumed by each protocol or each host, identifying the most active network hosts, monitoring link utilization and reliability, and so on. In Wireshark, you can use the **Statistics** menu to access traffic analysis tools.

Wireshark · Conversations · Ethernet

Ethernet · 19		IPv4 · 12		IPv6 · 9		TCP · 92		UDP · 163	
Address A	Address B	Packets	Bytes	Packets A → B	Bytes A → B	Packets B → A	Bytes B → A	Rel Start	Duration
10.1.0.2	10.1.0.101	52,173	505 M	16,583	494 M	35,590	10 M	112.913034	861.6283
10.1.0.2	10.1.0.102	4,396	11 M	1,980	11 M	2,416	475 k	323.374607	638.5910
10.1.0.1	10.1.0.2	809	157 k	401	78 k	408	78 k	0.000021	973.6849
10.1.0.102	10.1.0.255	84	8662	84	8662	0	0	80.060952	894.4482
10.1.0.101	239.255.255.250	13	5616	13	5616	0	0	258.131340	6.8096
10.1.0.102	239.255.255.250	8	2262	8	2262	0	0	321.757773	6.0453
10.1.0.101	10.1.0.255	7	1318	7	1318	0	0	80.758797	717.1660
10.1.0.102	255.255.255.255	2	684	2	684	0	0	578.264063	335.0809
10.1.0.1	10.1.0.255	3	578	3	578	0	0	88.573811	722.9047
10.1.0.2	10.1.0.255	2	486	2	486	0	0	92.770001	717.9873
10.1.0.101	224.0.0.251	2	138	2	138	0	0	258.358084	2.1196
10.1.0.102	224.0.0.252	2	128	2	128	0	0	578.258570	0.1087

☐ Name resolution
 ☐ Limit to display filter
 ☐ Absolute start time
 Conversation Types ▼
 Copy Follow Stream... Graph... Close Help

Using the Conversations option from Wireshark's Statistics tools.
(Screenshot courtesy of Wireshark.)

Wireshark · Protocol Hierarchy Statistics · Ethernet

Protocol	Percent Packets	Packets	Percent Bytes	Bytes	Bits/s
Frame	100.0	57626	100.0	517234742	5517 k
Ethernet	100.0	57626	0.2	806764	8606
Internet Protocol Version 6	0.7	390	0.0	15600	166
Internet Protocol Version 4	99.3	57204	0.2	1144080	12 k
User Datagram Protocol	0.5	283	0.0	2264	24
Simple Service Discovery Protocol	0.0	12	0.0	1380	14
NetBIOS Name Service	0.0	23	0.0	1150	12
NetBIOS Datagram Service	0.0	11	0.0	2158	23
Multicast Domain Name System	0.0	2	0.0	54	0
Link-local Multicast Name Resolution	0.0	2	0.0	44	0
Domain Name System	0.4	218	0.0	9278	98
Data	0.0	9	0.0	5616	59
Connectionless Lightweight Director...	0.0	4	0.0	811	8
Bootstrap Protocol	0.0	2	0.0	600	6
Transmission Control Protocol	98.8	56921	99.6	515215393	5496 k
Simple Mail Transfer Protocol	4.4	2515	0.7	3421454	36 k
NetBIOS Session Service	37.2	21423	77.9	402860633	4297 k
Kerberos	0.1	74	0.0	70741	754
Internet Message Access Protocol	5.0	2894	2.0	10273205	109 k
Hypertext Transfer Protocol	0.1	29	0.1	491155	5239
Data	0.0	18	0.0	18	0
Address Resolution Protocol	0.1	32	0.0	896	9

No display filter. Close Copy Help

Using the Protocol Hierarchy tool in Wireshark to view the most active protocols on a network link.
This sort of report can be used to baseline network activity. (Screenshot courtesy of Wireshark.)

Lesson 8.6

Traffic Monitoring



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 3.2 Given a scenario, use network monitoring technologies.
- 5.4 Given a scenario, troubleshoot common performance issues.
- 5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

Modern networks have demanding performance requirements that can only be evaluated and met by assessing a range of performance metrics. Understanding these metrics and how to record and measure them will help you to optimize a network to perform at its peak level.

As you study this lesson, answer the following questions:

- What is the difference in bandwidth, throughput, latency, and jitter?
- What tools can be used to report and test traffic metrics?
- How does traffic shaping ensure quality of service for real-time applications?

Common Performance Issues

Quality of service (QoS) protocols and appliances are designed to support real-time services. Applications such as voice and video that carry real-time data have different network requirements than the sort of data represented by file transfer. With “ordinary” data, it might be beneficial to transfer a file as quickly as possible, but the sequence in which the packets are delivered and the variable intervals between packets arriving do not materially affect the application. This type of data transfer is described as bursty.

While streaming video applications can have a high-bandwidth requirement in terms of the sheer amount of data to be transferred, bandwidth on modern networks is typically less of a problem than packet loss, latency, and jitter.

Bandwidth

Bandwidth is the amount of information that can be transmitted, measured in bits per second (bps), or some multiple thereof. Bandwidth expresses the available capacity of the link. When monitoring, you need to distinguish between the nominal data link/Ethernet bit rate, the throughput of a link at layer 3, and the goodput available to an application.

Bandwidth for audio depends on the sampling frequency (Hertz) and bit depth of each sample. For example, early digital telecommunications links were based on 64 Kbps channels. This was derived through the following calculation:

- The voice frequency range is 4,000 Hz. This must be sampled at twice the rate to ensure an accurate representation of the original analog waveform.
- The sample size is 1 byte (or 8 bits). Therefore, 8 KHz x 8 bits = 64 Kbps.

For VoIP, bandwidth requirements for voice calling can vary, but allowing 100 Kbps per call upstream and downstream should be sufficient in most cases.

Bandwidth required for video is determined by image resolution (number of pixels), color depth, and the frame rate, measured in frames per second (fps).

Bottlenecks

A **bottleneck** is a point of poor performance that reduces the productivity of the whole network. A bottleneck may occur because a device is underpowered or faulty. It may also occur because of user or application behavior. To identify the cause of a bottleneck, you need to identify where and when on the network overutilization or excessive errors occur. If the problem is continual, it is likely to be device related; if the problem only occurs at certain times, it is more likely to be user or application related.

Packet Loss

Packet loss is expected but only to a degree. The larger the network, the more likely you are to lose packets during heavy traffic periods. If you run a packet sniffer on the affected segment, high numbers of TCP retransmission and duplicate acknowledgments are strong indicators of packet loss. Knowing where and when the packet loss occurs can direct you to the device that is dropping the frames. Reasons packets are dropped can include the following:

- A server, router, or switch is overloaded.
- A power outage occurred.
- A firewall is blocking packets from a known destination.
- A malicious actor is interfering with network transmissions.
- Faulty firmware is causing packet processing errors.

Latency and Jitter

Problems with the timing and sequence of packet delivery are defined as latency and jitter. Latency is the time it takes for a transmission to reach the recipient, measured in milliseconds (ms). **Jitter** is defined as being a variation in the delay. Jitter manifests itself as an inconsistent rate of packet delivery. Jitter is also measured in milliseconds, using an algorithm to calculate the value from a sample of transit times.

Latency and jitter are not significant problems when data transfer is bursty, but real-time applications are much more sensitive to their effects because they manifest as echo, delay, and video slow down. If packets are delayed, arrive out of sequence, or are lost, then the receiving host must buffer received packets until the delayed packets are received. If packet loss or delay is so excessive that the buffer is exhausted, then noticeable audio or video problems (artifacts) are experienced by users.

You can test the latency of a link using tools such as `ping`, `pathping`, and `mtr`. You can also use `mtr` to calculate jitter. When assessing latency, you need to consider the Round Trip Time (RTT). VoIP is generally expected to require an RTT of less than 300 ms. Jitter should be 30 ms or less. The link should also not exhibit more than 1% packet loss.

Also, if you run a packet capture, Wireshark gives each packet a timestamp relevant to when the last frame was sent from the very first transmitted frame. This can find the delays in a TCP conversation during a session between two devices. You can plot a sequence graph (using the Statistics tab on Wireshark's menu) to visually represent how this delay behaves. The line should have a gradual, steady increase upward to the right. An optimal network should show small gaps between each transmission. The longer the gap and more jagged the graph, the more latency is being introduced.

Interface Statistics

To diagnose a performance issue due to congestion, bottlenecking, bandwidth, or packet loss, you must collect data and configure alerts for **interface statistics**, whether on a network adapter or switch or router port.

- **Utilization**—The data transferred over a period. This can either be measured as the amount of data traffic both sent and received (measured in bits or bytes per second or a multiple thereof) or calculated as a percentage of the available bandwidth.



You also need to differentiate between average utilization and peak utilization. If average utilization is around 80%, it may appear that there is sufficient bandwidth. However, if peak utilization often spikes to 100%, then that will manifest as delay and packet loss and may require that you upgrade the link. Monitoring the queue length can help to determine whether the link is a bottleneck.

- **Per-protocol utilization**—Packet or byte counts for a specific protocol. It is often useful to monitor both packet counts and bandwidth consumption. High packet counts will incur processing load on the CPU and system memory resources of the appliance, even if the size of each packet is quite small.
- **Error rate**—The number of packets per second that cause errors. Errors may occur as a result of interference or poor link quality causing data corruption in frames. In general terms, error rates should be under 1%; high error rates may indicate a driver problem if a network media problem can be ruled out.
- **Discards/drops**—An interface may discard incoming and/or outgoing frames for several reasons. Each interface is likely to class the type of discard or drop separately to assist with troubleshooting the precise cause.



Some vendors may use the terms “discard” for frames that are rejected because of errors or security policies and “drop” for frames that are lost due to high load, but often the terms are used interchangeably.

- **Retransmissions**—Errors and discards/drops mean that frames of data are lost during transmission between two devices. As a result, the communication will be incomplete, and the data will, therefore, have to be retransmitted to ensure application data integrity. If you observe high levels of retransmissions (as a percentage of overall traffic), you must analyze and troubleshoot the specific cause of the underlying packet loss, which could involve multiple aspects of network configuration and connectivity.

Flow Data

As well as monitoring individual interface statistics, diagnosing performance issues depends on detailed information about network **traffic flows**. A packet analyzer can be used to measure network traffic statistics, but trying to record each frame imposes a heavy processing overhead on the network tap or mirror port. Collecting just the packet metadata, rather than the whole packet payload, reduces the bandwidth required by the sniffer. Technologies such as Cisco's **NetFlow** ([cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow/index.html](https://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow/index.html)) gather traffic flow data only and report it to a structured database. These technologies can also use sampling to further reduce processing demands. NetFlow has been redeveloped as the IP Flow Information Export (IPFIX) IETF standard (tools.ietf.org/html/rfc7011).

Using NetFlow involves deploying three types of components:

- A NetFlow exporter is configured on network appliances (switches, routers, and firewalls). Each flow is defined on an exporter. A traffic flow is defined by packets that share the same characteristics, such as IP source, destination addresses, and protocol type. These five bits of information are referred to as a 5-tuple. A 7-tuple flow adds the input interface and IP type of service data. Each exporter caches data for newly seen flows and sets a timer to determine flow expiration. When a flow expires or becomes inactive, the exporter transmits the data to a collector.
- A NetFlow collector aggregates flows from multiple exporters. A large network can generate huge volumes of flow traffic and data records, so the collector needs a high-bandwidth network link and substantial storage capacity. The exporter and collector must support compatible versions of NetFlow and/or IPFIX. The most widely deployed versions of NetFlow are v5 and v9.
- A NetFlow analyzer reports and interprets information by querying the collector and can be configured to generate alerts and notifications. In practical terms, the collector and analyzer components are often implemented as a single product.

Application	Protocol	Client	Server	Duration	Breakdown	Actual Thpt	Tot
TargusDataSp...	UDP	mon10 33577	172.16.0.201 5201	00:30	Client	0 bps	37
STUN.WhatsAp...	UDP	_gateway 34220	mon10 2055	00:56	Client	0 bps	147.:
ICMP	ICMP	mon10	_gateway	00:56	Client	0 bps	6
DNS	UDP	mon10 32888	dc10.corp.515support.com domain	00:25	Client Server	0 bps	624
DNS	UDP	mon10 43382	dc10.corp.515support.com domain	00:25	Client Server	0 bps	6;
DNS	UDP	mon10 43995	dc10.corp.515support.com domain	00:25	Client Server	0 bps	624
DNS	UDP	mon10 37313	dc10.corp.515support.com domain	00:25	Client Server	0 bps	6;
DNS	UDP	mon10 50749	dc10.corp.515support.com domain	00:23	Client Server	0 bps	624
DNS	UDP	mon10 58804	dc10.corp.515support.com domain	00:23	Client Server	0 bps	624
DNS	UDP	mon10 51449	dc10.corp.515support.com domain	00:23	Client Server	0 bps	624

*ntopng community edition being used to monitor NetFlow traffic data.
(Screenshot used courtesy of ntop.)*

Traffic Testing Tools

You might not always be able to observe poor performance directly. If this is the case, you can use various testers to simulate traffic loads.



This type of test can have a substantial impact on performance. Schedule this type of testing when the network is otherwise unused.

Throughput Testers

One fairly simple way to measure network throughput is to transfer a large file between two hosts. To determine your network throughput using this method, simply divide the file size by the amount of time taken to copy the file. For example, if you transfer a 1 GB file in half an hour, the throughput can be calculated as follows:

- 1 gigabyte is 1,024x1,024x1,024 bytes (1,073,741,824 bytes or 8,589,934,592 bits).
- 8,589,934,592 bits in 1,800 seconds is 4,772,186 bits per second or 4.55 Mbps.

This method derives a value that is different from the nominal data rate. Because two hosts are transferring the files between one another, it is the Application layers that handle the file transfer. The intervening layers on both hosts add complexity (headers) and introduce inaccuracy, such as corrupt frames that have to be retransmitted.


Several software utilities, such as **iperf** (iperf.fr), Ttcp (linux.die.net/man/1/ttcp), and BWPing (bwping.sourceforge.io), can be used to automate this testing process. An instance of the tool is configured on two network hosts, and the tools measure the throughput achieved between the sender and the listener.

```
mon@mon10:~/Desktop$ iperf3 -c 172.16.0.201 -u -b100M -t 30 -i 10
Connecting to host 172.16.0.201, port 5201
[ 5] local 10.1.16.242 port 35889 connected to 172.16.0.201 port 5201
[ ID] Interval      Transfer    Bitrate    Total Datagrams
[ 5] 0.00-10.00 sec  119 MBytes  100 Mbits/sec  86319
[ 5] 10.00-20.00 sec  119 MBytes  100 Mbits/sec  86325
[ 5] 20.00-30.00 sec  119 MBytes  100 Mbits/sec  86326
-----
[ ID] Interval      Transfer    Bitrate    Jitter    Lost/Total Datagrams
[ 5] 0.00-30.00 sec  358 MBytes  100 Mbits/sec  0.000 ms  0/258970 (0%) sender
[ 5] 0.00-30.00 sec  358 MBytes  100 Mbits/sec  0.016 ms  0/258970 (0%) receiver
iperf Done.
```

iperf3 transfer report showing bitrate, jitter, and packet loss. (Screenshot courtesy of iperf.)

Top Talkers/Listeners

Top talkers are the interfaces generating the most outgoing traffic (in terms of bandwidth), while top listeners are the interfaces receiving the most incoming traffic. Identifying these hosts and the routes they are using is useful in identifying and eliminating performance bottlenecks. Most network analyzer software comes with filters or built-in reporting to identify top talkers or top listeners.



The screenshot shows the Wireshark 'Endpoints' report for Ethernet. The report is a table with columns for Address, Packets, Bytes, Tx Packets, Tx Bytes, Rx Packets, and Rx Bytes. The first row is highlighted in blue, indicating the highest traffic volume.

Address	Packets	Bytes	Tx Packets	Tx Bytes	Rx Packets	Rx Bytes
00:15:5d:01:ca:93	52,249	505 M	35,647	10 M	16,602	494 M
00:15:5d:01:ca:92	58,145	517 M	19,427	506 M	38,718	11 M
00:15:5d:01:ca:94	4,631	11 M	2,637	501 k	1,994	11 M
00:15:5d:01:ca:91	1,525	230 k	712	111 k	813	119 k
ff:ff:ff:ff:ff:ff	192	20 k	0	0	192	20 k
33:33:00:01:00:02	53	8798	0	0	53	8798
33:33:00:00:00:0c	21	8226	0	0	21	8226
01:00:5e:7f:ff:fa	21	7878	0	0	21	7878
33:33:00:00:00:fb	2	178	0	0	2	178
33:33:00:01:00:03	2	168	0	0	2	168
01:00:5e:00:00:fb	2	138	0	0	2	138
01:00:5e:00:00:fc	2	128	0	0	2	128
33:33:ff:44:dc:c6	1	86	0	0	1	86

The Endpoints report in Wireshark can be used to identify top talkers and top listeners. (Screenshot courtesy of Wireshark.)

Bandwidth Speed Testers

In addition to testing performance on a local network, you may also want to test Internet links using some type of **bandwidth speed tester**. There are many Internet tools available for checking performance. The two main classes are the following:

- **Broadband speed checkers**—These test how fast the local broadband link to the Internet is. They are mostly designed for SOHO use. The tool will test downlink and uplink speeds, will test latency using ping, and can usually compare the results with neighboring properties and other users of the same ISP.
- **Website performance checkers**—These query a nominated website to work out how quickly pages load. One of the advantages of an online tool is that you can test your site's response times from the perspective of customers in different countries.

Bandwidth Management

Latency and jitter on the Internet are difficult to control because of the number of different parties that are involved (both caller networks plus any ISP transit networks). On a local network, delay is typically caused by congestion and contention:

- Congestion is where the network infrastructure is not capable of meeting the demands of peak load and starts to queue or drop packets.
- Contention is the ratio between demand for a service and its available capacity. For example, if 100 video conferencing hosts each requiring 10 Mbps share a 1 Gbps link, the contention ratio is 1:1 ($100 * 10 \text{ Mbps} = 1 \text{ Gbps}$). If there are 200 handsets, the ratio is 2:1. Contention is a planning issue. You might not expect all 200 hosts to be running conferences at the same time, and so you may accept the 2:1 ratio. You would use monitoring to determine if the ratio changes (if there are more hosts or they start to require more bandwidth).

If you observe unacceptable levels of congestion or contention, you can either provision higher bandwidth links and/or faster switches and routers, or you can use some sort of bandwidth management mechanism. For example, if you are running VoIP over your network and someone decides to copy a 40 GB file down from a server, the file transfer could potentially disrupt VoIP call quality. Without QoS, switches and routers forward traffic based on best effort or first-in, first-out, meaning that frames or packets are forwarded in the order in which they arrive. A QoS system identifies the packets or traffic streams belonging to a specific application, such as VoIP, and prioritizes them over other applications, such as file transfer.

Differentiated Services

The **Differentiated Services (DiffServ)** framework classifies each packet passing through a device. Router policies can then be defined to use the packet classification to prioritize the delivery. DiffServ is an IP (layer 3) service tagging mechanism. It uses the Type of Service field in the IPv4 header (Traffic Class in IPv6). The field is populated with a 6-byte DiffServ Code Point (DSCP) by either the sending host or by the router. Packets with the same DSCP and destination are referred to as behavior aggregates and allocated the same Per Hop Behavior (PHB) at each DiffServ-compatible router.

DiffServ traffic classes are typically grouped into three types:

- Best Effort.
- Assured Forwarding (which is broken down into sub-levels).
- Expedited Forwarding (which has the highest priority).

IEEE 802.1p

While DiffServ works at layer 3, IEEE **802.1p** can be used at layer 2 (independently or in conjunction with DiffServ) to classify and prioritize traffic passing over a switch or wireless access point. 802.1p defines a tagging mechanism within the 802.1Q VLAN field (it is also often referred to as 802.1Q/p). The 3-bit priority field is set to a value between 0 and 7. Most vendors map DSCP values to corresponding 802.1p values. For example, 7 and 6 can be reserved for network control (such as routing table updates), 5 and 4 map to expedited forwarding levels for two-way communications, 3 and 2 map to assured forwarding for streaming multimedia, and 1 and 0 for “ordinary” best-effort delivery.



As well as invoking the priority tag, VLAN infrastructure is often used for traffic management on local networks. For example, voice traffic might be allocated to a different VLAN than data traffic.

Traffic Shaping

Quality of service (QoS) is distinct from class of service (CoS). CoS mechanisms, such as DiffServ and 802.1p, categorize protocols into groups requiring different service levels and provide a tagging mechanism to identify a frame or packet's class. QoS allows fine-grained control over traffic parameters. For example, if a network link is congested, DiffServ and 802.1p cannot address it, but a protocol such as Multiprotocol Label Switching (MPLS) with QoS functionality can reserve required bandwidth and predetermine statistics such as acceptable packet loss and maximum latency and jitter when setting up the link.

In terms of QoS, network functions are commonly divided into three planes:

- **Control plane**—Makes decisions about how traffic should be prioritized and where it should be switched.
- **Data plane**—Handles the actual switching of traffic.
- **Management plane**—Monitors traffic conditions.

Protocols, appliances, and software that can apply these three functions can be described as traffic shapers or bandwidth shapers. **Traffic shapers** delay certain packet types—based on their content—to ensure that other packets have a higher priority. This can help to ensure that latency is reduced for critical applications.

Simpler devices performing traffic policing do not offer the enhanced traffic management functions of a shaper. For example, typical traffic policing devices will simply fail to deliver packets once the configured traffic threshold has been reached (this is often referred to as tail drop). Consequently, there will be times when packets are being lost, while other times when the network is relatively idle, and the bandwidth is being underutilized. A traffic shaper will store packets until there is free bandwidth available. Hopefully, this leads to consistent usage of the bandwidth and few lost packets.



It is essential that the selected device is capable of handling high traffic volumes. As these devices have a limited buffer, there will be situations when the buffer overflows. Devices can either drop packets and in essence provide traffic policing, or else they must implement a dropping algorithm. Random Early Detection (RED) is one of several algorithms that can be implemented to help manage traffic overflow on the shaper.

Module 8

Summary

You should be able to explain the purpose of organizational processes and procedures and use network monitoring technologies to troubleshoot performance issues.

Guidelines for Managing Networks

Setting up a management and monitoring system for a network can be a complex process, involving the evaluation and testing of different products. Follow these guidelines to make effective use of network management and monitoring tools:

- Set up policies, procedures, and tools for configuration and change management.
- Create a network asset inventory and diagrams that show the physical and logical configuration.
- Deploy IP and port scanners to gain visibility into hosts attached to the network and protocol traffic passing over it.
- Deploy packet capture and protocol analyzer software to gain visibility into individual packets and per-host or per-protocol statistics.
- Select a log collection system that will provide the best compatibility with the endpoints used on your network, plus the reporting and management features that you require.
- Configure endpoints to provide information to the log collector. This could involve one or more different methods, such as the following:
 - Configure SNMP traps.
 - Configure remote logging to use syslog or a similar protocol.
 - Configure NetFlow/IPFIX exporters.
 - Deploy agents to the endpoints to perform log and performance counter collection and measurement.
- Identify metrics to use to monitor network interfaces, device health and performance, plus network traffic levels.
- Record baseline measurements for the selected metrics.
- Set up filters to alert and notify administrators when key thresholds are exceeded or when hosts fail heartbeat tests.
- Set up a process for responding to alerts, reviewing logs, and diagnosing trends. Use this analysis to plan deployment of traffic marking (DiffServ/802.1p) and traffic shaping/bandwidth management solutions.

Module 9

Explaining Network Security Concepts

Module Introduction

You have identified the basic components and concepts for deploying and monitoring a network, but a network implementation is not complete without security mechanisms. In this module, you will describe basic concepts related to network security. As a networking professional, it is part of your responsibility to understand these fundamental concepts so that you can support network security controls.

Module Objectives

In this module, you will do the following:

- Explain common security concepts.
- Distinguish risk, vulnerability, exploit, and threat.
- Explain the importance of audits and regulatory compliance.
- Summarize types of attacks and their impact on the network.

Lesson 9.1

Security Concepts



Exam Objectives Covered

4.1 Explain the importance of basic network security concepts.

In this lesson, you will describe basic concepts related to security terminology and network security audits and assessments. It's important to have a solid foundation and awareness of the industry terminology used when you are discussing network security.

As you study this lesson, answer the following questions:

- How can you use the CIA triad to help create security policies and select security controls?
- What is meant by the terms “risk,” “vulnerability,” “exploit,” and “threat”?
- What are the types of security audits and assessments, and how are they affected by regulatory compliance?
- How does encryption underpin logical security and access controls?

Common Security Terminology

Establishing computer and network security means developing processes and controls that protect data assets and ensure business continuity by making network systems and hosts resilient to different kinds of attack. The computer industry uses common security terminology to describe some basic principles.

The Confidentiality, Integrity, and Availability Triad


One of the foundational principles of computer security is that the systems used to store, transmit, and process data must demonstrate three properties, often referred to as the **CIA triad**:

- Confidentiality means that certain information should only be known to certain people.
- Integrity means that the data is stored and transferred as intended and that any modification is authorized.
- Availability means that information is accessible to those authorized to view or modify it.

Vulnerability, Threat, and Risk

To assess how a system can demonstrate the CIA triad properties, security teams must identify ways in which their systems could be compromised. These assessments involve vulnerabilities, threats, and risk:

- **Vulnerability**—A weakness that could be accidentally triggered or intentionally exploited to cause a security breach.
- **Threat**—The potential for someone or something to exploit a vulnerability and breach security. A threat may be intentional or unintentional. The person or thing that poses the threat is called a threat actor or threat agent. The path or tool used by a malicious threat actor can be referred to as the attack vector.
- **Risk**—The likelihood and impact (or consequence) of a threat actor exercising a vulnerability. Assessing risk helps you decide which vulnerabilities to prioritize patching and what additional security measures to implement.



Relationship between vulnerability, threat, and risk.

Security Audits and Assessments

Many tools and techniques are available to ensure that network systems demonstrate the three properties of the CIA triad. Selection and deployment of these tools is guided by security policies. Security policies ensure that an organization has evaluated the risks it faces and has put security controls in place to mitigate those risks. Making a system more secure is also referred to as hardening. Different security policies should cover every aspect of an organization's use of computer and network technologies, from procurement and change control to acceptable use. Various types of audits and assessments are used to plan and monitor the implementation of these policies.

Risk Management

Risk management is a process for identifying, assessing, and mitigating vulnerabilities and threats to the essential functions that a business must perform to serve its customers. Risk management is complex and treated very differently in companies and institutions of different sizes, and with different regulatory and compliance requirements. Most companies will institute enterprise risk management (ERM) policies and procedures, based on published frameworks.

Risk assessment is a subset of risk management where the company's systems and procedures are audited for risk factors. Separate assessments can be devised to perform an initial evaluation and ongoing monitoring of threats, vulnerabilities, and security posture.

Posture Assessment

There are many different ways of thinking about how IT services should be governed to fulfill overall business needs. Some organizations have developed IT service frameworks to provide best practice guides to implementing IT and cybersecurity. These frameworks can shape company policies and provide checklists of procedures, activities, and technologies that should ideally be in place.

Collectively, these procedures, activities, and tools can be referred to as security controls. A security control is something designed to give a system or data asset the properties of confidentiality, integrity, availability, and non-repudiation.

In theory, security controls or countermeasures could be introduced to address every risk factor. The difficulty is that security controls can be expensive, so you must balance the cost of the control with the cost associated with the risk. It is not possible to eliminate risk; rather, the aim is to mitigate risk factors to the point where the organization is exposed only to a level of risk that it can afford. The overall status of risk management is referred to as risk posture. Risk posture shows which risk response options can be identified and prioritized. **Posture assessment** is often performed with reference to an IT or security framework. The framework can be used to assess the organization's maturity level in its use of security policies and controls.

Cybersecurity audits are comprehensive reviews designed to ensure an organization's security posture aligns with established standards and best practices. There are various types of cybersecurity audits, including compliance audits, which assess adherence to regulations; risk-based audits, which identify potential threats and vulnerabilities in an organization's systems and processes; and technical audits, which delve into the specifics of the organization's IT infrastructure, examining areas such as network security, access controls, and data protection measures.

Process Assessment

Mitigating risk can involve a large amount of expenditure so it is important to focus efforts. Effective risk management must focus on mission essential functions that could cause the whole business to fail if they are not performed. Part of this process involves identifying critical systems and assets that support these functions. A **mission essential function (MEF)** is one that cannot be deferred. This means that the organization must be able to perform the function as close to continually as possible, and if there is any service disruption, the mission essential functions must be restored first.

Business impact analysis (BIA) is the process of assessing what losses might occur for a range of threat scenarios. For instance, if a denial of service (DoS) attack suspends an e-commerce portal for five hours, the business impact analysis will be able to quantify the losses from orders not made and customers moving permanently to other suppliers based on historic data. The likelihood of a DoS attack can be assessed on an annualized basis to determine annualized impact, in terms of costs. You then have the information required to assess whether a security control, such as load balancing or managed attack mitigation, is worth the investment.

Regulatory Compliance

Internal security audits assess risk factors or ensure that an organization has deployed sufficient controls to protect its systems and data assets. Regulatory compliance imposes externally determined requirements on companies in certain industries or when processing certain types of data. These regulations might dictate the type of controls that must be deployed, and the type and frequency of audits. An organization might be subject to a compliance audit run by external auditors to verify that it is meeting the requirements of the regulations.

Personal Data and the General Data Protection Regulation

International, national, and state legislation can impose regulations on the collection and processing of personal data. **Personally identifiable information (PII)** is data that can be used to identify, contact, locate, or describe an individual. A Social Security number (SSN) is a good example of PII. Others include name, date of birth, email address, telephone number, street address, biometric data, and so on. Some bits of information, such as a SSN, may be unique; others uniquely identify an individual in combination (for example, full name with birth date and street address).

Privacy is the concept that collection and processing of personal information be both secure and fair. The European Union's **General Data Protection Regulation (GDPR)** is one example of privacy legislation governing the collection and processing of PII. GDPR means that personal data cannot be collected, processed, or retained without the individual's informed consent unless there are other overriding considerations, such as public interest or other legal obligations. *Informed consent* means that the data must be collected and processed only for the stated purpose, and that purpose must be clearly described to the user in plain language, not legal jargon. GDPR gives data subjects the right to withdraw consent, and to inspect, amend, or erase data held about them. Failure to comply with GDPR rules can result in large fines.

Data Locality

Some states and nations may respect privacy more or less than others; and likewise, some nations may disapprove of the nature and content of certain data. They may even be suspicious of security measures such as encryption. When your data is stored or transmitted in other jurisdictions, or when you collect data from citizens in other states or other countries, you may not "own" the data in the same way as you'd expect or like to.

Data sovereignty refers to a jurisdiction preventing or restricting processing and storage from taking place on systems that do not physically reside within that jurisdiction. For example, GDPR protections are extended to any EU citizen while they are within EU or EEA (European Economic Area) borders. Data subjects can consent to allow a transfer, but there must be a meaningful option for them to refuse consent. If the transfer destination jurisdiction does not provide adequate privacy regulations (to a level comparable to GDPR), then contractual safeguards must be given to extend GDPR rights to the data subject.

Data sovereignty may require you to implement data locality policies and tools. Data locality establishes storage and processing boundaries based on national or state borders. Most cloud storage and processing solutions offer data locality tools. For example, if a healthcare database is hosted in the cloud, data locality could be configured to prevent an administrator from replicating it to any datacenter outside the United States.

Payment Card Industry Data Security Standard

Privacy legislation is typically horizontal, meaning that it applies across a whole range of industries. There are also vertical standards for various types of regulated data that are specific to particular industries or commercial activities, such as healthcare or payment processing. Organizations that handle regulated data must comply with relevant laws and regulations governing its protection. Compliance typically involves implementing appropriate security measures, data encryption, access controls, data breach notification procedures, and data handling protocols. Organizations may also need to establish data storage, retention, and destruction safeguards to meet regulatory requirements.

One example is the **Payment Card Industry Data Security Standard (PCI DSS)** for credit card information. An organization that directly processes credit card transactions must adopt the PCI DSS standard to safeguard the cardholder data environment (CDE). Payment card information comprises the card number, expiry date, and the three-digit card verification value (CVV). Cards are also associated with a PIN, but this should never be transmitted to or handled by the merchant. Abuse of the card may also require the holder's name and the address to which the card is registered.

Encryption

An access control system ensures that a computer system or network meets the goals of the CIA triad. Access control governs how subjects may interact with objects. Subjects are people, devices, software processes, or any other system that can request and be granted access to a resource. Objects are the resources. An object could be a network, server, database, app, or file. Subjects are assigned rights or permissions on resources.

Many access control solutions depend on some type of encryption, or more generally some type of cryptographic technology. Encryption is an example of a logical security system. Logical security means that the system depends on software components, rather than the physical security of locks and intruder alarms. There are two main types of cryptographic cipher or algorithm:

- An **encryption algorithm** converts a human-readable plaintext into a ciphertext. A ciphertext must be decrypted using a key linked to the initial encryption process before it can be read. This makes data confidential, so long as the key is only available to authorized persons.
- A **cryptographic hash algorithm** converts a variable length string into a fixed-length hash. This hash cannot be converted back to a plaintext. This can prove the integrity of data (verifying that it has not been modified). It is also used for password storage and in other authentication solutions.

When deploying a cryptographic system to protect data assets, consideration must be given to all the ways that information could potentially be intercepted. Data can be described as being in one of three states:

- **Data at rest**—The state in which data is in some sort of persistent storage media.
- **Data in transit** (or data in motion)—The state in which data is transmitted over a network.
- **Data in use** (or data in processing)—The state in which data is present in volatile memory, such as system RAM or CPU registers and cache.

Different cryptographic solutions are used to protect data in these states. Data in transit can be protected by Transport Layer Security (TLS) encryption, for example. Data at rest can be protected by self-encrypting drives, file system encryption, or database encryption.

Vulnerability and Exploit Types

As software, any access control system can be subject to vulnerabilities. Vulnerabilities can exist because of misconfigurations or poor practice, but many people understand the term to mean faults in software specifically. A software vulnerability is a design flaw that can cause the application security system to be circumvented or that will cause the application to crash. The most serious vulnerabilities allow the attacker to execute arbitrary code on the system, which could allow the installation of malware or allow the threat actor to disable or weaken a secure configuration. Typically, applications such as web servers, web browsers, web browser plug-ins, email clients, and databases are targeted.

An **exploit** is the specific code or method of using a vulnerability to gain control of a system or damage it in some way. Typically, software vulnerabilities can be exploited only in quite specific circumstances, but because of the complexity of modern software and the speed with which new versions must be released to market, almost no software is free from vulnerabilities.

Zero-Day Vulnerabilities and Exploits

Most vulnerabilities are discovered by software and security researchers, who notify the vendor to give them time to patch the vulnerability before releasing details to the wider public. A vulnerability that is exploited before the developer knows about it or can release a patch is called a **zero-day**. These can be extremely destructive, as it can take the vendor a lot of time to develop a patch, leaving systems vulnerable for days, weeks, or even years.



The term “zero-day” is usually applied to the vulnerability itself but can also refer to an attack or malware that exploits it.

Unpatched and Legacy Systems

While an exploit for a zero-day vulnerability can be extremely destructive, they are relatively rare events. A greater threat is the large number of unpatched or legacy systems in use. An unpatched system is one that its owner has not updated with OS and application patches; a legacy system is one where the software vendor no longer provides support or fixes for problems.

Vulnerability Assessment

A **vulnerability assessment** is an evaluation of a system’s security and ability to meet compliance requirements based on the configuration state of the system. Essentially, the vulnerability assessment determines if the current configuration matches the ideal configuration (the baseline). Vulnerability assessments might involve manual inspection of security controls but are more often accomplished through automated vulnerability scanners.

Deception Technologies

Deception and disruption technologies are powerful cybersecurity resilience tools that significantly increase the attacker's cognitive load and resource expenditure by forcing them to constantly adapt their tactics, techniques, and procedures (TTPs). A **honeypot** is a decoy computer system designed to attract attackers. By analyzing their attack strategies and tools, honeypots provide early warning of attack attempts and valuable insights into attacker behavior. A honeynet is an entire decoy network. This may be set up as an actual network or simulated using an emulator.

A honeypot or honeynet is more likely to be located in a protected but untrusted area between the Internet and the private network or on a closely monitored and filtered segment within the private network itself. This provides early warning and evidence of whether a threat actor has been able to penetrate to a given security zone.

Lesson 9.2

Network Threats and Attacks



Exam Objectives Covered

4.1 Summarize various types of attacks and their impact to the network.

Effective network security design requires an understanding of how threat actors can compromise defenses. Threat research produces analysis of common types of attacks, or threat actor tactics, techniques, and procedures (TTPs). As threat actors seek ways to overcome security controls, novel attack vectors and types are developed all the time. Each organization and service/app has an attack surface, representing all the points and vulnerabilities that a threat actor could potentially exploit. To assist with the selection and deployment of security controls, you must be able to evaluate the characteristics of threat actors and attack types and determine their potential impact on the network.

As you study this lesson, answer the following questions:

- What are threat types, and what methods do threat actors use to achieve their goals?
- What are attack types, and what threat actor objectives do they work toward?
- What are the types of denial of service (DoS) and distributed DoS (DDoS)?
- What are the main malware classifications, and how has the use of malware evolved to give threat actors new capabilities?

Threat Types and Assessment

Exploits for vulnerabilities are either developed by threat actors or exposed by unintentional weaknesses in procedures. Threat assessment is the process of identifying threat sources and profiling the types and capabilities of threat actors.

External Versus Internal Threats

An external threat actor or agent is one that has no account or authorized access to the target system. A malicious external threat must infiltrate the security system using malware and/or social engineering. Note that an external actor may perpetrate an attack remotely or on-premises (by breaking into the company's headquarters, for instance). It is the threat actor that is defined as external, rather than the attack method.

Conversely, an internal (or insider) threat actor is one that has been granted permissions on the system. This typically means an employee, but an insider threat can also arise from contractors and business partners.

Threat Research

Threat research is a counterintelligence gathering effort in which security companies and researchers attempt to discover the tactics, techniques, and procedures (TTPs) of threat actors.

The outputs from the primary research undertaken by security solutions providers and academics can take three main forms:

- **Behavioral threat research**—Narrative commentary describing examples of attacks and TTPs gathered through primary research sources.
- **Reputational threat intelligence**—Lists of IP addresses and domains associated with malicious behavior, plus signatures of known file-based malware.
- **Threat data**—Computer data that can correlate events observed on a customer's own networks and logs with known TTP and threat actor indicators.

Attack Types

A network can be attacked by many kinds of threat actors for many different reasons, and different attacks can have various kinds of impact. The goals of most types of adversaries will be to steal (exfiltrate) information from the network, to misuse network services (for fraud, for instance), or to compromise the availability of the network. Insider threat-type attacks may be launched with privileged access to the network, while external threats must find some way of accessing the network, perhaps by installing malware on a host system.

Footprinting and Fingerprinting Attacks

Footprinting and fingerprinting are **enumeration** or information gathering attacks. Footprinting allows a threat actor to discover the topology and general configuration of the network and security systems. Footprinting can be done by social engineering attacks—persuading users to give information or locating information that has been thrown out as trash, for instance. Port scanning specifically aims to enumerate the TCP or UDP application ports on which a host will accept connections.

Fingerprinting allows a threat actor to identify device and OS types and versions. When a host running a particular operating system responds to a port scan, the syntax of the response might identify the specific operating system. This fact is also true of application servers, such as web servers, FTP servers, and mail servers. The responses these servers make often include headers or banners that can reveal a great deal of information about the server. A threat actor can use this information to probe for known vulnerabilities.

Spoofing Attacks

The term “**spoofing**” covers a wide range of different attacks. Spoofing can include any type of attack where the threat actor disguises their identity, or in which the source of network information is forged to appear legitimate. Social engineering and techniques such as phishing and pharming, where the attacker sets up a false website in imitation of a real one, are types of spoofing attacks. It is also possible to abuse the way a protocol works or how network packets are constructed to inject false or modified data onto a network. ARP and DNS services are often used as vectors for this type of attack.

Denial of Service Attacks

A **denial of service (DoS)** attack causes a service at a given host to fail or to become unavailable to legitimate users. Resource exhaustion DoS attacks focus on overloading a service by using up CPU, system RAM, disk space, or network bandwidth. It is also possible for DoS attacks to exploit design failures or other vulnerabilities in application software. A physical DoS attack might involve cutting telephone lines or network cabling or switching off the power to a server. DoS attacks may be motivated by the malicious desire to cause trouble. They may also be part of a wider attack, such as the precursor to a spoofing or data exfiltration attack. DoS can assist these attacks by diverting attention and resources away from the real target. For example, a blinding attack attempts to overload a logging or alerting system with events.



When we describe attacks, there's a focus on purposefully malicious threat actors, but do be aware of inadvertent, accidental, and non-malicious threats. For example, a user might cause accidental DoS by connecting two wall ports and creating a switching loop. Users can also create inadvertent vulnerabilities. For example, shadow IT (devices or apps used in the workplace without authorization) could be vectors for exploits that aren't mitigated by security controls.

Distributed DOS Attacks and Botnets

A **distributed DoS (DDoS)** attack is launched simultaneously by multiple hosts. Some types of DDoS attacks simply aim to consume network bandwidth, denying it to legitimate hosts. Others cause resource exhaustion on the hosts processing requests, consuming CPU cycles and memory. This delays processing of legitimate traffic and could potentially crash the host system completely. For example, a SYN flood attack works by withholding the client's ACK packet during TCP's three-way handshake. The client's IP address is spoofed, meaning that an invalid or random IP is entered so the server's SYN/ACK packet is misdirected. A server can maintain a queue of pending connections. When it does not receive an ACK packet from the client, it resends the SYN/ACK packet a set number of times before timing out the connection. The problem is that a server may only be able to manage a limited number of pending connections, which the DoS attack quickly fills up. This means that the server is unable to respond to genuine requests.

Distributed Reflection DoS/Amplification Attacks

A more powerful TCP SYN flood attack is a type of distributed reflection DoS (DRDoS) or amplification attack. In this attack, the adversary spoofs the victim's IP address and attempts to open connections with multiple servers. Those servers direct their SYN/ACK responses to the victim server. This rapidly consumes the victim's available bandwidth.

One example of this technique is to bombard a victim network with responses to bogus DNS queries. One of the advantages of this technique is that while the request is small, the response to a DNS query can be made to include a lot of information, so this is a very effective way of overwhelming the bandwidth of the victim network with much more limited resources on the attacker's network. The Network Time Protocol (NTP) can be abused in a similar way.

Botnets

A **botnet** is a group of compromised hosts that can be used to launch DDoS and DRDoS attacks. A threat actor will first compromise one or two machines to use as handlers or herders. The handlers are used to compromise hundreds, thousands, or even millions of zombie hosts and install DDoS tools on them (the bots). To compromise a host, the attacker must install malware that opens a backdoor remote connection. The attacker can then use the malware to install bots and trigger the zombies to launch the attack at the same time. The network established between the handlers and the bots is called a **command and control (C2 or C&C)** network.



Any type of Internet-enabled device is vulnerable to compromise. This includes web-enabled cameras, SOHO routers, smart TVs, and other appliances. This is referred to as an internet of things (IoT) botnet.

Malware Attacks

Many of the intrusion attempts perpetrated against computer networks depend on the use of malicious software, or **malware**. Malware can be defined simply as software that does something bad, from the perspective of the system owner. There are many types of malware, but they are not classified in a rigorous way, so some definitions overlap or are blurred. Some malware classifications, such as Trojan, virus, and worm, focus on the vector used by the malware. The vector is the method by which the malware executes on a computer and potentially spreads to other network hosts. Another complicating factor with malware classification is the degree to which its installation is expected or tolerated by the user. The following categories describe some types of malware according to vector:

- **Viruses and worms**—These represent some of the first types of malware and spread without any authorization from the user by being concealed within the executable code of another process. Viruses infect files, while worms can infect processes running in system memory.
- **Trojan**—Malware concealed within an installer package for software that appears to be legitimate. This type of malware does not seek any type of consent for installation and is actively designed to operate secretly.
- **Potentially unwanted programs (PUPs)/Potentially unwanted applications (PUAs)**—Software installed alongside a package selected by the user or perhaps bundled with a new computer system. Unlike a Trojan, the presence of a PUP is not automatically regarded as malicious. It may have been installed without active consent or consent from a purposefully confusing license agreement. This type of software is sometimes described as grayware rather than malware.

Other classifications are based on the payload delivered by the malware. The payload is an action performed by the malware other than simply replicating or persisting on a host. Examples of payload classifications include spyware, rootkit, remote access Trojan (RAT) or backdoor, and ransomware.

As malware has continued to be developed for criminal intent and security software became better able to detect and block known viruses, worms, and Trojans, malware code and techniques have become more sophisticated. The term “fileless” has gained prominence to refer to these modern types of malware. Fileless is not a definitive classification, but it describes a collection of common behaviors and techniques:

- Fileless malware does not write its code to disk. The malware uses memory-resident techniques to run in its own process, within a host process or dynamic link library (DLL), or within a scripting host. This does not mean that there is no disk activity at all, however. The malware may change registry values to achieve persistence (executing if the host computer is restarted). The initial execution of the malware may also depend on the user running a downloaded script, file attachment, or Trojan software package.
- Fileless malware uses lightweight **shellcode** to achieve a backdoor mechanism on the host. The shellcode is easy to recompile in an obfuscated form to evade detection by scanners. It is then able to download additional packages or payloads to achieve the threat actor’s objectives. These packages can also be obfuscated, streamed, and compiled on the fly to evade automated detection.
- Fileless malware may use “live off the land” techniques rather than compiled executables to evade detection. This means that the malware code uses legitimate system scripting tools, notably PowerShell and Windows Management Instrumentation (WMI), to execute payload actions. If they can be executed with sufficient permissions, these environments provide all the tools the attacker needs to perform scanning, reconfigure settings, and exfiltrate data.

The terms “**advanced persistent threat (APT)**” and “advanced volatile threat (AVT)” can be used to describe this general class of modern fileless/live off the land malware. Another useful classification is low-observable characteristics (LOC) attack. The exact classification is less important than the realization that adversaries can use any variety of coding tricks to effect intrusions and that their tactics, techniques, and procedures to evade detection are continually evolving.

Lesson 9.3

Spoofing Attacks



Exam Objectives Covered

4.2 Summarize various types of attacks and their impact to the network.

Spoofing encompasses a very wide variety of attack types. When considering threats to the network, it is helpful to use the OSI model to determine what vulnerabilities and potential exploits exist at different layers. Spoofing attacks that corrupt forwarding at layer 2 (Data Link) can have substantial impact and be difficult to mitigate.

As you study this lesson, answer the following questions:

- What are the impacts to the network from on-path spoofing attacks?
- What are some common attacks that use MAC and ARP spoofing or flooding?
- How can a MAC address be spoofed if it is hard coded on the NIC?
- How can an attacker manipulate an ARP cache to redirect frames to capture all outbound data?
- How can an attacker manipulate VLANs to gain access that wouldn't normally be permitted?

On-path Attacks

An **on-path** attack is a specific type of spoofing attack where a threat actor compromises the connection between two hosts and transparently intercepts and relays all communications between them. The threat actor might also have the opportunity to modify the traffic before relaying it.



On-path attacks are also known by the term "Man-in-the-Middle (MitM)." Such terms are non-inclusive and/or use inappropriate or vague metaphors and are deprecated in the latest CompTIA exam objectives documents. The terms "Manipulator in the Middle," "Machine in the Middle," and "Adversary in the Middle (AiTM)" are also used as replacements.

MAC Spoofing and IP Spoofing

A host can arbitrarily select any MAC and/or IP address and attempt to use it on the network. While each network interface has a burned-in MAC address, this can be changed to any arbitrary value using packet crafting software. A threat actor might exploit this to spoof the value of a valid MAC or IP address to try to circumvent an access control list or impersonate a legitimate server. For this type of attack to succeed, the threat actor must normally disable the legitimate host or there will be duplicate addresses on the network, which will have unpredictable results.



IP spoofing is also used in most denial of service (DoS) attacks to mask the origin of the attack and make it harder for the target system to block packets from the attacking system. In this type of spoofing, the threat actor does not care that they will not receive replies, so it is different from an on-path attack.

ARP Spoofing and ARP Poisoning

ARP spoofing is a common means of perpetrating an on-path attack. It works by broadcasting unsolicited ARP reply packets, also known as gratuitous ARP replies, with a source address that spoofs a legitimate host or router interface. Because ARP has no security, all devices in the same broadcast domain as the rogue host trust this communication and update their MAC:IP address cache table with the spoofed address. The ARP cache is said to be poisoned. Because the threat actor broadcasts endless ARP replies, it overwhelms the legitimate interface.

No.	Time	Source	Destination	Protocol	Length	Info
6	10.022521400	Microsoft_01:ca:4a	Microsoft_01:ca:76	ARP	42	10.1.0.102 is at 00:15:5d:01:ca:76
7	10.032593900	Microsoft_01:ca:4a	Microsoft_01:ca:77	ARP	42	10.1.0.2 is at 00:15:5d:01:ca:77
8	10.032605300	Microsoft_01:ca:4a	Microsoft_01:ca:76	ARP	42	10.1.0.101 is at 00:15:5d:01:ca:76
9	18.219200600	10.1.0.101	10.1.0.2	TCP	66	1702 → 80 [SYN] Seq=0 Win=65535 Len=0
10	18.220473400	10.1.0.101	10.1.0.2	TCP	66	[TCP Out-Of-Order] 1702 → 80
11	18.223616200	10.1.0.2	10.1.0.101	TCP	66	80 → 1702 [SYN, ACK] Seq=0 Ack=1702 Win=65535 Len=0
12	18.228456800	10.1.0.2	10.1.0.101	TCP	66	[TCP Retransmission] 80 → 1702
13	18.228797700	10.1.0.101	10.1.0.2	TCP	54	1702 → 80 [ACK] Seq=1 Ack=1702 Win=0 Len=0
14	18.229264100	10.1.0.101	10.1.0.2	HTTP	433	GET / HTTP/1.1
15	18.238162600	10.1.0.101	10.1.0.2	TCP	54	1702 → 80 [ACK] Seq=1 Ack=1702 Win=0 Len=0
16	18.238250400	10.1.0.101	10.1.0.2	TCP	433	[TCP Retransmission] 1702 → 80
17	18.239342200	10.1.0.2	10.1.0.101	HTTP	412	HTTP/1.1 302 Redirect (text/html)
18	18.244530700	10.1.0.2	10.1.0.101	TCP	412	[TCP Retransmission] 80 → 1702
19	18.245021200	10.1.0.101	10.1.0.2	TCP	54	1702 → 80 [ACK] Seq=380 Ack=1702 Win=0 Len=0
20	18.252481800	10.1.0.101	10.1.0.2	TCP	54	[TCP Dup ACK 19#1] 1702 → 80
21	18.255190400	10.1.0.101	10.1.0.2	TCP	66	1703 → 443 [SYN] Seq=0 Win=65535 Len=0
22	18.260503200	10.1.0.101	10.1.0.2	TCP	66	[TCP Retransmission] 1703 → 443
23	18.261065300	10.1.0.2	10.1.0.101	TCP	66	443 → 1703 [SYN, ACK] Seq=0 Ack=1703 Win=65535 Len=0
24	18.268454300	10.1.0.2	10.1.0.101	TCP	66	[TCP Retransmission] 443 → 1703

▶ Frame 9: 66 bytes on wire (528 bits), 66 bytes captured (528 bits) on interface 0	
▶ Ethernet II, Src: Microsoft_01:ca:77 (00:15:5d:01:ca:77), Dst: Microsoft_01:ca:4a (00:15:5d:01:ca:4a)	
▶ Destination: Microsoft_01:ca:4a (00:15:5d:01:ca:4a)	
▶ Source: Microsoft_01:ca:77 (00:15:5d:01:ca:77)	
Type: IPv4 (0x0800)	
▶ Internet Protocol Version 4, Src: 10.1.0.101, Dst: 10.1.0.2	
▶ Transmission Control Protocol, Src Port: 1702, Dst Port: 80, Seq: 0, Len: 0	

0000	00 15 5d 01 ca 4a 00 15 5d 01 ca 77 08 00 45 00	..w..E..
0010	00 34 1c ca 40 00 80 06 c9 91 0a 01 00 65 0a 01	.4..@... ..e..
0020	00 02 06 a6 00 50 dc 52 ee 41 00 00 00 00 80 02	P.R.A.....
0030	ff ff 89 1d 00 00 02 04 05 b4 01 03 03 08 01 01	
0040	04 02	..

Observing ARP poisoning in a Wireshark packet capture. (Screenshot courtesy of Wireshark.)

The usual target will be the subnet's default gateway. If the attack is successful, all traffic destined for remote networks will be sent to the attacker. The threat actor can then perform an on-path attack to monitor the communications and continue to forward them to the router to avoid detection. The attacker could also modify the packets before forwarding them. ARP poisoning could also perform a DoS attack by not forwarding the packets.

ARP spoofing can be difficult to detect without closely monitoring network traffic. However, attempts at ARP spoofing are likely to cause sporadic communications difficulties, such as an unreachable default gateway. In such cases, performing network captures and examining ARP packets may reveal the poison packets, as will examining local ARP caches for multiple IP addresses mapping to the same MAC address.



Technically, ARP spoofing is the broadcast of the unsolicited ARP replies, while ARP poisoning is the injection of spoofed MAC:IP mappings into the victim cache. The terms are often just used interchangeably, however. Be aware that ARP poisoning could include other methods of injecting fake mappings, such as the local host being infected with malware.



While IPv6 does not use ARP, it is also vulnerable to layer 2 spoofing if the unencrypted Neighbor Discovery (ND) Protocol is used. Abuse of this can be used for router advertisement (RA) spoofing.

MAC Flooding Attack

Where ARP poisoning is directed at hosts, **MAC flooding** is used to attack a switch. The intention of the attacker is to exhaust the memory used to store the switch's MAC address table. The switch uses the MAC address table to determine which port to use to forward unicast traffic to its correct destination. Overwhelming the table can cause the switch to stop trying to apply MAC-based forwarding and flood unicast traffic out of all ports, working as a hub. This makes sniffing network traffic easier for the threat actor.

VLAN Hopping Attacks

VLAN hopping is an attack designed to send traffic to a VLAN other than the one the host system is in. This exploits the default VLAN feature of 802.1Q. Default VLANs are designed to provide compatibility with non-VLAN capable switches. The attacker, using a device placed in the default VLAN, crafts a frame with two VLAN tag headers. The first trunk switch to inspect the frame strips the first header, and the frame gets forwarded to the target VLAN. Such an attack can only send packets one way but could be used to perform a DoS attack against a host on a different VLAN.

A VLAN hopping attack can also be launched by attaching a device that spoofs the operation of a switch to the network and negotiating the creation of a trunk port. As a trunk port, the attacker's device will receive all inter-VLAN traffic.

Another common switch attack is known as a Spanning Tree Protocol (STP) manipulation attack. STP is normally configured on a network with several switches. The primary purpose is to prevent switching loops. To make STP work, a single switch is designated as the root bridge. If an attacker can become the root bridge, they are then able to see a variety of frames that they normally wouldn't see. To perpetrate this attack, the attacker inserts their switch into the tree and manipulates it to appoint their switch as the root bridge. By doing this, they can use a sniffer to collect data traversing the network.

Lesson 9.4

Rogue System Attacks



Exam Objectives Covered

4.2 Summarize various types of attacks and their impact to the network.

Spoofing attacks can also be launched at the Application layer. Such attacks often involve the use of a rogue system. A rogue system is hardware or software that spoofs legitimate services, such as DNS or DHCP.

As you study this lesson, answer the following questions:

- What are impacts on the network that arise from rogue DHCP services?
- What spoofing and poisoning attack types can compromise DNS services?

Rogue Devices and Services

When a device or service on your network isn't under the administrative control of the network staff, it's called a rogue. Rogue devices and services are often completely malicious. They exist for the sole purpose of stealing sensitive information such as credit card numbers and passwords. Most legitimate network hardware and services can be exploited through rogues. Rogue devices and services could include wireless access point, DHCP servers, DNS servers, and so on. Ordinary PCs and laptops can be maliciously configured to run any type of service, or they could use packet crafting software to perpetrate spoofing attacks.

Threats and vulnerabilities can also arise from unintentional sources. In the context of rogue devices and services, one example of an unintentional insider threat is the concept of **shadow IT**, where users purchase or introduce computer hardware or software to the workplace without the sanction of the IT department and without going through a procurement and security analysis process. The problem of shadow IT is exacerbated by the proliferation of cloud services and mobile devices, which are easy for users to obtain. Shadow IT creates a new unmonitored attack surface for malicious adversaries to exploit.

Rogue DHCP

The Dynamic Host Configuration Protocol (DHCP) provides IP addressing autoconfiguration to hosts. If a Windows client fails to obtain a DHCP lease, it defaults to using an address in the Automatic Private IP Addressing (APIPA) range of 169.254.0.0/16. It will be limited to communication with other APIPA hosts on the same network segment (broadcast domain). Linux hosts will use the 169.254.0.0/16 range if they have Zeroconf support, leave the IP address set to 0.0.0.0, or disable IPv4 on the interface.



APIPA is Microsoft terminology. Standards documentation refers to this address range as IPv4 link local (IPv4LL). Zeroconf is a standards-based approach to technologies that allows hosts to obtain a usable network configuration and discover services automatically and without the use of DHCP or DNS server infrastructure.

Clients have no means of preferring a DHCP server. If two DHCP servers are running on the same subnet, clients could end up with an incorrect IP configuration because they have obtained a lease from a rogue server. A rogue DHCP server may be deployed accidentally (forgetting to disable a DHCP server in an access point or router, for instance) or may be used by a malicious threat actor to subvert the network. A threat actor would normally use a rogue server to change the default gateway and/or DNS resolver addresses for the subnet and route communications via their machine. This is a means of using DHCP to facilitate an on-path attack.

A DHCP starvation attack uses bogus requests to use up leases in a legitimate DHCP server's address pool. An exhausted DHCP scope means legitimate hosts cannot obtain a lease. A DHCP starvation attack might be a denial of service (DoS) mechanism or be used to force legitimate hosts to obtain a lease from a rogue DHCP server.

DNS Attacks

The Domain Name System (DNS) resolves requests for named host and services to IP addresses. Name resolution is a critical addressing method on the Internet and on private networks. There are many potential attacks against DNS. On the public Internet, attacks might use typosquatting techniques to cause victims to confuse malicious sites with legitimate ones. DNS can be exploited in a DRDoS attack. Threat actors can also directly target public DNS services as a means of performing DoS against a website or cloud resource. Finally, a threat actor might be able to hijack a public DNS server and insert poisoned records, directing victims to rogue websites.

On a private network, a DNS attack is likely to mean some sort of DNS spoofing or **DNS poisoning**. These DNS attacks compromise the process by which clients query name servers to locate the IP address for a domain name.



As with ARP spoofing/poisoning, DNS spoofing and poisoning attacks are often taken to mean the same thing, but technically spoofing is using false DNS requests or replies or running a rogue DNS service, while poisoning is manipulating cached records.

DNS-Based On-Path Attacks

If the threat actor has access to the same local network as the victim, the attacker can use ARP poisoning to respond to DNS queries from the victim with spoofed replies. This might be combined with a denial of service attack on the victim's legitimate DNS server. A rogue DHCP could be used to configure clients with the address of a DNS resolver controlled by the threat actor.

DNS Client Cache Poisoning

Before DNS was developed in the 1980s, name resolution took place using a text file named HOSTS. Each name:IP address mapping was recorded in this file, and systems administrators had to download the latest copy and install it on each Internet client or server manually. Even though most name resolution now functions through DNS, the HOSTS file is still present and most operating systems check the file before using DNS or load it into a cache of known name:IP mappings, and the client only contacts a DNS server if the name is not cached. Therefore, if an attacker is able to place a false name:IP address mapping in the HOSTS file and effectively poison the DNS cache, they will be able to redirect traffic. The HOSTS file requires administrator access to modify. In UNIX and Linux systems it is stored as `/etc/hosts`, while in Windows it is placed in `%SystemRoot%\System32\Drivers\etc\hosts`. The presence of suspect entries in the HOSTS file is an indicator that the machine has been compromised.

DNS Server Cache Poisoning

DNS server cache poisoning aims to corrupt the records held by the DNS server itself. One attack method involves getting the victim name server to respond to a recursive query from the attacking host. A recursive query compels the DNS server to query the authoritative server for the answer on behalf of the client. The attacker's DNS, masquerading as the authoritative name server, responds with the answer to the query, but also includes a lot of false domain:IP mappings for other domains that the victim DNS accepts as genuine. The `nslookup` or `dig` tool can be used to query the name records and cached records held by a server to discover whether any false records have been inserted.

```

HOSTNAME    www.web.local    yes    Hostname to hijack
INTERFACE   no              The name of the interface
NEWADDR     192.168.2.192   yes    New address for hostname
RECONS      192.168.2.254   yes    The nameserver used for reconnaissance
RHOST       192.168.1.1     yes    The target address
SNAPLEN     65535          The number of bytes to capture
SRCADDR     Real           The source address to use for sending t
he queries (Accepted: Real, Random)
SRCPORT     0              The target server's source query port (
0 for automatic)
TIMEOUT     500            The number of seconds to wait for new d
ata
TTL         46348     The TTL for the malicious host entry
XIDS        0          The number of XIDS to try for each quer
y (0 for automatic)

msf auxiliary(bailiwicked_host) > run

[-] Failure: This hostname is already in the target cache: www.web.local
[-] Cache entry expires on 2017-09-17 09:08:17 -0700... sleeping.
^C[-] Auxiliary interrupted by the console user
[*] Auxiliary module execution completed
msf auxiliary(bailiwicked_host) > set hostname updates.web.local
hostname => updates.web.local
msf auxiliary(bailiwicked_host) > run

[*] Targeting nameserver 192.168.1.1 for injection of updates.web.local. as 192.
168.2.192
[*] Querying recon nameserver for web.local.'s nameservers...
[*] Got an NS record: web.local. 604800 IN NS ns.web.lo
cal.
[*] Querying recon nameserver for address of ns.web.local....
[*] Got an A record: ns.web.local. 604800 IN A 192.168.
1.1
[*] Checking Authoritativeness: Querying 192.168.1.1 for web.local....
[*] ns.web.local. is authoritative for web.local., adding to list of nameser
vers to spoof as
[*] Calculating the number of spoofed replies to send per query...
[*] race calc: 100 queries | min/max/avg time: 0.0/0.0/0.0 | min/max/avg repli
es: 0/1/0
[*] The server did not reply, giving up.
[*] Auxiliary module execution completed
msf auxiliary(bailiwicked_host) >

```

Attempting to poison a DNS server cache—This attack has failed.

Lesson 9.5

Social Engineering



Exam Objectives Covered

4.2 Summarize various types of attacks and their impact to the network.

People—employees, contractors, suppliers, and customers—represent part of the attack surface of any organization. A person with permissions on the system is a potential target of social engineering. Being able to compare and contrast social engineering techniques will help you to lead security awareness training and to develop policies and other security controls to mitigate these risks.

As you study this lesson, answer the following questions:

- What are the different types of social engineering attacks, and why are they so common and effective?
- What vulnerabilities are exposed by using password-based authentication?

Social Engineering Attacks


Threat actors can use a diverse range of techniques to compromise a security system. A prerequisite of many types of attacks is to obtain information about the network and security system. **Social engineering** (or hacking the human) refers to a collection of techniques and tricks designed to gain some type of unauthorized access to premises or data. Passive social engineering takes advantage of the unintentional actions of others to gather information or gain access to a secure facility. Active social engineering involves direct interaction with users, asking them to disclose information or take actions.

Impersonation (pretending to be someone else) is one of the basic social engineering techniques. The classic impersonation attack is for the threat actor to phone into a department pretending to be calling from IT support, claim they have to adjust something on the user's system remotely, and get the user to reveal their password. For this attack to succeed, the approach must be persuasive and establish trust. Social engineering might also use intimidation or hoaxes as a means of eliciting information.

Phishing Attacks

Phishing is a combination of social engineering and spoofing. It persuades or tricks the target into interacting with a malicious resource disguised as a trusted one, traditionally using email as the vector. A phishing message might try to convince the user to perform some action, such as installing disguised malware or allowing a remote access connection by the attacker.

Other types of phishing campaigns use a spoof website set up to imitate a bank or e-commerce site or some other web resource that should be trusted by the target. The attacker then emails users of the genuine website to inform them that their account must be updated or with a hoax alert or alarm. This message will contain a disguised link that actually leads to the spoofed site. When the user authenticates with the spoofed site, their login credentials are captured.



Example phishing email—on the right, you can see the message in its true form as the mail client has stripped out the formatting (shown on the left) designed to disguise the nature of the links.

Shoulder Surfing Attack

A threat actor can learn a password or PIN (or other secure information) by watching the user type it. This is referred to as a **shoulder surfing** attack. Despite the name, the attacker may not have to be in close proximity to the target—they could use high-powered binoculars or CCTV to directly observe the target remotely.

Tailgating and Piggybacking Attacks

Tailgating is a means of entering a secure area without authorization by following closely behind the person that has been allowed to open the door or checkpoint. Piggybacking is a similar situation but means that the attacker enters a secure area with an employee's permission. For instance, an attacker might impersonate a member of the cleaning crew and request that an employee hold the door open while they bring in a cleaning cart or mop bucket. Another technique is to persuade someone to hold a door open, using an excuse such as "I've forgotten my badge (or key)." Alternatively, piggybacking may be a means of an insider threat actor to allow access to someone without recording it in the building's entry log.

Dumpster Diving Attack

Dumpster diving refers to combing through an organization's (or individual's) garbage to try to find useful documents (or even files stored on discarded removable media).




A threat actor might stage attacks over a long period. Initial attacks may only aim at compromising low-level information and user accounts, but this low-level information can be used to attack more sensitive and confidential data and better protected management and administrative accounts.

Password Attacks

On-path and malware attacks can be difficult to perpetrate. Many network intrusions occur because a threat actor is able to obtain credentials to access the network. Also, when a threat actor gains some sort of access via an on-path or malware attack, they are likely to attempt to escalate privileges to gain access to other targets on the network by harvesting credentials for administrative accounts.

Passwords or password hashes can be captured by obtaining a password file or by sniffing the network. If the protocol uses cleartext credentials, then the threat actor can simply read the cleartext password from the captured frames.



If authentication credentials are transmitted in cleartext, such as the unencrypted version of the IMAP mailbox access protocol, it is a simple matter for the credentials to be intercepted via packet sniffing. (Screenshot courtesy of Wireshark.)



A password might be sent in an encoded form, such as Base64, which is simply an ASCII representation of binary data. This is not the same as encryption. The password value can easily be derived from the Base64 string.

In most cases, a password is stored and transmitted securely by making a cryptographic hash of the string entered by the user. A cryptographic hash algorithm, such as Secure Hash Algorithm (SHA) or Message Digest v5 (MD5), produces a fixed-length string from a variable-length string. This means that, in theory, no one except the user (not even the system administrator) knows the password, because the plaintext should not be recoverable from the hash.

Password cracking software uses various methods to work out the plaintext password string from a cryptographic hash:

- **Dictionary**—The software matches the hash to those produced by ordinary words found in a dictionary. This could also include information such as user and company names, pet names, or any other data that people might naively use as passwords.
- **Brute force**—The software tries to match the hash against one of every possible combination it could be. If the password is short (under eight characters) and non-complex (using only letters, for instance), a password might be cracked in minutes. Longer and more complex passwords increase the amount of time the attack takes to run.

A threat actor might obtain password hashes from a protocol such as SMB with no encryption configured. The risks posed by cracking software mean that it is more secure to use end-to-end encryption, such as IPSec or Transport Layer Security (TLS). This means that all payload data is encrypted, and a network sniffer cannot even recover the password hashes.

Module 9

Summary

You should be able to explain common security concepts relating to risk, vulnerabilities, exploits, threats, and attacks.

Guidelines for Supporting Security Planning and Auditing

Follow these guidelines to support security audits and assessments:

- Establish security policies and deploy security controls that address the CIA triad.
- Deploy assessment and monitoring processes and tools to evaluate vulnerabilities, threats, and risk:
 - Overall risk and posture assessment for mission essential functions (MEF) to produce business impact analysis, business continuity plans, and security policies, such as privileged access management and vendor assessment.
 - Vulnerability assessment to analyze systems for misconfigurations and missing patches.
 - Threat assessment to develop awareness of tactics, techniques, and procedures (TTPs) and obtain threat data feeds for automated detection via SIEM.
- Develop awareness of common attacks such as footprinting/fingerprinting, spoofing, DoS/DDoS, on-path, DNS spoofing/poisoning, VLAN hopping, ARP spoofing/poisoning, rogue DHCP/DNS, malware, password, and social engineering.

Module 10

Applying Network Security Features

Module Introduction

Authentication and authorization policies and systems enforce access control. Access control ensures that an organization's data is processed and stored securely. During your career in network support, you will often have to configure accounts and permissions and troubleshoot issues arising from access control.

Module Objectives

In this module, you will do the following:

- Explain identity and access management concepts.
- Distinguish protocols and standards used for authentication and directory management.
- Use defense in depth techniques to ensure that only policy-compliant devices can connect to the network.
- Apply security rules, such as ACLs and content filtering, to manage network traffic.

Lesson 10.1

Authentication



Exam Objectives Covered

4.1 Explain the importance of basic security concepts.

4.3 Given a scenario, apply network security features, defense techniques, and solutions.

Strong authentication is the first line of defense to secure network resources. As a network professional, to effectively manage authentication on your network, you will need to understand these different systems and what each one can provide for your organization.

As you study this lesson, answer the following questions:


- What is an authentication factor, and what different factors are available?
- What is the difference between local and remote authentication?
- What is the relationship between authentication, authorization, and single sign-on?
- How do certificates and PKI support authentication?
- What roles do RADIUS, TACACS+, and SAML play in authentication?

Access Control

Access control governs how subjects/principals may interact with objects. When implemented on a computer system or network, access control is a type of logical security. Modern access control is typically implemented as an **identity and access management (IAM)** system. IAM comprises four main processes:

- **Identification**—Creating an account or ID that uniquely represents the user, device, or process on the network.
- **Authentication**—Proving that a subject is who or what it claims to be when it attempts to access the resource. An authentication factor determines what sort of credential the subject can use. For example, people might be authenticated by providing a password; a computer system could be authenticated using a token such as a digital certificate.
- **Authorization**—Determining what rights subjects should have on each resource, and enforcing those rights. An authorization model determines how these rights are granted. For example, in a discretionary model, the object owner can allocate rights. In a mandatory model, rights are predetermined by system-enforced rules and cannot be changed by any user within the system.

- **Accounting**—Tracking authorized usage of a resource or use of rights by a subject and alerting when unauthorized use is detected or attempted.



Differences among identification, authentication, authorization, and accounting. (Images © 123RF.com.)

For example, if you are setting up an e-commerce site and want to enroll users, you need to select the appropriate controls to perform each function:

- **Identification**—Ensure that customers are legitimate. For example, you might need to ensure that billing and delivery addresses match and that they are not trying to use fraudulent payment methods.
- **Authentication**—Ensure that customers have unique accounts and that only they can manage their orders and billing information.
- **Authorization**—Rules to ensure customers can place orders only when they have valid payment mechanisms in place. You might operate loyalty schemes or promotions that authorize certain customers to view unique offers or content.
- **Accounting**—The system must record the actions a customer takes (to ensure that they cannot deny placing an order, for instance).



Remember that these processes apply both to people and to systems. For example, you need to ensure that your e-commerce server can authenticate its identity when customers connect to it using a web browser.

Authentication Methods

An account defines a subject on the computer or network system. Assuming that an account has been created securely (the identity of the account holder has been verified), authentication verifies that only the account holder is able to use the account and that the system may be used only by account holders. Authentication is performed when the account holder submits credentials to the system to request access. These are compared to the credentials stored on the system. If they match, the account is authenticated.



Account creation should be managed by onboarding policies. These define procedures for checking identity, issuing credentials securely, and allocating appropriate permissions. Similarly, when an employee changes roles or leaves the company, an offboarding process should be applied to either adjust permissions or ensure that the account is disabled or deleted.

The type of data used to create a credential is called an authentication factor. Authentication factors fall into the following categories:

- **Knowledge factor**—Something you know (such as a password).
- **Ownership factor**—Something you have (such as a smart card).
- **Human or biometric factor**—Something you are (such as a fingerprint).
- **Behavioral factor**—Something you do (such as making a signature).
- **Location factor**—Somewhere you are, such as only being able to log into an account from a specific location, known as geofencing.
- **Time factor**—Somewhen you are (such as only being permitted to start a session during work hours or using an access token before it expires).

An authentication technology or mechanism is considered strong if it combines the use of more than one authentication data type (**multifactor**). Single-factor authentication systems can quite easily be compromised. For example, a password could be written down or shared, or compromised by a social engineering attack, a smart card could be lost or stolen, and a biometric system could be subject to high error rates. Behavioral, location, and time factors are not specific or reliable enough to be used as single factors, but can supplement other factors to make the authentication system stronger.

Two-factor authentication combines something such as a smart card or biometric mechanism with a knowledge factor, such as a password or personal identity number (PIN). Three-factor authentication combines three of the possible technologies. An example of this would be a smart card with an integrated fingerprint reader. This means that to authenticate, the user must possess the card, the user's fingerprint must match the template stored on the card, and the user must input a PIN.



Multifactor authentication requires a combination of different technologies. For example, requiring a PIN along with date of birth may be stronger than entering a PIN alone, but it is not multifactor.

Local Authentication

One of the most important features of an operating system is the authentication provider. The local authentication provider is the software architecture and code that underpins the mechanism by which the user is authenticated before starting a shell. This is usually described as a login (Linux) or a logon or sign-in (Microsoft). Knowledge-based authentication, using a password or PIN, is the default authentication provider for most operating systems.

Knowledge-based authentication relies on cryptographic hashes. A cryptographic hash is a function that converts any string to a unique, fixed-length code. The function should ensure that the code cannot be converted back into the plaintext string.

```
[s]tatus [p]ause [b]ypass [c]heckpoint [q]uit => s

Session.....: hashcat
Status.....: Running
Hash.Type.....: NetNTLMv2
Hash.Target.....: ADMINISTRATOR::515support:2f8cbd19fd1bfac9:881c5503...000000
Time.Started.....: Mon Jan  6 11:25:16 2020 (1 min, 38 secs)
Time.Estimated...: Sat Jan 11 07:49:57 2020 (4 days, 20 hours)
Guess.Mask.....: ?1?1?1?1?1?1?1 [8]
Guess.Charset....: -1 pPaAsSwWoOrRdD0123456789$, -2 Undefined, -3 Undefined, -4
Undefined
Guess.Queue.....: 1/1 (100.00%)
Speed.#1.....: 364.1 kH/s (11.09ms) @ Accel:128 Loops:32 Thr:1 Vec:8
Recovered.....: 0/1 (0.00%) Digests, 0/1 (0.00%) Salts
Progress.....: 34233472/152587890625 (0.02%)
Rejected.....: 0/34233472 (0.00%)
Restore.Point....: 2176/9765625 (0.02%)
Restore.Sub.#1...: Salt:0 Amplifier:1824-1856 Iteration:0-32
Candidates.#1....: $87r8678 -> dSDoRS12
```

Password credentials are stored as cryptographic hashes (such as the Hash. Target value shown in the screenshot) that cannot normally be converted back to plaintext strings. The hashcat utility attempts to recover passwords by matching hashes through dictionary or brute force methods.

A password is not usually transmitted or stored in a credential database as a plaintext because of the risk of compromise. Instead, the password is stored as a cryptographic hash. When a user enters a password to log in, an authenticator converts what is typed into a hash and transmits that to an authority. The authority compares the submitted hash to the one in the database and authenticates the subject only if they match.

Windows Authentication

Windows authentication involves a complex architecture of components (docs.microsoft.com/en-us/windows-server/security/windows-authentication/credentials-processes-in-windows-authentication), but the following three scenarios are typical:

- **Windows local sign-in**—The Local Security Authority (LSA) compares the submitted credential to a hash stored in the Security Accounts Manager (SAM) database, which is part of the registry. This is also referred to as interactive logon.
- **Windows network sign-in**—The LSA can pass the credentials for authentication to a network service. The preferred system for network authentication is based on Kerberos, but legacy network applications might use NT LAN Manager (NTLM) authentication.

- **Remote sign-in**—If the user's device is not connected to the local network, authentication can take place over some type of virtual private network (VPN) or web portal.

Linux Authentication

In Linux, local user account names are stored in `/etc/passwd`. When a user logs in to a local interactive shell, the password is checked against a hash stored in `/etc/shadow`. Interactive login over a network is typically accomplished using Secure Shell (SSH). With SSH, the user can be authenticated using cryptographic keys instead of a password.

A pluggable authentication module (PAM) is a package for enabling different authentication providers, such as smart card login (tecmin.com/configure-pam-in-centos-ubuntu-linux). The PAM framework can also be used to implement authentication to network servers.


Single Sign-On and Kerberos

A **single sign-on (SSO)** system allows the user to authenticate once to a local device and be authorized to access compatible application servers without having to enter credentials again. For example, a user could log in to a Windows computer using a Microsoft account and be able to access OneDrive, Teams, Office 365 in Outlook, and other linked Microsoft and non-Microsoft web services, without having to sign in again.

One means of implementing SSO is the Kerberos framework. **Kerberos** provides SSO authentication to Active Directory, as well as compatibility with other, non-Windows operating systems. Kerberos was named after the three-headed guard dog of Hades (Cerberus) because it consists of three parts: Client (which requests services), Server (from which the service is requested) and a Key Distribution Center (KDC)—to vouch for their identity.


There are two services that make up a KDC: the Authentication Service and the Ticket Granting Service.

The Authentication Service is responsible for authenticating user logon requests. More generally, users and services can be authenticated; these are collectively referred to as principals. For example, when you sit at a Windows domain workstation and log on to the domain (Kerberos documentation refers to realms rather than domains, which is Microsoft's terminology), the first step of logon is to authenticate with a KDC server (implemented as a domain controller).



Kerberos Authentication Service. (Images © 123RF.com.)

When authenticated, the KDC server presents the user with a Ticket Granting Ticket. To access resources within the domain, the client requests a service ticket (a token that grants access to a target application server) by supplying the Ticket Granting Ticket to the Ticket Granting Service (TGS).



Kerberos Ticket Granting Service. (Images © 123RF.com.)

Digital Certificates and PKI

A protocol such as Kerberos can be used with smart cards. A smart card is programmed with an encryption key pair and a digital certificate, issued by the authenticating domain. Digital certificates are also used to authenticate machines when using Transport Layer Security (TLS). A certificate can be installed on a web server or email server to validate its identity and establish a secure transmission channel.

Digital certificates depend on the concept of public key cryptography. Public key cryptography, also referred to as asymmetric encryption, solves the problem of distributing encryption keys when you want to communicate securely with others, authenticate a message that you send to others, or authenticate yourself to an access control system. With asymmetric encryption, you generate a key pair. The private key in the pair remains a secret that only you know. The public key can be transmitted to other subjects. The private key cannot be derived from the public key. The key pair can be used in the following ways:


- When you want others to send you confidential messages, you give them your public key to use to encrypt the message. The message can then only be decrypted by your private key, which you keep known only to yourself. Due to the way asymmetric encryption works, the public key cannot be used to decrypt a message, even though it was used to encrypt it in the first place.



As encryption using a public key is relatively slow; rather than encrypting the whole message using a public key, more typically, the public key is used to encrypt a symmetric encryption key for use in a single session and exchange it securely. The symmetric session key is then used to encrypt the actual message. In a symmetric cipher, the same key can perform both encryption and decryption.

- When you want to authenticate yourself to others, you create a signature using your private key. You give others your public key to use to verify the signature. As only you know the private key, everyone can be assured that only you could have created the signature.

The basic problem with public key cryptography lies in proving the identity of the owner of a public key. The system is vulnerable to attacks where a threat actor is able to substitute your public key for their own. **Public key infrastructure (PKI)** aims to prove that the owners of public keys are who they say they are. Under PKI, anyone distributing public keys should obtain a digital certificate. The validity of the certificate is guaranteed by a certificate authority (CA). A digital certificate is essentially a wrapper for a subject's (or end entity's) public key. As well as the public key, it contains information about the subject and the certificate's issuer or guarantor. The certificate is digitally signed to prove that it was issued to the subject by a particular CA.



Digital certificate details. (Screenshot used with permission from Microsoft.)

In some circumstances, using PKI can be too difficult or expensive to manage. Any machine, web server, or program code can be deployed with a **self-signed certificate**. For example, the web administrative interfaces of SOHO routers are often only protected by a self-signed certificate. Self-signed certificates can also be useful in development and test environments. The operating system or browser will mark self-signed certificates as untrusted, but a user can choose to override this. The nature of self-signed certificates makes them very difficult to validate. They should not be used to protect critical hosts and applications.

Key Management

Key management refers to operational considerations for the various stages in the lifecycle of an encryption key or key pair. A key's lifecycle may involve the following stages:

- **Key Generation**—Creates an asymmetric key pair or symmetric secret key of the required strength, using the chosen cipher.
- **Storage**—Prevents unauthorized access to a private or secret key and protects against loss or damage.
- **Revocation**—Prevents use of the key if it is compromised. If a key is revoked, any data that was encrypted using it should be re-encrypted using a new key.
- **Expiration and Renewal**—Gives the certificate that validates the key a "shelf-life" to increase security. Every certificate expires after a certain period. Certificates can be renewed with the same key pair or with a new key pair.

A decentralized key management model means that keys are generated and managed directly on the computer or user account that will use the certificate. This does not require any special setup and so is easy to deploy. It makes the detection of key compromise more difficult, however.

Some organizations prefer to centralize key generation and storage using a tool such as a key management system. In one type of cryptographic key management system, a dedicated server or appliance is used to generate and store keys. When a device or app needs to perform a cryptographic operation, it uses the Key Management Interoperability Protocol (KMIP) to communicate with the server.


Federated Identity and SAML

Federation is the notion that a network needs to be accessible to more than just a well-defined group of employees. In business, a company might need to make parts of its network open to partners, suppliers, and customers. The company can manage its employee accounts easily enough. Managing accounts for each supplier or customer internally may be more difficult. Federation means that the company trusts accounts created and managed by a different network. As another example, in the consumer world, a user might want to use both Google Workspace and Facebook. If Google and Facebook establish a federated network for the purpose of authentication and authorization, then the user can log on to Facebook using their Google credentials or vice versa.

An on-premises network can use technologies such as Active Directory and Kerberos, because the administration of accounts and devices can be centralized. When implementing federation, authentication and authorization design comes with more constraints and additional requirements to ensure interoperability between different platforms. Web applications might not support Kerberos, while third-party networks might not support direct federation with Active Directory. The design for these cloud networks likely requires the use of other standard protocols or frameworks for interoperability between web applications.

These interoperable federation protocols use claims-based identity. While the technical implementation and terminology is different, the overall model is similar to that of Kerberos SSO:

1. A service provider (SP) establishes a trust relationship with an identity provider (IdP).
2. The principal attempts to access a service provider.
3. The service provider redirects the principal to the IdP.
4. The principal authenticates with the identity provider.
5. If authentication is successful, the principal obtains a claim, in the form of some sort of token or document signed by the IdP.
6. The principal presents the claim to the service provider. The SP can validate that the IdP has signed the claim because of its trust relationship with the IdP.



Federated identity management overview. (Images © 123RF.com.)

The service provider can now connect the authenticated principal to its own accounts database to determine its permissions and other attributes. It may be able to query attributes of the user account profile held by the IdP if the principal has authorized this type of access.

A federated network or cloud needs specific protocols and technologies to implement user identity assertions and transmit claims between the principal, the service provider, and the identity provider. **Security Assertion Markup Language (SAML)** is one such solution. SAML assertions (claims) are written in eXtensible Markup Language (XML). Communications are established using HTTP/HTTPS and the Simple Object Access Protocol (SOAP). The secure tokens are signed using the XML signature specification. The use of a digital signature allows the relying party to trust the identity provider.



Most other federated identity frameworks, such as OAuth, use the service provider terminology. In SAML, the service provider is referred to as a relying party (RP).

An example of a SAML implementation is Amazon Web Services (AWS) which functions as a SAML relying party. This allows companies using AWS to develop cloud applications to manage their customers' user identities and provide them with permissions on AWS without having to create accounts for them on AWS directly.

Remote Authentication

Local authentication takes place when a user tries to start an interactive session directly on a computer. The authentication system might use Kerberos or another SSO provider to allow the local host and user to access a network, but the session is still started locally. Remote authentication means that a host runs a remote access server or terminal server that accepts login requests initiated via another host over a network. Remote authentication is typically used in two scenarios:


- Authenticating with a cloud provider or web host or joining a virtual private network (VPN). With a VPN, the remote user connects to a remote access server on the perimeter of the private network.
- Authenticating with a different host over a private network. Administrators commonly need to manage switches, routers, and servers. Rather than go to the device and start a local console session, they use Secure Shell (SSH) or Remote Desktop Protocol (RDP) to start a session over the network from their management workstation or laptop. The target device must be running an SSH server service or RDP terminal access server.

With remote authentication, storing credentials on the remote access server is a risk. If the access server is on the network edge, it is more vulnerable to attacks. Also, if there are multiple access servers, it is difficult to synchronize accounts, credentials, and SSO authorizations between them. To mitigate these issues, remote access usually uses an authentication, authorization, and accounting (AAA) architecture. AAA uses the following components:

- **Supplicant**—The device requesting access, such as a user's PC or laptop.
- **Network access server (NAS) or network access point (NAP)**—Edge network appliances, such as switches, access points, and VPN gateways. These are also referred to as AAA clients or authenticators.
- **AAA server**—The authentication server, positioned within the local network. This server either holds a database of accounts and credentials or has access to a directory server that can authenticate requests and issue SSO authorizations. There are two main types of AAA server: RADIUS and TACACS+.

Remote Authentication Dial-In User Service

Remote Authentication Dial-In User Service (RADIUS) is very widely used for client device access over VPNs. There are several RADIUS server and client products. RADIUS typically uses UDP ports 1812 and 1813. Each RADIUS client must be configured with the IP address of the RADIUS server plus the same shared secret.



Configuring an OPNsense security appliance as a RADIUS client. The OPNsense appliance is working as a virtual private network (VPN) access server. It uses the RADIUS server at 10.1.16.1 to authenticate VPN users. The client must be configured with the same shared secret as the server. (Screenshot used with permission from OPNsense.)

Terminal Access Controller Access Control System

Terminal Access Controller Access Control System (TACACS+) is a similar protocol to RADIUS but designed to be more flexible and reliable. TACACS+ was developed by Cisco but is also supported on many of the other third-party and open source RADIUS server implementations. Where RADIUS is often used for network access control over end user devices, TACACS+ is often used in authenticating administrative access to routers and switches. It uses TCP over port 49, and the reliable delivery offered by TCP makes it easier to detect when a server is down.

Also, authentication, authorization, and accounting functions are discrete. Many device management tasks require reauthentication (similar to having to reenter a password for sudo or UAC) and per-command authorizations and privileges for users, groups, and roles. TACACS+ supports this workflow better than RADIUS.



sudo ("superuser do") in Linux and User Account Control (UAC) in Windows allow a user to temporarily elevate permissions without having to fully sign out and in to use a different account.

Lesson 10.2

Authorization and Account Management



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 4.1 Explain the importance of basic security concepts.

When a user is given an account, the account is authorized to perform action on network objects, such as being able to read a file. These accounts, authorizations, and objects are typically centrally managed using a directory service. Being able to explain the importance of these access management concepts will help you to perform more effectively in a network support role.

As you study this section, answer the following questions:

- What rules should govern how users are authorized to perform actions on the network?
- What role does a directory server play in access management?


Authorization and Role-Based Access Control

A user account that has been authenticated can be allocated rights and permissions on networks, computers, and data. This process is referred to as authorization. An access control model describes the principles that govern how authorization works to give users rights on network systems and data.

Discretionary access control (DAC) is based on the primacy of the resource owner. In a DAC model, every resource has an owner. The owner creates a file or service although ownership can be assigned to another user. The owner has full control over the resource, and they can modify its access control list (ACL) to grant rights to others.

Role-based access control (RBAC) means that an organization defines its authorizations in terms of the tasks that an employee or service must be able to perform. Each set of permissions is a role. Each principal (user or service account) is allocated to one or more roles. Under this system, the right to modify the permissions assigned to each role is reserved to a system owner. Therefore, the system is nondiscretionary as each principal cannot modify the ACL of a resource, even though they can change the resource in other ways. Principals gain rights implicitly (through being assigned to a role) rather than explicitly (being assigned the right directly).

The concept of a security group account goes some way toward turning a discretionary system into a role-based one. Rather than assigning rights directly to user accounts, the system owner assigns user accounts to security group accounts. Principals gain rights by being made a member of a security group. A principal can be a member of multiple groups and can therefore receive rights and permissions from several sources.



Using security groups to assign privileges. (Images © 123RF.com.)

RBAC can be partially implemented by mapping security groups onto roles, but they are not identical schemes. Membership in security groups is largely discretionary (assigned by administrators rather than determined by the system). Also, ideally, a principal should only inherit the permissions of a role to complete a particular task rather than retain them permanently. Administrators should be prevented from escalating their own privileges by assigning roles to their own accounts arbitrarily or boosting a role's permissions.

Privileged Access Management

A privileged account is one that can make significant configuration changes to a host, such as installing software or disabling a firewall or other security system. Privileged accounts also have rights to log on to network appliances and application servers.

Privileged access management (PAM) refers to policies, procedures, and technical controls to prevent the malicious abuse of privileged accounts by internal threat actors and to mitigate risks from weak configuration control over authorizations. These controls identify and document privileged accounts, giving visibility into their use and managing the credentials used to access them.

Some other general principles of PAM include least privilege and separation of duties:

- **Least privilege** means that a user is granted sufficient rights to perform their job and no more. This mitigates risk if the account should be compromised and fall under the control of a threat actor. Authorization creep refers to a situation where a user acquires more and more rights, either directly or by being added to security groups and roles. Least privilege should be ensured by closely analyzing business workflows to assess what privileges are required and by performing regular account audits.

- **Separation of duties** is a means of establishing checks and balances against the possibility that critical systems or procedures can be compromised by insider threats. Duties and responsibilities should be divided among individuals to prevent ethical conflicts or abuse of powers.


Lightweight Directory Access Protocol

Directory services are the principal means of implementing privilege management and authorization on an enterprise network.

When an authenticated user logs on to a network, the server security service generates an access key for the user. This contains the username and group memberships of the authenticated user. Whenever the user attempts to access a resource, their access key is provided as identification. The server's security service matches username and group memberships from the access key with entries in the access list, and from this, it calculates the user's access privileges.

All this information is stored in a directory. A directory is like a database, where an object is like a record, and things that you know about the object (attributes) are like fields. For products from different vendors to be interoperable, most directories are based on the same standard. The main directory standard is the X.500 series of standards. As X.500 is complex, most directory services are implementations of the **Lightweight Directory Access Protocol (LDAP)**. LDAP is not a directory standard, but a protocol used to query and update an X.500-like directory. LDAP is widely supported in current directory products, most notably in Windows Active Directory. Insecure LDAP messaging uses TCP and UDP port 389 by default.

In an X.500, each object has a unique identifier called a distinguished name. A distinguished name is made up of attribute value pairs, separated by commas. The most specific attribute is listed first, and successive attributes become progressively broader. This most specific attribute is also referred to as the relative distinguished name, as it uniquely identifies the object within the context of successive (parent) attribute values.



*Browsing objects in an Active Directory LDAP schema.
(Screenshot used with permission from Microsoft.)*

The types of attributes, what information they contain, and the way object types are defined through attributes (some of which may be required and some optional) are described by the directory schema. For example, the distinguished name of a web server operated by Widget in London might be:

```
CN=WIDGETWEB, OU=Marketing, O=Widget, L=London,
ST=London, C=UK, DC=widget, DC=example
```

LDAP Secure

Like many TCP/IP protocols, LDAP provides no security, and all transmissions are in plaintext, making it vulnerable to sniffing and spoofing attacks. Also, a server that allows anonymous access by unauthenticated clients is vulnerable to overloading by denial of service attacks. Authentication, referred to as binding to the server, can be implemented in the following ways:

- **Simple bind**—The client must supply its distinguished name (DN) and password, but these are passed as plaintext.
- **Simple Authentication and Security Layer (SASL)**—This framework allows a client and server to negotiate authentication and encryption parameters to make a connection over TCP port 389 secure. The client and server negotiate the use of a supported authentication mechanism, such as Kerberos. The STARTTLS command can be used to require certificate-based encryption (sealing) and message integrity (signing).
- **LDAP Secure (LDAPS)**—The server is installed with a digital certificate, which it uses to set up a secure tunnel for the user credential exchange. LDAPS uses port 636.

Generally, two levels of access will need to be granted on the directory: read-only access (query) and read/write access (update). This is implemented using an access control policy, but the precise mechanism is vendor-specific and not specified by the LDAP standards documentation.

Unless it is hosting a public service, the LDAP directory server should also only be accessible from the private network. This means that LDAP ports (389 over TCP and UDP) should be blocked by a firewall from access over the public interface.



LDAPS is sometimes referred to as LDAP over Secure Sockets Layer (SSL). SSL is an older, deprecated version of TLS. LDAPS would now always be configured to use TLS, not SSL.

Lesson 10.3

Network Hardening



Exam Objectives Covered

4.3 Given a scenario, apply network security features, defense techniques, and solutions.

When devising authentication and authorization policies, it is sometimes too easy to focus on user accounts. Networks are also posed to risks from vulnerable and compromised computers and services. Policies and technologies that focus on protecting individual computing nodes are referred to as endpoint security. Endpoint security helps to establish multiple layers of access controls and protections, referred to as defense in depth. As a network technician, one of your roles will be ensuring that network hosts are put in a secure configuration and fully patched against vulnerabilities before being granted network access.

As you study this lesson, answer the following questions:

- What is the relationship between defense in depth and endpoint security?
- What are the critical elements of device and service hardening policies?

Defense in Depth

Firewalls try to establish a secure barrier at the network perimeter. This barrier is designed to subject any connections between the internal private network and external public or third-party networks to access controls. For example, a host on a public network would only be permitted to join the private network if it authenticates over a virtual private network (VPN). This system of focusing on the boundary between the public and private network and trusting everything that has connected via internal switches is called the perimeter security model.

The proliferation of mobile devices with wireless or cellular data access and cloud services, plus the better recognition of insider threat and vulnerabilities to malware, has eroded confidence in a solely perimeter-based security model. Network security design must address the concept of **defense in depth**. This refers to placing security controls throughout the network, so that all access attempts are authenticated, authorized, and audited.

Logical security controls governing access management, deception/honeypot strategies, and identity and access management (IAM) are all important parts of defense in depth. In addition to these, endpoint security is a set of procedures and technologies designed to restrict both remote and local network access at a device level and to ensure that each endpoint device is hardened to mitigate vulnerabilities.



The network perimeter is the border between the private network and external, public networks. This is often also referred to as the network edge. However, as the perimeter security model has eroded, the concept of where the network edge lies has expanded to include access switches and wireless access points. These would previously have been considered as “internal” under a perimeter-based security model.

Device and Service Hardening

As part of a defense in depth strategy, you need to think about making each host and network infrastructure device secure against tampering or abuse. It can be tempting to think of network devices such as switches and routers as self-contained. In fact, these devices often run complex firmware and host numerous services to enable remote management and configuration. Deploying systems in a secure configuration is known as device **hardening**. Some of the policies that will make up a secure configuration involve the following:

- **Change default passwords/credentials**—Devices such as wireless access points, switches, and routers sometimes ship with a default management password such as password, admin, or the device vendor's name. These should be changed on installation.
- **Enforce password complexity/length requirements**—Passwords for network infrastructure must be highly resistant to guessing and cracking attacks:
 - **Length**—No passwords should be less than eight characters. However, as critical infrastructure, passwords for network appliances should be 14+ characters.
 - **Complexity**—Requiring multiple character classes (mixing letters, case, digits, and symbols) is deprecated by NIST's latest guidance, but it is still required by many organizations' password policies.
 - **Avoiding common passwords**—The number of successful attacks against web servers and company networks has led to huge databases of credentials being posted online. Analysis of these databases shows how many users—even administrative users—rely on trivially simple passwords, such as 123456 or password. These password database dumps give attackers a useful dictionary to work with when trying to crack credentials. Any password that could be matched to a dictionary term is completely insecure and must not be used.
- **Configure role-based access**—The default administrator, superuser, or root account has unrestricted access to the device. If the credentials for this account are shared, the risk of compromise is greatly magnified. Role-based access means that a limited set of permissions is configured for different administrative groups, such as separating permissions for configuring the system to those for configuring logging and auditing. This separation of duties reduces impacts from the compromise of any single account.
- **Disable unneeded network services**—Any services or protocols that are not used should be disabled. This reduces the attack surface of a network appliance or OS. Attack surface means the range of things that an attacker could possibly exploit in order to compromise the device. It is particularly important to disable unused administration interfaces.

- **Disable insecure protocols**—Sniffing attacks can be mitigated by encrypting the channel over which communications takes place. This means that even if the eavesdropper can listen to the message, they cannot understand it without obtaining the encryption key. It is important to understand which protocols are insecure in terms of using unencrypted channels. This is particularly important when using a channel to authenticate. Insecure protocols should be deprecated, and secure protocols should be used instead. For example, the original versions of SNMP are unencrypted. To implement secure SNMP, either configure SNMPv3, which supports encryption, or use an encapsulation protocol such as IPSec to encrypt SNMP traffic.

Lesson 10.4

Switch Security



Exam Objectives Covered

4.3 Given a scenario, apply network security features, defense techniques, and solutions.

Hardening techniques are applied to individual hosts. There also need to be security mechanisms configured on switches to ensure that only authorized and configuration-compliant computers are permitted to access the network. Configuring switch security and network access control mechanisms will be an important part of your role as a network technician.

As you study this lesson, answer the following questions:

- What are the different types of network access control, and what parameters can they evaluate before a device is permitted to communicate on the network?
- What are the differences between port security, 802.1X, and Extensible Authentication Protocol?
- What are the roles of AAA servers, such as RADIUS and TACACS+, in network access control?
- What switch port security and monitoring features protect against other kinds of network attacks?

Network Access Control and Port Security

Network access control (NAC) is a system for authenticating endpoints before they can fully connect to the network. This is principally designed to mitigate risks from rogue devices and services. A basic type of NAC can be implemented by configuring **port security** mechanisms.

Disable Unneeded Switch Ports

Access to the physical switch ports and switch hardware should be restricted to authorized staff, using a secure server/equipment room and/or lockable hardware cabinets. This still leaves wall ports in work areas exposed, however, and is still vulnerable to rogue administrators (or theft of the equipment room keys). To prevent the attachment of unauthorized devices, a switch port can be disabled using the management software or isolated to a VLAN with no route to the network (a sinkhole VLAN). On a Cisco switch, these configuration settings will generally be applied using some version of a `switchport` command or subcommand. As another option, the patch cable can be physically removed from the port. Completely disabling ports in this way can introduce a lot of administrative overhead and scope for error. Also, it doesn't provide complete protection, as an attacker could unplug a device from an enabled port and connect their own laptop. Consequently, more sophisticated methods of ensuring port security have been developed.

MAC Filtering

Configuring **MAC filtering** on a switch means defining which MAC addresses are permitted to connect to a particular port. This can be done by creating a static lock list of valid MAC addresses or by specifying a limit to the number of permitted addresses. For example, if port security is enabled with a maximum of two MAC addresses, the switch will record the first two MACs to connect to that port but then drop any traffic from machines with different network adapter IDs that try to connect. This dynamic method is often referred to as sticky MACs. Addresses are dropped from the table if they go unused for a specified amount of time.

If a host attempts to connect with a MAC address that violates policy, the switch port enters a violation state:

- Protect mode means the port drops frames from the invalid source address but keeps the interface open otherwise. Protect mode can only be used with sticky MACs.
- Restrict mode drops frames and logs and alerts violations but also keeps the interface open.
- Shutdown mode disables the port and sends alerts. The port must be manually re-enabled using the `no shutdown` command. This is the default mode.

```

VB1F1(config)#int f0/2
VB1F1(config-if)#switchport port-security
VB1F1(config-if)#switchport port-security maximum 2
VB1F1(config-if)#switchport port-security mac-address sticky
VB1F1(config-if)#switchport port-security violation restrict
...
VB1F1#show port-security interface f0/2
Port Security           : Enabled
Port Status             : Secure-up
Violation Mode          : Restrict
Aging Time              : 0 mins
Aging Type              : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses   : 2
Total MAC Addresses     : 2
Configured MAC Addresses : 0
Sticky MAC Addresses    : 2
Last Source Address:Vlan : 0001.963A.81C3:112
Security Violation Count : 5

```

Configuring port security with up to two learned MACs and violation policy set to restrict. Note that violations have been reported, indicating that multiple additional hosts have been connected to this port. (Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)


Extensible Authentication Protocol and IEEE 802.1X

The **Extensible Authentication Protocol (EAP)** provides a framework for deploying multiple types of authentication protocols and technologies when an endpoint device needs to be authenticated before it can join the network. This pre-authentication requirement arises in a number of scenarios:

- When the user is accessing a wireless network and needs to authenticate with the network directory server.
- When a device is connecting to a network via a switch and network policies require the user to be authenticated before the device is allowed to communicate.
- When the user is connecting to the network over a public network via a virtual private network (VPN).

EAP allows many different authentication methods, but some of them use a digital certificate on the server and/or client machines. These certificates allow the machines to establish a trust relationship and create a secure tunnel to transmit the user credential, or to perform smart card authentication without a user password.

Where EAP implements a particular authentication factor and mechanism, the **IEEE 802.1X Port-Based Network Access Control (NAC)** standard provides the means of using an EAP method when a device connects to an Ethernet switch port, wireless access point, or VPN gateway. 802.1X uses authentication, authorization, and accounting (AAA) architecture. If the AAA protocol is RADIUS, the switch is configured as a RADIUS client by specifying the IP address or host name of the RADIUS server and setting a shared secret. The RADIUS server is positioned in a secure zone within the private network. The RADIUS server stores (or can obtain) account details and can validate authentication credentials. The switch does not have to store any authentication credentials. The switch forwards authentication data between the RADIUS server and the supplicant device. The RADIUS server uses the shared secret to validate RADIUS clients.



RADIUS authentication with EAP overview. (Images © 123RF.com.)

Basic NAC solutions can authenticate a client on the basis of machine certificates and/or user passwords. More sophisticated solutions can enforce a health policy. A health policy means that the client must submit an attestation report. This secure report proves that the client is running an authorized OS and has up-to-date patches and security scanner configurations.

Port Guards

There are various switch port security features to mitigate attacks on network infrastructure:

- A malicious host may use a spoofed MAC address to try to perform ARP cache poisoning against other hosts on the network and perpetrate an on-path attack. A switch port security feature such as dynamic ARP inspection (DAI) prevents a host attached to an untrusted port from flooding the segment with gratuitous ARP replies. ARP inspection maintains a trusted database of IP:ARP mappings. It also ensures that ARP packets are validly constructed and use valid IP addresses.
- Configuring **DHCP snooping** causes the switch to inspect DHCP traffic arriving on access ports to ensure that a host is not trying to spoof its MAC address. It can also be used to prevent rogue DHCP servers from operating on the network. With DHCP snooping, only DHCP offers from ports configured as trusted are allowed.
- Neighbor Discovery (ND) Inspection and **Router Advertisement (RA) Guard** perform similar functions to DAI and DHCP snooping for IPv6 networks. Most hosts have IPv6 enabled by default and disabling it can often cause unexpected problems. Consequently, these switch protections should be enabled to mitigate spoofing and on-path attacks over IPv6.
- When configuring VLANs, ensure that the default VLAN uses a different ID than any other user accessible VLAN. This mitigates against double tagging attacks.
- Ensure that ports allowed to be used as trunks are predetermined in the switch configuration and that access ports are not allowed to auto-configure as trunk ports. This mitigates against VLAN hopping attacks.
- To mitigate risks from attacks on spanning tree and root bridge selection, make sure that attackers can't easily guess which bridge ID number is being used by the legitimate root bridge. Set up Bridge Protocol Data Units Guard, or BPDU Guard, to allow an interface to put itself into blocking state when it receives a BPDU packet meant to change the root bridge switch. Enable root guard on the ports not being used as trunk lines. This keeps ports in their assigned roles. If one of these ports receives a BPDU frame, an error is logged and that port is blocked, thwarting the attacker's attempt to change the root bridge.

Port Mirroring

If it is operating normally, a switch forwards unicast traffic only to the specific port connected to the intended destination interface. This prevents sniffing of unicast traffic by hosts attached to the same switch. There are circumstances in which capturing and analyzing network traffic is a legitimate activity, however, and **port mirroring** provides the facility to do this. Port mirroring copies all packets sent to one or more source ports to a mirror (or destination) port. On a Cisco switch, this is referred to as a switched port analyzer (SPAN).

```

cumulus@cumulus:mgmt:~$ net add port-mirror session 1 ingress span src-port swp5-8 dst-port swp4
cumulus@cumulus:mgmt:~$ net commit
--- /etc/cumulus/switchd.d/port-mirror.conf      2021-08-06 18:48:09.494000000 +0000
+++ /run/nclu/port_mirror/port-mirror.conf      2021-08-06 18:49:43.620000000 +0000
@@ -9,11 +9,16 @@
#
# [session_n]
# session-id = n
# mirror.session.n.direction = (ingress | egress)
# mirror.session.n.src = <swpx, bond>
# mirror.session.n.dest = (swpx | <src-ip> <dst-ip>)
# mirror.session.n.type = (span | erspan | none)
#
# Default is all sessions off
# mirror.session.all.type = none
-mirror.session.all.type = none
+[session_1]
+session-id = 1
+mirror.session.1.src = swp5-8
+mirror.session.1.direction = ingress
+mirror.session.1.dest = swp4
+mirror.session.1.type = span

net add/del commands since the last "net commit"
=====

```

User	Timestamp	Command
cumulus	2021-08-06 18:48:09.496506	net add port-mirror session 1 ingress span src-port swp5-8 dst-port swp4
cumulus	2021-08-06 18:49:43.621259	net add port-mirror session 1 ingress span src-port swp5-8 dst-port swp4

```

cumulus@cumulus:mgmt:~$ net show port-mirror session 1
session-id direction type src      dest
-----
1 ingress span swp5-8 swp4

```

*Configuring port mirroring on a Cumulus VX switch.
(Screenshot courtesy of Nvidia.)*

The mirror port would be used by management or monitoring software, such as a packet sniffer, network analyzer, or intrusion detection system (IDS) sensor. Either ingress or egress traffic, or both, can be captured. Optionally, in order to avoid overloading the monitoring system, packets may be filtered based on criteria such as protocol ID or TCP/UDP port number.



Port mirroring demands a lot of processing and can lead to the switch hardware becoming overloaded and consequently crashing. If possible, test any security solution that requires port mirroring under typical loads before deploying it on a production network.

Lesson 10.5

Network Security Rules



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 4.3 Given a scenario, apply network security features, defense techniques, and solutions.
- 5.3 Given a scenario, troubleshoot common issues with network services.

Appliances such as firewalls and proxy servers enforce security rules to ensure authorized use of the network. They perform a filtering function to analyze the properties of connection requests and then allow, deny, and/or log them as appropriate. These rules require careful configuration. If they are too loose, your network will be exposed to a wider range of threats. If they are too restrictive, services and workflows could be disrupted. It is important you know the types of rules that can be configured to apply defense techniques and solutions.

As you study this section, answer the following questions:

- What distinguishing features of firewalls and proxies make them candidates for different types of network filtering?
- What distinguishes security rule types, such as access control lists (ACLs), uniform resource locator (URL) filtering, and content filtering?
- What is the difference between a stateful and a stateless ACL?
- What issues and impacts can arise from misconfigured security rules?


Security Rules and ACL Configuration

Logical network segments can be established using either a physical or virtual LAN (VLAN), switching topology mapped to IP subnets. Traffic between these segments must be routed. A firewall can be deployed to filter traffic entering or leaving a segment.



We're focusing on firewalls that protect routed network traffic here. Do be aware that firewalls can also be deployed in different ways. These include layer 2 inline or "bump in the wire" appliances and host or personal firewalls. A firewall can also be deployed as a virtual appliance, typically in cloud environments.

The security rules for what traffic is permitted and what should be blocked are configured as access control lists (ACLs). Design of firewall ACLs is guided by the principle of least access. This is similar to the principle of least privilege; only allow the minimum amount of traffic required for the operation of network services that should be permitted on the network segment and no more. The rules in a firewall's ACL are processed from top to bottom. If traffic matches a rule that allows the packet, then it is allowed to pass. Conversely, if a match is made to a block rule, the traffic is dropped. Once the firewall matches traffic to a rule, it stops processing subsequent rules. Consequently, the most specific rules are placed at the top of the ACL. If traffic does not match any rule, a firewall can be configured to block the traffic by default. This is called an **implicit deny**. If the firewall is not configured for default implicit deny, an **explicit deny** all rule can be added manually to the end of the ACL.



Sample firewall ruleset configured on OPNsense. This ruleset allows any ICMP traffic, HTTP/HTTPS traffic being forwarded to a local server (172.16.0.201), and SMTP traffic sent to the firewall (it is operating an SMTP mail gateway). (Screenshot used with permission from OPNsense.)

Each rule can specify whether to block or allow traffic based on parameters, often referred to as a tuple. For example, in the screenshot, the firewall imposes a 5-tuple rule, with matches against Protocol, Source address, Source port, Destination address, and Destination port.



As well as allowing and blocking, rules can be configured to log matches. Log-only rules are often used as a means of testing a new rule.

Most firewalls apply stateful rules. For example, a stateful firewall can differentiate between new and established connections. If a rule allows a new connection to be made by hosts on the public network to a server on the private network, a stateful firewall could automatically allow server replies via the established connection, but not allow the server to create new connections. To apply stateful rules, the firewall must be capable of maintaining a state table of established connections. If the firewall is stateless, an explicit rule would have to be configured for replies.

Proxy Servers

A typical network firewall filters traffic that is routed through its interfaces. By contrast, a **proxy server** forwards requests and responses on behalf of its clients. A proxy implies more of a break in the communications flow. Rather than inspecting traffic as it passes through, the proxy deconstructs each packet, performs analysis, and then rebuilds the packet and forwards it on, providing it conforms to the rules. The proxy could also perform address translation to convert between private and public addressing schemes.




Note that you should consider proxying as a function, rather than a class of device. It is possible to configure proxy server software on general-purpose computer hardware, but equally most firewall appliances will be capable of working as a proxy. Similarly, a proxy can be configured as a virtual appliance.

Forward Proxies

A forwarding proxy server provides for protocol-specific outbound traffic. For example, you might deploy a web proxy that enables hosts on a private network to connect to websites and secure websites on the Internet. A proxy server must understand the application it is servicing. A web proxy must be able to parse and modify HTTP and HTTPS requests and replies (and potentially HTML too). Some proxy servers are application-specific; others are multipurpose. A multipurpose proxy is one configured with filters for multiple protocol types, such as HTTP, FTP, and SMTP.

The main benefit of a proxy server is that clients connect to a specified point within the perimeter network for web access. This provides for a degree of traffic management and security. In addition, most web proxy servers provide caching engines, whereby frequently requested webpages and image assets are retained on the proxy, negating the need to refetch those files for subsequent requests.

Proxy servers can generally be classed as nontransparent or transparent. A nontransparent proxy means that the client must be configured with the proxy server address and port number to use it. The port on which the proxy server accepts client connections is often configured as port 8080. A transparent (or “forced” or “intercepting”) proxy intercepts client traffic without the client having to be reconfigured.



Configuring transparent proxy settings for the proxy server running on the OPNsense security appliance. (Screenshot used with permission from OPNsense.)

Reverse Proxies

A reverse proxy server provides for protocol-specific inbound traffic. Security and performance factors might make it unwise to allow hosts on a public network to connect directly to application servers. Instead, you can deploy a reverse proxy and configure it to listen for client requests from a public network (the Internet) and create the appropriate request to the application server. The proxy is said to publish the application. Typical applications for reverse proxy servers include publishing a web server, publishing messaging or conferencing applications, and enabling POP/IMAP mail retrieval.


A reverse proxy might handle the encryption/decryption and authentication on behalf of the application servers, reducing the overhead on those servers. It can also perform caching to improve performance. Reverse proxies could also be configured to perform load balancing across an application server pool.

Content Filtering

An ACL-type security rule applies basic Network or Transport layer filtering. By contrast, **content filtering** is capable of applying Application layer filters based on HTTP data. It could also apply more general business rules, such as time-of-day restrictions. Most firewalls and proxies now support some level of content filtering.

Content filtering plays a pivotal role in safeguarding an organization's network by blocking users from accessing malicious or inappropriate websites. A content filtering firewall or proxy can restrict access based on various criteria:

- **Uniform Resource Locator (URL) filtering**—A URL contains protocol (http or https, for instance), domain name, server path, and optionally query parameter components. A URL filter can scan the URL embedded in an HTTP request and allow or block it. The filter might use simple domain filtering to block access to specific URLs that are known to host malicious content, are inappropriate, or violate the company's Internet usage policy. It could use regular expression pattern matching to filter by keywords or path and query parameters.
- **Content and reputation-based filtering**—This filtering leverages continually updated databases that score websites based on their observed behavior and history. Sites known for hosting malware, engaging in phishing attacks, or distributing spam, for instance, would have a poor reputation score and could be automatically blocked. These databases can also sort websites into categories, such as social networking, gambling, adult content, webmail, or file sharing. Allowing or denying each category provides a simpler way to enforce web usage policies than allowing or denying individual URLs.



*Web filter content categories using the IPFire open source firewall.
(Screenshot used with permission from IPFire.)*



Transport Layer Security (TLS) poses major issues for proxies and content filters. A proxy cannot inspect or modify application data in encrypted traffic, but it cannot decrypt the traffic without breaking the TLS handshake between the client and the website. To perform TLS inspection, the proxy has to generate an enterprise certificate for each domain. The client trusts this certificate as it is issued by an enterprise CA but still matches the domain it has requested. The proxy establishes its own TLS tunnel with the website, forwarding the client's requests (if they conform to policy).

Misconfigured Firewall and ACL Issues

A misconfigured firewall or access control list (ACL) can cause different kinds of network access problems.

Access Denied Issues

One type of firewall, ACL, or content filter misconfiguration causes blocked services, ports, or addresses that should be allowed through. This will cause an application or protocol to fail to function correctly. For example, the firewall might be blocking TCP or UDP ports that are supposed to be open, or it might be allowing the ports but denying access to an IP network or host address that is supposed to be able to connect. Also, consider that advanced firewalls can apply additional filtering criteria, such as evaluating process/service executable names/locations or authorizations based on user accounts or group memberships.

A deny type of error will usually be easy to identify, as users will report incidents connected with the failure of the data traffic. With such incidents, firewall configuration will always be a likely cause, so will be high on the list to investigate. Diagnosis can be confirmed by trying to establish the connection from both inside and outside the firewall. If it connects from outside the firewall but not from inside, this will confirm the firewall to be the cause of the issue.

Multiple Firewall Issues

Another potential issue is where there are both network-based and host-based firewall settings to navigate in the communication path. There could be a host firewall running on the client, on the server, or on both. To diagnose an issue with a host firewall, attempt the connection with the host firewall disabled. If the connection attempt succeeds, then the network firewall ACL is allowing the packets, but the host firewall is configured to block them. If the connection attempt fails, investigate the network firewall ACL first. You can also inspect the firewall's log files to discover what rules have been applied to block traffic at a particular time.

Security Violation Issues

The other possible outcome of a badly configured firewall is that packets may be allowed through that should be blocked. This is a more serious outcome, because the result is to open the system to security vulnerabilities. It is also not necessarily so easily detected, as it does not typically cause anything to stop functioning. As no incidents usually arise from this outcome (except in the case that a vulnerability is exploited), it is not a scenario that is subject to troubleshooting. Rather, it underlines the need for regular firewall audits and thorough change control processes to deal with firewall change requests.

Module 10

Summary

You should be able to explain identity and access management concepts and apply device hardening and network access control techniques and solutions.

Guidelines for Applying Network Access Control Solutions

Follow these guidelines to support the use of IAM, network access control, and device hardening techniques and technologies:

- Deploy a multifactor scheme to authenticate users and devices. Consider the use of PKI and certificates to support the ownership factor.
- Deploy a directory server as a database of network users and resources.
- Consider using an SSO authentication and authorization framework, such as Kerberos. Evaluate requirements for federated identity management and SSO, using an interoperable framework, such as SAML.
- When deploying network appliances and servers, change default device credentials and ensure that accounts are secured with strong passwords or MFA. Configure fine-grained permissions to support role-based access and enforcement of least privilege management practices.
- Use only secure channels for administration traffic or any other protocol where credentials need to be submitted.
- Configure services according to the device's baseline and disable any services which are not required.
- Ensure that only the necessary IP ports (TCP and UDP ports) to run permitted services are open and that access to a port is controlled by a firewall ACL if appropriate.
- Use wall port and switch port security techniques to prevent attachment of unauthorized devices.
- To implement remote authentication or NAC with 802.1X/EAP/AAA, deploy a RADIUS or TACACS+ server to the internal network with a static IP address, shared secret, and accepted authentication methods. Configure network access devices (AAA clients) with the IP of the server and same shared secret.
- Configure security rules on firewalls and proxies to protect network segments and filter communications with public networks. Consider appropriate choices for security rules, from ACLs, URL filtering, and content filtering.

Module 11

Supporting Network Security Design

Module Introduction

The idea of an office network with web, file, and messaging services is a familiar one, but it is not the only use of networking. Networks are also used to support industrial and fulfillment processes, for example. The types and variety of services run over a network have a direct impact on its design, and on the security design especially. To support these diverse networks, you must be able to explain segmentation and security zone concepts. You should also be able to describe the technologies used to ensure the physical security of a network site.

Module Objectives

In this module, you will do the following:

- Explain the importance of network segmentation and use of trusted and untrusted zones.
- Describe security implications of internet of things (IoT) and industrial internet of things (IIoT).
- Explain the importance of physical security.

Lesson 11.1

Zone-Based Security



Exam Objectives Covered

4.3 Given a scenario, apply network security features, defense techniques, and solutions.

A zone-based security model groups hosts with the same level of trust into separate network segments. Traffic between segments is subject to policies and rules that are enforced by security appliances such as firewalls, proxy servers, and intrusion detection/prevention systems. These appliances perform a filtering function to analyze the properties of connection requests and then allow, deny, and/or log them as appropriate. Understanding the concept of zone-based security is vital for the application of effective defense techniques and solutions, especially when it comes to configuring perimeter networks.

As you study this lesson, answer the following questions:

- What technologies are used to implement zones and network segmentation enforcement?
- What determines whether a zone is fully trusted, partially trusted, or untrusted?
- What are the functions of the two firewalls in a screened subnet?
- Which types of hosts are placed within a screened subnet?

Network Security Zones

Effective placement of security appliances depends on segmenting the network into clearly defined areas. At layers 2 and 3, **network segmentation enforcement** is applied using a combination of virtual LANs and subnets. Each segment is a separate broadcast domain. Any traffic between segments must be routed.

In security terms, the main unit of a logically segmented network is a zone. A zone is an area of the network where the security configuration is the same for all hosts within it. Put another way, a zone is an area where all hosts have the same level of trust. Network traffic between zones should be restricted by policies and rules. These rules are enforced by a security device—typically a firewall or proxy.

Trust largely depends on the extent to which devices, user accounts, services, and traffic in the zone are managed and monitored. A zone with hosts that are highly trusted will have a minimal attack surface, because permitted traffic is strictly defined and extensive security controls are deployed to minimize threats and vulnerabilities. Conversely, hosts with low trust might expose a large attack surface, because they need to establish diverse connections with other zones with different trust levels.

For example, an organization could identify the following general security zones to use as a basis for writing security policies and rules:

- **Private server administrative networks**—Devices are subject to strict hardening and configuration management policies. Hosts, user accounts, and traffic with permission to operate in the zone are continually monitored to ensure compliance with security policies.
- **Private client network**—Devices are subject to security policies and monitoring, but the diverse range of technologies and permissions to use public networks make the zone less than fully trusted.
- **Guest**—Unmanaged devices are allowed to connect, subject to some restrictions and monitoring. This zone is typically untrusted and would not be allowed access to trusted networks.
- **Public server network**—Devices are fully managed but accept connections from unmanaged public clients. Consequently, hosts within this zone are only partially trusted.
- **Public**—The zone is unmanaged and therefore untrusted.

Zones with different trust levels and security rules would typically be configured to protect the integrity and confidentiality of different asset groups within the organization. For example, servers storing financial records can be their own VLAN, and marketing servers could be another VLAN. If something like a remote access Trojan were introduced in one VLAN, it should not be able to spread to other VLANs without also being able to pass through the firewall protecting each zone.

Perimeter Networks

One important distinction between different security zones is whether a host is Internet facing. An Internet-facing host accepts or initiates connections from or to hosts on the public Internet. Internet-facing hosts are placed in a perimeter network zone. The basic principle of a perimeter network zone is that traffic cannot pass through it directly. A perimeter network enables external clients to access data on private systems, such as web servers, without compromising the security of the internal network.


If communication is required between hosts on either side of the perimeter network, a host within it can be configured to act as a proxy. For example, if a host on the local network requests a connection with a web server on the Internet, a proxy in the network perimeter takes the request and checks it. If the request is valid, it retransmits it to the destination. To the external host, all communications seem to be initiated by the proxy. The external host has no direct connectivity with the LAN client device.

Servers that provide public access services should be placed in a perimeter network. These would typically include web servers, mail and other communications servers, proxy servers, and remote access servers. The hosts in the perimeter are not fully trusted by the internal network because of the possibility that they could be compromised from the Internet.

Screened Subnets


To configure a perimeter network, two different security configurations must be enabled: one on the external interface and one on the internal interface. There are several ways of implementing this as a topology of physical or virtual security appliances.

A **screened subnet** uses two firewalls placed on either side of the perimeter network zone. The screening firewall restricts traffic on the external/public interface and allows permitted traffic to the hosts in the perimeter zone subnet. The internal firewall filters communications between hosts in the perimeter and hosts on the LAN. This firewall is often described as the choke firewall. A choke point is a purposefully narrow gateway that facilitates better access control and easier monitoring.



Screened subnet using multiple firewalls. The border firewall has a WAN interface but no direct connection to the LAN. Instead, it routes filtered traffic to the choke firewall. The choke firewall has a LAN interface, plus two interfaces in screened subnets implementing a guest network and a public-facing app server network. (Images © 123RF.com.)

A perimeter network can also be established using one router/firewall appliance with three (or more) network interfaces, referred to as triple homed. One interface is the public one, another is the perimeter subnet, and the third connects to the LAN. Routing and filtering rules determine what forwarding is allowed between these interfaces. This can achieve the same sort of configuration as a screened subnet.



Screened subnet using a single firewall. This is directly connected to the LAN.
(Images © 123RF.com.)




Various types of Internet-facing zones or hosts are also popularly referred to as a demilitarized zone (DMZ). This is vague marketing terminology that does not accurately describe the purpose or configuration of a perimeter network. Hosts in a perimeter network remain fully managed by a private organization. Filtered public access is permitted, but there is no "demilitarization" in the sense of making the zone in any way neutral or jointly operated.

Intrusion Detection and Prevention Systems

An **intrusion detection system (IDS)** performs real-time analysis of either network traffic or system and application logs. Where a firewall or proxy applies security rules, an IDS is configured with signature patterns. Each pattern represents a known type of malicious activity. If a pattern is matched in a traffic stream, the IDS raises an alert. Like antivirus software, the IDS must be kept up to date with the latest signature patterns.

An IDS might also be capable of anomaly-based detection. Anomaly-based detection first defines a baseline of normal network traffic and then monitors it. It then looks for anything that falls outside that baseline. The main drawback is that anomaly-based detection generates high levels of false positives, where legitimate traffic is flagged as malicious. An IDS is often also configured with automated threat data, such as lists of IP addresses and domains that are associated with threat actors.

Like a packet analyzer, an IDS must be configured with a sniffer to read frames from a mirrored port or TAP. Placement of the sniffer must be carefully considered to meet security goals. Typically, an IDS is positioned behind a firewall to monitor traffic entering and exiting a security zone. The aim is to detect suspicious traffic that the firewall has not blocked, providing defense in depth. This type of passive sensor does not slow down traffic and is undetectable by the attacker (it does not have an IP address on the monitored network segment).



```

25 #####
26 # This file contains a sample snort configuration.
27 # You should take the following steps to create your own custom configuration:
28 #
29 # 1) Set the network variables.
30 # 2) Configure the decoder
31 # 3) Configure the base detection engine
32 # 4) Configure dynamic loaded libraries
33 # 5) Configure preprocessors
34 # 6) Configure output plugins
35 # 7) Customize your rule set
36 # 8) Customize preprocessor and decoder rule set
37 # 9) Customize shared object rule set
38 #####
39
40 #####
41 # Step #1: Set the network variables. For more information, see README.variables
42 #####
43
44 # Setup the network addresses you are protecting
45 ipvar HOME_NET 10.1.0.0/24
46
47 # Set up the external network addresses. Leave as "any" in most situations
48 ipvar EXTERNAL_NET !$HOME_NET
49
50 # List of DNS servers on your network
51 ipvar DNS_SERVERS $HOME_NET
52
53 # List of SMTP servers on your network
54 ipvar SMTP_SERVERS $HOME_NET
55
56 # List of web servers on your network
57 ipvar HTTP_SERVERS $HOME_NET
58
59 # List of sql servers on your network
60 ipvar SQL_SERVERS $HOME_NET
61

```

Configuration file for the Snort IDS.

Compared to the passive logging/alerting functionality of an IDS, an **intrusion prevention system (IPS)** can provide an active response to any network threats that it matches. One typical preventive measure is to end the session by sending a TCP reset packet to the attacking host. Another option is for the sensor to apply a temporary filter on the firewall to block the attacker's IP address (shunning). Other advanced measures include throttling bandwidth to attacking hosts, applying complex firewall filters, and even modifying suspect packets to render them harmless. Finally, the appliance may be able to run a script or third-party program to perform some other action not supported by the IPS software itself.

IPS functionality is now very commonly built into firewall appliances and proxy servers. An IPS-enabled firewall is inline with the network, meaning that all traffic passes through it (also making them a single point of failure if there is no fault tolerance mechanism). This obviously means that they need to be able to cope with high bandwidths and process each packet very quickly to avoid slowing down the network.



Network IDS/IPS can be combined with host-based IDS/IPS. These run as agents on end systems to monitor application processes, data files, and log files for suspicious activity. Advanced IDS/IPS suites analyze information from multiple sensors to identify suspicious traffic flows and host activity.

Lesson 11.2

Internet of Things



Exam Objectives Covered

4.1 Explain the importance of basic network security concepts.

Many people and businesses are deploying internet of things (IoT) devices in their homes and offices, and some businesses depend on the underlying embedded systems technology for manufacturing and fulfillment. In this lesson, you will examine how these technologies can be integrated securely with or alongside corporate data networks.

As you study this lesson, answer the following questions:

- What are embedded and IoT systems and devices?
- What differences are there between consumer and industrial IoT devices and networks?
- What is the difference between ICS and SCADA?
- What impact does the use of IoT or IIoT have on network segmentation enforcement?

IoT Devices

An embedded system is a complete computer system that is designed to perform a specific, dedicated function. These systems can be as contained as a microcontroller in an intravenous drip-rate meter or as large and complex as the network of control devices managing a water treatment plant. Embedded systems can be characterized as static environments. A PC is a dynamic environment. The user can add or remove programs and data files, install new hardware components, and upgrade the operating system. A static environment does not allow or require such frequent changes.

In terms of security this can be ideal, because unchanging environments are typically easier to protect and defend. Static computing environments pose their own risks, however. A static environment is often an unknown environment to security administrators. Unlike an OS environment such as Windows, there may be little support for identifying and correcting security issues.

The term **internet of things (IoT)** is used to describe a global network of embedded systems used as or in personal devices, home appliances, home control systems, vehicles, and other items that have been equipped with sensors, software, and network connectivity. These features allow these types of objects to communicate and pass data between themselves and other traditional systems like computer servers. This is often referred to as machine to machine (M2M) communication.

Consumer-Grade Smart Devices

Smart devices are used to implement home automation systems. An IoT smart device network will generally use the following types of components:

- **Hub/control system**—IoT devices usually require a communications hub to facilitate wireless networking. There must also be a control system, as most IoT devices are headless, meaning they have no terminal interface. A headless hub could be implemented as a smart speaker operated by voice control or use a smartphone/PC app for configuration.
- **Smart devices**—IoT endpoints implement the function, such as a smart light bulb, refrigerator, thermostat/heating control, or doorbell/video entry phone. These devices are capable of compute, storage, and network functions that are all potentially vulnerable to exploits. Most smart devices use a Linux or Android kernel. Because they're effectively running mini-computers, smart devices are vulnerable to some of the standard attacks associated with web applications and network functions. Integrated peripherals such as cameras or microphones could be compromised to facilitate surveillance.

Physical Access Control Systems and Smart Buildings

A physical access control system (PACS) is a network of monitored locks, intruder alarms, and video surveillance cameras. A building automation system (BAS) or smart building for offices and datacenters can include PACS, but also network-based configuration and monitoring of heating, ventilation, and air conditioning (HVAC); fire control; power and lighting; and elevators and escalators. These subsystems are implemented by programmable logic controllers (PLCs) and various types of sensors that measure temperature, air pressure, humidity, room occupancy, and so on.

Industrial Embedded Systems

IoT devices and other embedded systems are used within many sectors of industry, including energy generation and distribution, mining and refining raw materials, fabrication and manufacturing, and logistics (moving and delivering components and goods).

Industrial internet of things (IIoT) systems have different priorities than IT systems. Often, hazardous electromechanical components are involved, so safety is the overriding priority. Industrial processes also prioritize availability and integrity over confidentiality—reversing the CIA triad as the AIC triad.

Workflow and Process Automation Systems

An **industrial control system (ICS)** provides mechanisms for workflow and process automation. An ICS controls machinery used in critical infrastructure, such as power suppliers, water suppliers, health services, telecommunications, and national security services. An ICS that manages process automation within a single site is usually referred to as a distributed control system (DCS).

An ICS comprises plant devices and equipment with embedded programmable logic controllers (PLCs). The PLCs are linked by a cabled network to actuators that operate valves, motors, circuit breakers, and other mechanical components, plus sensors that monitor some local state, such as temperature. Output and configuration of a PLC is performed by one or more human-machine interfaces (HMIs). An HMI might be a local control panel or software running on a computing host. PLCs are connected within a control loop, and the whole process automation system can be governed by a control server. Another important concept is the data historian, which is a database of all the information generated by the control loop.

Supervisory Control and Data Acquisition

A **supervisory control and data acquisition (SCADA)** system takes the place of a control server in large-scale, multiple-site ICSs. SCADA typically run as software on ordinary computers, gathering data from and managing plant devices and equipment with embedded PLCs, referred to as field devices. SCADA typically use WAN communications, such as cellular data networks, to link the SCADA server to field devices.

IoT Networks

Each device in an IoT network is identified with some form of unique serial number or code embedded within its own operating or control system and can interoperate within Internet infrastructure, either directly or via an intermediary. As these devices tend to be small and often either unpowered or dependent on battery power, the standard Ethernet, cellular, and Wi-Fi networking products that connect computers are not always suitable for use. Other networking standards and products have been developed to facilitate IoT networks.

Operational Technology Networks

A cabled network for industrial systems is referred to as an **operational technology (OT)** network. This term is purposefully used to distinguish the requirements of an industrial network from those of an “ordinary” IT data network. An OT network typically uses either serial data protocols or industrial Ethernet. Industrial Ethernet is optimized for real-time, deterministic transfers. OT networks could also use vendor-developed data link and networking protocols, as well as specialist application protocols.

Cellular Networks

A cellular network for IoT enables long-distance communication over the same system that supports mobile phones and smartphones. This is also called baseband radio, after the baseband processor that performs the function of a cellular modem. There are several baseband radio technologies:

- **Narrowband-IoT (NB-IoT)**—This refers to a low-power version of the Long Term Evolution (LTE) or 4G cellular standard. The signal occupies less bandwidth than regular cellular. This means that data rates are limited (20–100 kbps), but most sensors need to send small packets with low latency, rather than making large data transfers. Narrowband also has greater penetrating power, making it more suitable for use in inaccessible locations, such as tunnels or deep within buildings, where ordinary cellular connectivity would be impossible.
- **LTE Machine Type Communication (LTE-M)**—This is another low-power system but supports higher bandwidth (up to about 1 Mbps).

IoT Network Security

Internet of things (IoT) devices might be present in dedicated embedded systems and smart building networks and/or as individual smart devices installed in employee workplaces. Placement and segmentation issues for embedded and IoT systems are best considered by dividing them into three principal groups: consumer-grade devices, smart building technology, and industrial systems.

Consumer-Grade Smart Devices

IoT devices are likely to use a Wi-Fi network for connectivity and administration. Consumer-grade smart devices and home automation products can be poorly documented, and patch management/security response processes of vendors can be inadequate. When they are designed for residential use, IoT devices can suffer from weak defaults. They may be configured to “work” with a minimum of configuration effort. There may be recommended steps to secure the device that the customer never takes.

In a corporate workspace, the main risk from smart device placement is that of shadow IT, where employees deploy a network-enabled device without going through a change and configuration management process. A vulnerability in the device would put it at risk of being exploited as an access point to the network. These devices also pose a risk for remote working, where the employee joins the corporate VPN using a home wireless network that is likely to contain numerous undocumented vulnerabilities and configuration weaknesses.

These risks can be mitigated by regular audits and through employee security awareness training. Ensure that administrative interfaces are secured and that device configuration and management is assigned to appropriate organizational roles. Include all IoT devices in patch and vulnerability management audits.

Smart Buildings

By contrast with consumer-grade components, there should be less scope for compromise in the entry mechanisms and climate/lighting control components of a properly designed smart building system. Management and monitoring of the system should be performed over isolated network segments. Configuration management and change control processes should ensure that no weak configurations are introduced and that vendor advisories are tracked for any known vulnerabilities or exploits so that these can be patched or mitigated.

ICS/SCADA

While an ICS or SCADA is typically implemented as a dedicated OT or wireless WAN network, there may be points where these networks are linked to a corporate data network. Historically, these vulnerable links and bridging hosts have been exploited by threat actors. There are risks both to embedded systems from the data network and to corporate data assets and systems from the embedded network. Where possible, isolate management and monitoring traffic for embedded systems to minimize access to and from the corporate data network. If OT and IT networks cannot be completely isolated, links to them must be closely monitored and subject to access controls.

Lesson 11.3

Physical Security



Exam Objectives Covered

4.1 Explain the importance of basic network security concepts.

For a network to be secure, access to the premises must be controlled by physical security measures. Additionally, there might be access controls for different zones within a site. An understanding of procedures and hardware that improve the physical security of site premises will help reduce the risk of intrusion.

As you study this lesson, answer the following questions:

- What are the roles of prevention and detection in physical security?
- What types of equipment are available to aid in physical security?


Locks

Prevention-type physical controls are ones that stop an intruder from gaining unauthorized access, if they work effectively. Where an area is controlled by being enclosed by walls or fencing, access is channeled through defined points of entry, such as doors and gates. These entry points can be protected by types of electronic lock.

Badge Reader and Biometric Locks

Various types of access control hardware or electronic locks can be deployed to enable users to authenticate quickly at access points:

- **Badge reader**—A photographic ID badge showing name and access level is one of the cornerstones of building security. A smart badge comes with an integrated chip and data interface that stores the user's key pair and digital certificate. The user presents the card and enters a PIN, and then the card uses its cryptographic keys to authenticate securely via the entry point's badge reader. A smart badge is either contact based, meaning that it must be physically inserted into a reader, or contactless, meaning that data is transferred using a tiny antenna embedded in the card. The ISO has published various ID card standards to promote interoperability, including ones for smart cards (ISO 7816 for contact and ISO 14443 for contactless types).
- **Biometric**—An electronic lock may also be integrated with a biometric scanner. A biometric device is activated by human physical features, such as a fingerprint, voice, retina, or signature. Each user's biometric is recorded as a template and stored on an authentication server. To gain access, the user's biometric is scanned again by a fingerprint reader or iris/retina scanner and compared to the template scan.



*Two types of electronic lock with biometric reader (left) and badge/card reader (right).
(Images © 123RF.com.)*

Rack System Locks

The access control hardware measures that can be deployed to prevent unauthorized entry to sites, buildings, and floors or zones within a building can also be used to manage access to IT assets.

Installing equipment within secure cabinets or enclosures provides mitigation against insider attack and attacks that have broken through the perimeter security mechanisms. These can be supplied with key-operated or electronic locks. It is also possible to provision lockable brackets and drawers to protect or isolate individual elements within a rack.

Some datacenters may contain racks with equipment owned by different companies (colocation). These racks can be installed inside cages so that technicians can only physically access the racks housing their own company's servers and appliances.



*Colocation cages. (Image © Chris Dag and shared with CC BY
2.0 flickr.com/photos/chrisdag/865711871.)*

Cameras

Detection-based controls provide an important additional layer of defense in the event that prevention-based controls fail to work. For example, surveillance is another layer of security designed to improve the resilience of perimeter gateways. Effective surveillance mechanisms ensure that attempts to penetrate a barricade are detected. Surveillance may be focused on perimeter areas or within security zones themselves. Surveillance can be performed by security guards or via video. Camera-based surveillance is a cheaper means of monitoring than maintaining separate guards at each gateway or zone.

A security camera is either fixed or can be operated using pan-tilt-zoom (PTZ) controls. Different cameras suit different purposes. If you want to record the image of every person entering through an access control vestibule, a fixed, narrow focal length camera positioned on the doorway will be perfectly adequate. If you want to survey a large room and pick out individual faces, a camera with PTZ is required.



*Pan-tilt-zoom CCTV installed to monitor a server room.
(Image by Dario Lo Presti © 123RF.com.)*

The cameras in a closed-circuit television (CCTV) network are typically connected to a multiplexer using coaxial cabling. The multiplexer can then display images from the cameras on one or more screens, allow the operator to control camera functions, and record the images to tape or hard drive. Newer camera systems may be linked in an IP network, using regular data cabling. Small IP cameras can use Power over Ethernet (PoE), avoiding the need to provision a separate power circuit.


Geofencing

Geolocation is the use of network attributes to identify (or estimate) the physical position of a device. The device uses location services to determine its current position. Location services can make use of two systems:

- **Global Positioning System (GPS)**—This is a means of determining the device's latitude and longitude based on information received from satellites via a GPS sensor.
- **Indoor Positioning System (IPS)**—This works out a device's location by triangulating its proximity to other radio sources, such as cell towers, Wi-Fi access points, and Bluetooth/RFID beacons.

Geofencing is the practice of creating a virtual boundary based on real-world geography. Geofencing can be a useful tool with respect to controlling the use of camera or video functions or applying context-aware authentication.

An organization may use geofencing to create a perimeter around its office property, and subsequently, limit the functionality of any devices that exceed this boundary. An unlocked smartphone could be locked and forced to reauthenticate when entering the premises, and the camera and microphone could be disabled. The device's position is obtained from location services, while policies and enforcement are facilitated by Mobile Device Management (MDM) software.



Restricting device permissions such as camera and screen capture using Intune. (Screenshot used with permission from Microsoft.)



Geofencing uses physical attributes reported by a device's location services, but it is considered a logical security control as the permissions are managed by a software suite.

Module 11

Summary

You should be able to explain the importance of network segmentation and physical security, and describe the security implications of IoT and IIoT.

Guidelines for Supporting Network Security Design

Follow these guidelines to support network segmentation and physical security:

- Identify requirements for different types of security appliances, based on the following factors:
 - Using firewalls or proxy servers to establish perimeter security in a screened subnet topology.
 - Using firewalls to protect internal zones or individual hosts.
- Deploy intrusion detection/prevention behind firewalls to identify threats that have passed filtering.
- Develop a strategy for managing IoT/IIoT devices in segmented network zones.
- Assess requirements for site access controls and monitoring.

Module 12

Configuring Wireless Networks

Module Introduction

Wireless media technologies have distinct advantages for businesses over cabling. They can be easier to install to existing premises, and they support the device mobility that users require from laptop or smartphone-based access to networks. Wireless technology implementations offer various advantages, but you need to understand their limitations and security issues to support them properly in your network environments.

Module Objectives

In this module, you will do the following:

- Summarize wireless standards.
- Install and configure secure wireless networks.
- Troubleshoot wireless networks.

Lesson 12.1

Wireless Concepts and Standards



Exam Objectives Covered

- 1.5 Compare and contrast transmission media and transceivers.
- 2.3 Given a scenario, select and configure wireless devices and technologies.

Wireless connectivity is a core feature of most network environments today because it supports users' need for mobility using different devices. To support a wireless network, you must be able to summarize the features of Wi-Fi, cellular, and satellite technologies.


As you study this lesson, answer the following questions:

- What are the differences between 802.11 standards, cellular, and satellite wireless technologies?
- How do wireless networks control media access?
- What are frequency bands and channels, and how do they affect wireless performance?

IEEE 802.11 Wireless Standards

IEEE **802.11 standards**, better known as Wi-Fi, define the Physical layer media by which data is encoded into a radio carrier signal using a modulation scheme. The properties of radio waves include amplitude (the height of peaks and troughs), frequency (the number of peaks per unit of time), and phase (the angle of a wave at a point in time). Modulation changes one or more of those properties to encode a signal. As well as modulation schemes, Wi-Fi standards use different carrier methods to provide sufficient resistance to interference from noise and other radio sources.

A wireless radio transmitting and receiving within a particular range of frequencies with the same modulation scheme is a half-duplex shared access medium (a physical bus). 802.11 uses **carrier sense multiple access with collision avoidance (CSMA/CA)** to cope with contention. Under CSMA/CA, when a station receives a frame, it performs error checking. If the frame is intact, the station responds with an acknowledgment (ACK). If the ACK is not received, the transmitting station resends the frame until timing out. 802.11 also defines a virtual carrier sense flow control mechanism to further reduce the incidence of collisions. A station broadcasts a request to send (RTS) with the source and destination and the time required to transmit. The receiving station responds with a clear to send (CTS), and all other stations in range do not attempt to transmit within that period.



The CSMA/CA media access method. (Images © 123RF.com.)

Wi-Fi typically uses a logical star topology to establish a wireless local area network (WLAN). A device called an access point (AP) implements the center of the star, mediating connections between client device radios, which are referred to as stations.

The original 802.11 Wi-Fi standard worked at just 1 Mbps, but like the 802.3 Ethernet standard, it has been revised many times, with each iteration specifying different signaling and transmission mechanisms. Products conforming to the various standards can be certified by the Wi-Fi Alliance (wi-fi.org).

IEEE 802.11a and 5 GHz Channel Bandwidth

A wireless radio is configured to use a **frequency band**. Support for a given Wi-Fi standard determines which bands are available. Within each band, the radio is configured to use a specific range of frequencies, referred to as a **channel**.

Wi-Fi standards can use different frequency bands. These frequency bands have different performance properties. The two most commonly used are 2.4 GHz and 5 GHz:

- 2.4 GHz is better at propagating through solid surfaces, making it ideal for providing the longest signal range. However, the 2.4 GHz band does not support many individual channels and is often congested, both with other Wi-Fi networks and other types of wireless technology, such as Bluetooth. Consequently, with the 2.4 GHz band, there is increased risk of interference, and the maximum achievable data rates are typically lower than with 5 GHz.
- 5 GHz is less effective at penetrating construction materials or solid furniture and so does not support the maximum ranges achieved with 2.4 GHz standards. However, the band supports more individual channels and suffers less from congestion and interference, meaning it supports higher data rates at shorter ranges.


The IEEE 802.11a standard specifies use of the 5 GHz frequency band and a multiplexed carrier scheme called Orthogonal Frequency Division Multiplexing (OFDM). 802.11a has a nominal data rate of 54 Mbps.

The 5 GHz band is subdivided into 25 non-overlapping channels, each of which is 20 MHz wide. However, some of the channels work in a frequency range also used by radar. The **802.11h** standard specifies a Dynamic Frequency Selection (DFS) method to scan for radar signals and prevent an access point from using channels that would cause interference. The exact use of channels can be subject to different regulations in different countries. Regulatory impacts also include a strict limit on power output, constraining the range of Wi-Fi devices.

IEEE 802.11b/g and 2.4 GHz Channel Bandwidth

The 802.11b standard uses the 2.4 GHz frequency band and was released in parallel with 802.11a. It standardized the use of the carrier method Direct Sequence Spread Spectrum (DSSS), along with Complementary Code Keying (CCK) signal encoding. While in some ways DSSS was an inferior technology to OFDM—with a nominal data rate of just 11 Mbps—802.11b products were quicker to market and became better established than 802.11a.

The 2.4 GHz band is subdivided into up to 14 channels, spaced at 5 MHz intervals from 2,412 MHz up to 2,484 MHz. Because the spacing is only 5 MHz and DSSS needs 22 MHz channel bandwidth, 802.11b channels overlap quite considerably. This means that co-channel interference is a real possibility unless non-overlapping channels are chosen (1, 6, and 11, for instance). Also, in the Americas, regulations permit the use of channels 1–11 only, while in Europe channels 1–13 are permitted, and in Japan all 14 channels are permitted.



Channel overlap in the 2.4 GHz band.

The 802.11g standard offered a relatively straightforward upgrade path from 802.11b. Like 802.11a, 802.11g uses OFDM, but in the 2.4 GHz band used by 802.11b and with the same channel layout. This made it straightforward for vendors to offer 802.11g devices that could offer backward support for legacy 802.11b clients. 802.11g has a nominal data rate of 54 Mbps. When in 802.11b compatibility mode, it drops back to using DSSS.

IEEE 802.11n, MIMO, and Channel Bonding

The 802.11n standard increases bandwidth by multiplexing the signals from two to four separate antennae (a radio chain) using a collection of technologies generally referred to as **multiple input multiple output (MIMO)**. The configuration of an 802.11n radio chain is identified by AxB:C notation, where A is the number of transmit antennae, B is the number of receive antennae, and C is the number of simultaneous transmit and receive streams. The maximum possible is 4x4:4, but common configurations are 2x2:2 or 3x3:2. For example, a 4x4:4 access point could allocate two streams carrying different data to a 2x2:2 client, increasing bandwidth. This is referred to as spatial multiplexing.

Having more transmit and receive antennae can also be used to improve signal reliability, rather than boosting bandwidth. If the same data stream is sent by two or three transmit antennae, the receiver can combine them to derive a stronger signal and increase range at a given data rate. Similarly, multiple receive antennae can derive a stronger signal, even if there is only one transmit stream. This is referred to as spatial diversity. For example, 2x2:2 and 2x3:2 radio chains have the same throughput, but the 2x3:2 chain could make more use of spatial diversity to increase range.

802.11n products can use channels in the 2.4 GHz band and the 5 GHz band. 802.11n also allows two adjacent 20 MHz channels to be combined into a single 40 MHz channel, referred to as **channel bonding**. Due to the restricted bandwidth of 2.4 GHz, on a network with multiple access points, channel bonding is a practical option only in the 5 GHz band. The 5 GHz band has a wider frequency range, so it can provide up to 25 non-overlapping channels. However, those channels are not necessarily contiguous, which slightly reduces the options for bonded channels. DFS and local regulatory requirements can also impact channel availability.

	U-NII-1	U-NII-2	U-NII-2 Extended	U-NII-3
20 MHz	36 40 44 48 52 56 60 64		100 104 108 112 116 120 124 128 132 136 140 144	149 153 157 161 165
40 MHz	38 46 54 62		102 110 118 126 134 142	151 159
80 MHz	42 58		106 122 138	155
160 MHz	50		114	

Dynamic Frequency Selection (DFS) Range

Bonded channel options in the 5 GHz Unlicensed National Information Infrastructure (U-NII) sub-bands. Channels within the DFS range may be disabled if the site is near a radar transmitter.



Some 802.11n client adapters may support only the 2.4 GHz band. An access point (AP) or adapter that can support both is referred to as dual band. A dual band AP can support both 2.4 GHz and 5 GHz bands simultaneously. This allows legacy clients to be allocated to the 2.4 GHz band.

The data rate for 802.11n is 72 Mbps per stream. Assuming the maximum number of four spatial streams and optimum conditions, the nominal data rate could be as high as 600 Mbps for a 40 MHz bonded channel. 802.11n can work in high throughput (HT)/greenfield mode for maximum performance or HT mixed mode for compatibility with older standards (801.11a-ht, 802.11b-ht, and 802.11g-ht). Mixed mode reduces overall WLAN performance, as it involves the transmission of legacy identification and collision avoidance frames (HT protection) but not to the extent that 802.11n devices are reduced to, for example, 802.11g data rates. Operating in greenfield mode is likely to cause substantial interference if there are legacy WLANs operating nearby on the same channel(s).

In recent years, Wi-Fi standards have been renamed with simpler digit numbers. 802.11n is now officially designated as Wi-Fi 4.

Wi-Fi 5 and Wi-Fi 6

The Wi-Fi 5 (or 802.11ac) and Wi-Fi 6 (802.11ax) standards continue the development of Wi-Fi technologies to increase bandwidth and support modern network access patterns.

Wi-Fi 5 (802.11ac)

Wi-Fi 5 is designed to work only in the 5 GHz band. The 2.4 GHz band can be used for legacy standards (802.11g/n) in mixed mode. The aim for Wi-Fi 5 is to get throughput like that of Gigabit Ethernet or better. It supports more channel bonding (up to 80 or 160 MHz channels), up to eight spatial streams, rather than four, and denser modulation (at close ranges). The way Wi-Fi 5 uses the radio spectrum is designated as very high throughput (VHT).

Wi-Fi 5 access points are marketed using AC values, such as AC5300. The 5300 value represents 1,000 Mbps over a 40 MHz channel in the 2.4 GHz band and two 2,167 Mbps streams over 80 MHz channels in the 5 GHz band.



While there aren't 802.11ac standards for 2.4 GHz, vendors use proprietary extensions to claim higher maximum throughput than 802.11n's 600 Mbps.

Wi-Fi 6 (802.11ax)

Wi-Fi 6 uses advanced modulation and signal encoding to improve the amount of data sent per packet by about 40%. The way Wi-Fi 6 uses the radio spectrum is designated as high efficiency (HE) to reflect these improvements. The aim for Wi-Fi 6 is to approximate 10G connection speeds (AX11000). Like Wi-Fi 5, Wi-Fi 6 products are branded using the combined throughput. For example, AX6000 is nominally 1,148 Mbps on the 2.4 GHz radio and 4,804 over 5 GHz. Wi-Fi 6 also specifies use of a new 6 GHz frequency band, which is required to achieve the highest data rates.

In Wi-Fi 6, the OFDM with multiple access (OFDMA) modulation scheme allows sub-carriers or tones to be allocated in groups of different sizes, referred to as resource units (RUs), each of which can communicate in parallel. Where small RUs are used, this reduces throughput but provides more opportunities for a larger number of devices to transmit. The effect is to reduce latency where numerous small data packets are being transmitted. This technology provides better support for IoT devices. Stations that require more bandwidth can be assigned larger RUs. RUs can also be assigned based on class of service parameters, such as prioritizing voice over IP (VoIP) traffic. It also allows an access point to support legacy Wi-Fi 4/5 stations efficiently.

Multuser MIMO and Band Steering

In basic 802.11 operation modes, bandwidth is shared between all stations because of the CSMA/CA contention protocol. An AP can communicate with only one station at a time; multiple station requests go into a queue. Wi-Fi 5 and Wi-Fi 6 products address this problem using beamforming or **Multuser MIMO (MU-MIMO)**.

Multuser MIMO

Downlink MU-MIMO (DL MU-MIMO) allows the AP to use its multiple antennae to process a spatial stream of signals in one direction separately from other streams. This means that groups of stations on a different alignment can connect simultaneously and also obtain more bandwidth. For example, if four stations are positioned north, south, east, and west of a 4x4:4 AP, the AP should be able to allow each of them to connect at close to the maximum speed. If another station is added to the north, those two northern stations will share the available bandwidth along that beam path. Both stations and AP must support MU-MIMO. Where Wi-Fi 5 supports up to four stations communicating in parallel over 5 GHz only, Wi-Fi 6 can support up to eight in 2.4 GHz, 5 GHz, and 6 GHz bands, giving it better performance in congested areas.

With DL MU-MIMO, only the AP can initiate beamforming, so it is only available on the downlink from AP to station (not station to AP). Wi-Fi 6 supports uplink MU-MIMO (UL MU-MIMO), allowing stations to initiate beamforming with the access point.



For both Wi-Fi 5 and Wi-Fi 6, improvements are released to market in waves. For example, UL MU-MIMO was released in wave 2 Wi-Fi 6 products, which also added support for the 6 GHz frequency band.



MU-MIMO and OFDMA are different but complementary technologies. MU-MIMO makes use of spatial streams, where OFDMA makes flexible use of subcarriers within a channel. Both can work together to increase parallelism (supporting communication with more devices simultaneously).

Band Steering

Many Wi-Fi devices have dual-band (2.4 GHz and 5 GHz) or tri-band (2.4, 5, and 6 GHz) radios. In a site where multiple access points support dual-band or tri-band networks, a client device will use a combination of beacon messages from the access point and a measure of signal strength to determine which band to use. In some circumstances, the network designer might want to exert more control over this process. **Band steering** allows an access point to make it more likely that a client will connect to the 5 GHz or 6 GHz band than the 2.4 GHz band. It does this by reducing the number of beacons used to advertise the 2.4 GHz network. The goal is to restrict the use of the 2.4 GHz band to devices with no 5/6 GHz capability.



Another approach is to use different network names for 2.4 GHz and 5/6 GHz networks. Alternatively, a network designer may prefer to allow client devices to select the best band. The AP does not use signal strength as a factor in band steering, so it can lead to poor performance in some circumstances.

Cellular Technologies

Where Wi-Fi is typically operated as private infrastructure, **cellular radio** is operated by telecommunications providers. A cellular radio establishes a connection with the nearest available cell or base station. Each base station has an effective range of up to 5 miles (8 km). The base station links the device to global telecommunications networks.

Cellular digital communications standards are described as belonging to a generation. 2G and 3G cellular networks were implemented by different technologies in different parts of the world and suffered from low data rates. Converged 4G and 5G standards are supported by telecommunications providers worldwide. As well as faster mobile speeds, 4G and 5G can provide fixed-wireless broadband solutions for homes and businesses, and support IoT networks. 4G and 5G devices must be installed with a Subscriber Identity Module (SIM) card or chip issued by the network provider.

4G/Long Term Evolution

The **Long Term Evolution (LTE)** 4G standard has a maximum downlink of 150 Mbps in theory, but around 20 Mbps is typical of real-world performance. LTE Advanced (LTE-A) specifies a 300 Mbps downlink, but this aspiration is not matched by real-world performance. The current typical performance for LTE-A is up to 90 Mbps.

5G

Compared to earlier cellular technologies, 5G can use a broader radio spectrum, from low (sub-1 GHz) to medium/high (6 GHz to 40 GHz). Low bands, such as 900 MHz or 1,900 MHz, have greater range and penetrating power. High bands, also referred to as millimeter wave (mmWave), require close range (a few hundred feet) and cannot penetrate walls or windows. Consequently, design and rollout of 5G services is relatively complex. Rather than a single large antenna serving a large wireless cell, 5G involves installing hundreds of smaller antennae to form an array that can take advantage of multipath and beamforming to overcome the propagation limitations of the spectrum. This is also referred to as massive MIMO.

In theory, 5G has a maximum peak rate of 20 Gbps. As with 4G, real-world speeds are nowhere near the peak rate, ranging from about 50 Mbps to 300 Mbps at time of writing.



Cellular data rates vary widely from country to country and from region to region. The rates given here are only illustrative.

Satellite Technologies

Satellite systems use microwave dishes aligned to orbital satellites that can either relay signals between sites directly or via another satellite. The widespread use of satellite television receivers allows for domestic Internet connectivity services over satellite connections. Satellite services for business are also expanding, especially in rural areas where DSL or cable services are unlikely to be available.

Geostationary Orbital Satellite Internet Access

Historically, data communications satellites were placed in high geostationary orbit. One drawback of this is increased latency. The signal must travel over thousands of miles more than terrestrial connections, introducing a delay of many times what might be expected over a land link. For example, if accessing an Internet web server over DSL involves a 10–20 ms round trip time (RTT) delay on the link, accessing the same site over a satellite link could involve a 600–800 ms RTT delay. This is an issue for real-time applications, such as video conferencing, VoIP, and multiplayer gaming.



RTT is the two-way latency, or the time taken for a probe to be sent and a response to be received.

To create a satellite Internet connection, the ISP installs a very small aperture terminal (VSAT) satellite dish antenna at the customer's premises and aligns it with the orbital satellite. Because the satellite does not move relative to the dish, there should be no need for any realignment. The antenna is connected via coaxial cabling to a digital video broadcast satellite (DVB-S) modem. The transfer rates available vary between providers and access packages, but 2 or 6 Mbps up and 30 Mbps down would be typical.

Low Earth Orbital Satellite Internet Access

A different type of service uses an array of satellites positioned in low Earth orbit (LEO). LEO satellites support better bandwidth (around 70–100 Mbps at the time of writing) and are lower latency (100–200 ms RTT). The drawback is that the satellites move relative to the surface of Earth. The customer's premises antenna must be provisioned with a motor so that it can periodically realign with the array. The dish construction uses a technology called phased array to connect to different satellites as they pass overhead and minimize the amount of mechanical realignment required. The antenna must have a clear view of the whole sky.

Global Positioning System

Satellites are also used to implement the Global Positioning System (GPS). GPS allows a device with a suitable sensor to triangulate its position using signals from orbital satellites. As this triangulation process can be slow, most smartphones and laptops use Assisted GPS (A-GPS) to obtain coordinates from the nearest cell tower and adjust for the device's position relative to the tower. A-GPS uses cellular data. GPS satellites are operated by the US government. Some GPS sensors can use signals from other satellites, operated by the EU (Galileo), Russia (GLONASS), or China (BeiDou).

Lesson 12.2

Enterprise Wireless Network Design



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 2.3 Given a scenario, select and configure wireless devices and technologies.
- 3.1 Explain the purpose of organizational processes and procedures.

Designing a wireless network to meet the requirements of multiple types of devices and different office layouts can be a complex task. Planning the installation using survey tools can ensure that range and interference issues are accounted for at the outset and do not become support issues later.

As you study this lesson, answer the following questions:

- What are the functions of SSIDs, BSSIDs, and ESSIDs, and how do they compare to BSAs and ESAs?
- How can an infrastructure network cover a wider area than the range of a single access point?
- How do you measure wireless signal strength?
- Why might different antenna types be used in a wireless network design?
- What are the uses of ad hoc, mesh, and point to point wireless network types?

Infrastructure Network Type


At the Data Link layer, Wi-Fi uses similar framing and MAC addressing concepts to Ethernet. Each radio has a base MAC address, assigned by the vendor. For example, a dual band adapter will have one MAC address for the 2.4 GHz radio and one for the 5 GHz radio. Each radio is referred to as a station, though the term can also be used to refer to a wireless client device.

Each station can be configured to join a WLAN through the network name or **Service Set Identifier (SSID)**. An SSID can be up to 32 bytes in length and for maximum compatibility should only use ASCII letters and digits plus the hyphen and underscore characters. If the SSID is broadcast, the user can select it from a list. If it is not broadcast, the SSID must be added manually to the client's wireless configuration. When a station attempts to connect, it probes to establish 802.11 standards support and security configuration, completes any authentication required of it, and then requests to associate with the network.

Most sites use the infrastructure network type. An infrastructure topology means that each station is configured to connect to the WLAN via an access point (AP). This makes a logical star topology. Each client station requires a wireless adapter compatible with the standard(s) supported by the AP. The AP mediates communications between client devices using a specific frequency and channel. In 802.11 documentation, this arrangement is referred to as an infrastructure Basic Service Set (BSS). A virtual MAC address derived from one of the AP's radios is used as the **Basic Service Set Identifier (BSSID)**. The BSSID identifies the access point hosting a BSS.

Any given access point could operate multiple BSSs on its 2.4 GHz, 5 GHz, and 6 GHz radios. These BSSs could be configured with separate network names (SSIDs), or the same network name, depending on how many logical networks are required. Each network can have different security properties. For example, you might configure a guest network with no authentication, and a home network that requires a passphrase to join.


If a client connects to an SSID underpinned by multiple BSSs, it tries to determine which BSSID offers the best signal.



Access point hosting multiple networks. The CORPNET and GUEST networks are offered on both 2.4 GHz and 5 GHz bands. The IoT network is only offered on the 2.4 GHz band. Each BSS (a network name on a particular band) is identified by a BSSID MAC address. (Image © 123RF.com.)

As well as facilitating communications between stations, an access point can provide a bridge to a cabled network segment. The cabled network is referred to as a distribution system (DS). On an enterprise WLAN, each access point is cabled to a switch. The distribution system can be used to network multiple access points.

Multiple APs connected to the same distribution system can host an Extended Service Set (ESS). An ESS is a group of basic service sets that are all configured with the same SSID and security information. Each BSS uses a different channel within each frequency band. When deployed like this, the network name is more properly called an **Extended SSID (ESSID)**. An ESSID has the same format as an SSID.



WLAN configuration in infrastructure mode. (Images © 123RF.com.)



This infrastructure isn't limited to running a single ESS. Each ESS could be mapped to a VLAN to segment the traffic as it is carried around the switched network.

Range and Signal Strength

A device supporting the Wi-Fi standard should have an indoor range of at least 30 m (100 feet). 2.4 GHz radios support better ranges than 5 GHz ones, and 802.11n and later standards improve range compared to earlier standards. Outdoor range can be double or triple indoor range. Each station uses the Dynamic Rate Switching/Selection (DRS) mechanism to determine an appropriate data rate based on the signal quality. If the signal is strong, the station will select the highest available data rate (determined by support for a given iteration of the 802.11 standard); if the signal is weak, the station will reduce the data rate.

Radio signals pass through solid objects, such as brick or drywall, but can be weakened or blocked by particularly thick walls or those of dense concrete or metal construction. Other radio-based devices can cause interference as can devices such as fluorescent lighting, microwave ovens, cordless phones, and (in an industrial environment), power motors and heavy machinery. Bluetooth uses the 2.4 GHz frequency range but a different modulation technique, so interference is possible but not expected.

Signal strength is represented as the ratio of a measurement to 1 milliwatt (mW), where 1 mW is equal to 0 dBm. dB and dBm units can be combined to analyze losses and gains in signal strength along a communications path. For example, if you transmit a radio signal at 1 mW and use an antenna to boost the signal, the effective power is:

$$0 \text{ dBm} + 3 \text{ dB} = 2 \text{ mW} = \sim 3 \text{ dBm}$$

Because 0 dBm is 1 mW, a negative value for dBm represents a fraction of a milliwatt. For example, -30 dBm is 0.001 mW; -60 dBm is 0.000001 mW. Wi-Fi devices are all constrained by regulations governing radio spectrum use and output only small amounts of power.

The **Received Signal Strength Indicator (RSSI)** is the strength of the signal from the transmitter as measured at the client end. When measuring RSSI, dBm will be a negative value with values closer to zero representing better performance. -30 dBm is considered to be a perfect signal. A value around -65 dBm represents a good signal, while anything worse than -80 dBm is likely to suffer packet loss or be dropped. The RSSI must exceed the minimum receiver sensitivity.



Depending on the vendor, RSSI might be measured directly in dBm or might be an index value related to a scale of dBm measurements. RSSI indices can be measured as 0–60, 0–127, or as 0–255. On a client, this index is displayed as a number of bars of signal strength on the adapter icon.

The comparative strength of the data signal to the background noise is called the signal-to-noise ratio (SNR). Noise is also measured in dBm, but here values closer to zero are less welcome as they represent higher noise levels. For example, if signal is -65 dBm and noise is -90 dBm, the SNR is the difference between the two values, expressed in dB (25 dB). If noise is -80 dBm, the SNR is 15 dB and the connection will be much, much worse.

Wireless Surveys and Heat Maps

Range, interference, and signal strength constitute a complex set of factors that must be considered when designing a wireless network. Additionally, the designer must account for the number of devices in close proximity (device density), how much they move around the site, and special traffic requirements, such as voice or video conferencing. A **wireless survey** is a critical planning tool to ensure that the WLAN delivers acceptable data rates to the supported number of devices in all the physical locations expected.

The area served by a single AP is referred to as a basic service area (BSA). The area where stations can roam between access points to stay connected to the same ESSID is an extended service area (ESA). A wireless site survey ensures that these areas are properly sized.


A wireless survey is performed first by examining the blueprints or floor plans of the premises to understand the layout and to identify features that might produce radio frequency interference (RFI). This can be backed up by a visual inspection that may reveal things that are not shown on the blueprints, such as thick metal shelving surrounding a room that needs to have WLAN access. Each AP mounting point needs a network port and power jack, so it will help to obtain plans that show the locations of available ports.



Alternatively, Power over Ethernet (PoE) allows a switch to supply power to the AP over data cabling.

The next step is to create a new plan on which you will mark the basic service areas and associated APs and booster antennae. The idea here is to place APs close enough together to avoid “dead zones”—areas where connectivity is difficult or data transfer rates are below an acceptable tolerance level—but far enough apart that one AP does not interfere with another or that one AP is overutilized and a nearby one underutilized.

Position an AP in the first planned location, then use a laptop with a wireless adapter and a wireless survey tool to record signal strength and supported data rate at various points in the intended BSA. Many tools can show the signal strength within a particular channel obtained in different locations as a graphical **heat map**. The heat map would show areas with a strong signal in greens and yellows with warning oranges and reds where signal strength drops off. This step is then repeated for each planned location. Neighboring APs should be configured with non-overlapping channels to avoid interfering with one another. It may also be necessary to adjust the transmit power of an AP to size its BSA appropriately. Transmit power is a setting that configures the device to use less than its maximum output.



Heat map generated by Ekahau Site Survey. (Image © Ekahau Inc.)



The network design might call for different power levels for different bands. For example, you might lower power on the 2.4 GHz band to reduce its range and make it more likely that clients will connect on the 5 GHz band. Alternatively, you might configure band steering, or use different ESSIDs for 2.4 GHz and 5 GHz networks.

Wireless Roaming

Clients can **roam** within an extended service area (ESA). An ESA is created by installing APs with the same ESSID and security configuration connected by a wired network, or **distribution system (DS)**. The access points are configured with different channels so that where BSAs overlap, there is no interference. When the client detects that it is no longer receiving a good signal, it checks for another signal with the same ESSID on other channels or on a different frequency band, and if there is a stronger signal, it disassociates from the current AP. The station can then reassociate with the new AP. Depending on the roaming infrastructure and security type, the station may have to reauthenticate, or if 802.11r fast roaming is supported, it may be able to use its existing authentication status to generate security properties for the new association.

Roaming is supposed to be seamless, but in practice reestablishing the connection can often cause time-out problems for applications. To improve mobility, there needs to be a balance between determining what constitutes a “good” signal and the rate at which a client tries to associate with different APs. Many adapters support a roaming “aggressiveness” setting that can be configured to prevent a Wi-Fi adapter “flapping” between two APs or (conversely) to prevent a client from remaining associated with a more distant AP when it could achieve better bandwidth through one closer to it.

SSID Broadcast and Beacon Frame

A WLAN is typically configured to advertise its presence by broadcasting the SSID. This allows a user to connect to a named network from a list reported by the client adapter. If SSID broadcast is suppressed, the user must configure the connection to the network manually. A beacon is a special management frame broadcast by the AP to advertise the WLAN. The beacon frame contains the SSID/ESSID (unless broadcast is disabled), the BSSID, supported data rates and signaling, plus encryption/authentication requirements. The interval at which the beacon is broadcast (measured in milliseconds) can be modified. The default is usually 100 ms. Increasing the interval reduces the overhead of broadcasting the frame but delays joining the network and can hamper roaming between APs.



Even if SSID broadcast is suppressed, it is fairly easy for a network sniffer to detect it as clients still use it when connecting with the AP.

Wireless Distribution System

You can also configure access points to create wireless-only networks in areas where it is not possible to run cabling. This is referred to as a wireless distribution system (WDS). You must set the APs to use the same channel, SSID, and security parameters. The APs are configured in WDS/repeater mode. One AP is configured as a base station, while the others are configured as remote stations. The base station can be connected to a cabled segment. The remote stations must not be connected to cabled segments. The remote stations can accept connections from client stations and forward all traffic to the base station.



WDS support and implementation can vary between manufacturers. If you are implementing WDS, it is usually best to use APs from the same vendor.

Wireless Controllers

An enterprise network might require the use of tens or hundreds of access points (APs). If APs are individually managed, this can lead to configuration errors on specific APs and can make it difficult to gain an overall view of the wireless deployment, including which clients are connected to which APs and which clients or APs are producing the most traffic.



A wireless controller, an enterprise-level appliance capable of supporting up to 1,500 APs and 20,000 clients. (Image © 123RF.com.)

Rather than configure each device individually, enterprise wireless solutions allow for centralized management and monitoring of the APs on the network. This is typically achieved through the use of a dedicated hardware device called a **wireless controller**.

An AP whose firmware contains enough processing logic to be able to handle clients without the use of a wireless controller is known as an **autonomous AP**, while one that requires a wireless controller to function is a **lightweight AP**. Cisco wireless controllers usually communicate with the APs using the Lightweight Access Point Protocol (LWAPP). LWAPP allows an AP configured to work in lightweight mode to download an appropriate SSID, standards mode, channel, and security configuration. Alternatives to LWAPP include the derivative Control And Provisioning of Wireless Access Points (CAPWAP) protocol or a proprietary protocol.


As well as autoconfiguring the appliances, a wireless controller can aggregate client traffic and provide a central switching and routing point between the WLAN and wired LAN. It can assign clients to separate VLANs. Automated VLAN pooling ensures that the total number of stations per VLAN is kept within specified limits, reducing excessive broadcast traffic.

Antenna Types

The **antenna type** determines the propagation pattern or shape of the radio waves transmitted. Most wireless radios are fitted with omnidirectional vertical rod-type antennae. An omnidirectional antenna receives and sends signals in all directions more or less equally. Access points with omnidirectional antennae should ideally be ceiling mounted for best coverage. The propagation pattern is shaped like a torus (donut), rather than a sphere, and radiates more powerfully in the horizontal plane than it does in the vertical plane. Locating the antenna above head height will minimize interference from obstructing furniture by allowing line of sight to most connecting devices, but positioning it too high (above around 25 ft) will reduce signal strength, especially for stations directly below the antenna. You can obtain APs with downtilt omnidirectional antennae for use on high ceilings.

To extend the signal to a particular area, you can use a unidirectional antenna focused in a single direction. Both the sender and receiver must use directional antennae, or one will be able to receive signals but not send responses. Unidirectional antenna types include the Yagi (a bar with fins) and parabolic (dish or grid) form factors. Unidirectional antennae are useful for point to point wireless bridge connections.

The increase in signal strength obtained by focusing the signal is referred to as the gain and is measured in dBi (decibel isotropic). The amount of directionality, referred to as the beamwidth, is measured in degrees. A pair of 10-degree antennae is highly directional and will require more exact alignment than a pair of 90-degree antennae.



A variety of generic antenna types: from left to right, a vertical rod antenna, a Yagi antenna, a parabolic/dish antenna, and a parabolic grid antenna.

Polarization refers to the orientation of the wave propagating from the antenna. If you imagine a rod-type antenna, when the rod is pointed up relative to the floor, the wave is horizontally polarized; if you orient the rod parallel to the floor, the wave is vertically polarized. To maximize signal strength, the transmission and reception antennae should normally use the same polarization. This is particularly important when deploying unidirectional antennae for a point to point link. Some antennae are dual-polarized, meaning that they can be installed in either orientation. Dual-polarized antennae are also the best way to support mobile devices, as these can be held by their user in a variety of orientations.

Other Wireless Network Types

While most enterprise and SOHO networks are configured in infrastructure mode, there are also wireless network types that allow stations to establish peer-to-peer links.

Ad Hoc Topology

In an ad hoc topology, the wireless adapter allows connections to and from other devices. In 802.11 documentation, this is referred to as an **Independent Basic Service Set (IBSS)**. This topology does not require an access point. All the stations within an ad hoc network must be within range of one another. An ad hoc network might suit a small workgroup of devices, or connectivity to a single device, such as a shared printer, but it is not scalable to large network implementations.



IBSS is not supported by the updated WDI driver model in the latest versions of Windows (docs.microsoft.com/en-us/windows-hardware/drivers/network/wdi-features-not-carried-over-in-wdi). Peer-to-peer connections are more likely to use Wi-Fi Direct. Wi-Fi Direct allows a device such as a printer to operate a limited type of access point to allow clients to send print jobs wirelessly.

Mesh Topology

The 802.11s standard defines a **wireless mesh network (WMN)**. There are also various proprietary mesh protocols and products. Unlike an ad hoc network, nodes in a WMN are capable of discovering one another and peering, forming a Mesh Basic Service Set (MBSS). The mesh stations can perform path discovery and forwarding between peers using a routing protocol, such as the Hybrid Wireless Mesh Protocol (HWMP).

These features make a mesh topology more scalable than an ad hoc topology because the stations do not need to be within direct radio range of one another—a transmission can be relayed by intermediate stations. Mesh topologies are becoming increasingly popular and are the foundation of most internet of things (IoT) networks.

Point to Point

A point to point link means a physical and logical connection between two devices. A wireless point to point link is usually used as a means of bridging two locations when it is not possible to connect them using cables. For example, two office buildings could be connected by installing highly directional dish or Yagi antennae on the roofs. The antennae would be carefully aligned to point directly at one another. The antennae could be connected to access points configured in bridge mode. Traffic that needs to be sent from one office to another would pass over the wireless bridge.

Lesson 12.3

Wireless Security



Exam Objectives Covered

- 2.3 Given a scenario, select and configure wireless devices and technologies.
- 4.1 Explain the importance of basic network security concepts.
- 4.2 Summarize various types of attacks and their impact to the network.

Wireless networking is popular with users but also poses considerable risk to the whole network unless it is secured with access controls. In this topic, you will identify different wireless security methods and their configuration requirements.

As you study this lesson, answer the following questions:

- What is the strongest wireless encryption method, and what are the risks of using weaker encryption standards?
- What are the differences between personal, enterprise, and open authentication modes, and what is the impact on encryption?
- How can users' personal devices and guest networks be supported securely?
- What techniques do threat actors use to exploit vulnerabilities in wireless networks?

Wi-Fi Encryption Standards

As well as the site design, a wireless network must be configured with security settings. Without encryption, anyone within range can intercept and read packets passing over the wireless network. The choice of which security settings to apply is determined by device support for the various Wi-Fi encryption standards, by the type of authentication infrastructure, and by the purpose of the WLAN. The encryption standard determines the cryptographic protocols that are supported, the means of generating the encryption key, and available methods for authenticating wireless stations when they try to associate with the network.

The first version of **Wi-Fi Protected Access (WPA)** was designed to fix critical vulnerabilities in the earlier wired equivalent privacy (WEP) standard. Like WEP, version 1 of WPA uses the RC4 stream cipher to encrypt traffic but adds a mechanism called the Temporal Key Integrity Protocol (TKIP) to try to mitigate the various attacks against WEP that had been developed.

Personalize settings for each band or enable Smart Connect to configure the same settings for all bands.

OFDMA: ☒ Enable ?

Smart Connect: ☐ Enable ?

2.4GHz: ☒ Enable [Sharing Network](#)

Network Name (SSID): ☐ Hide SSID

Security:

Version:

Encryption:

Password:

Transmit Power:

Channel Width:

Channel:

Mode:

5GHz: ☒ Enable [Sharing Network](#)

Network Name (SSID): ☐ Hide SSID

Security:

Version:

Password:

Transmit Power:

Channel Width:

Channel:

Mode:

Configuring a TP-LINK SOHO access point with wireless encryption and authentication settings. In this example, the 2.4 GHz band allows legacy connections with WPA2-Personal security, while the 5 GHz network is for 802.11ax (Wi-Fi 6) capable devices using WPA3-SAE authentication. (Screenshot used with permission from TP-Link Technologies.)

Neither WEP nor the original WPA version are considered secure enough for continued use. They can be exploited by various types of replay attack that aim to recover the encryption key. WPA2 uses the Advanced Encryption Standard (AES) cipher deployed within the Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP). AES replaces RC4, and CCMP replaces TKIP. CCMP provides authenticated encryption, which is designed to make replay attacks harder.

Weaknesses have also been found in WPA2, however, which has led to its replacement by WPA3. The main features of WPA3 are as follows:

- **Simultaneous Authentication of Equals (SAE)**—WPA2 uses a four-way handshake to allow a station to authenticate its credential, exchange a key to use for data encryption, and establish an association with an access point. This four-way handshake mechanism is vulnerable to manipulations that allow a threat actor to recover the key. WPA3 replaces the four-way handshake with the more secure SAE mechanism.
- **Updated cryptographic protocols**—WPA3 replaces AES CCMP with the stronger AES Galois Counter Mode Protocol (GCMP) mode of operation.
- **Protected management frames**—Management frames are used for association and authentication and for disassociation and deauthentication messages as devices join and leave the network. These frames can be spoofed and misused in various ways under WPA and WPA2. WPA3 mandates use of encryption for these frames to protect against key recovery attacks and DoS attacks that force stations to disconnect.

- **Wi-Fi Enhanced Open**—An open Wi-Fi network is one with no passphrase. Any station can join the network. In WPA2, this also means that all traffic is unencrypted. WPA3 encrypts this traffic. This means that any station can still join the network, but traffic is protected against sniffing.

Personal Authentication

In order to secure a network, you need to be able to confirm that only valid users are connecting to it. Wi-Fi authentication comes in three types: personal, open, and enterprise. Within the personal authentication category, there are two methods: pre-shared key authentication (PSK) and Simultaneous Authentication of Equals (SAE).

WPA2 Pre-Shared Key Authentication

In WPA2, **pre-shared key (PSK)** authentication uses a passphrase to generate the key that is used to encrypt communications. It is also referred to as group authentication because a group of users shares the same secret. When the access point is set to WPA2-PSK mode, the administrator configures a passphrase of between eight and 63 characters. This is converted to a type of hash value, referred to as the pairwise master key (PMK). The same secret must be configured on the access point and on each node that joins the network. The PMK is used as part of WPA2's four-way handshake to derive various session keys.

WPA2-PSK authentication has been shown to be vulnerable to attacks that attempt to recover the passphrase. At a minimum, the passphrase must be at least 14 characters long to try to mitigate risks from cracking.

WPA3 Personal Authentication

While WPA3 still uses a passphrase to authenticate stations in personal mode, it changes the method by which this secret is used to agree session keys. The scheme used is also referred to as Password Authenticated Key Exchange (PAKE). In WPA3, the **Simultaneous Authentication of Equals (SAE)** protocol replaces the four-way handshake.




The configuration interfaces for access points can use different labels for these methods. You might see WPA2-Personal and WPA3-SAE rather than WPA2-PSK and WPA3-Personal, for example. Additionally, an access point can be configured for WPA3 only or with support for legacy WPA2 (WPA3-Personal Transition mode).

Enterprise Authentication

The main problems with personal modes of authentication are that distribution of the key or passphrase cannot be secured properly and that users may choose insecure phrases. Personal authentication also fails to provide accounting, as all users share the same credential.

As an alternative to personal authentication, WPA's **enterprise authentication** method implements IEEE 802.1X to use an Extensible Authentication Protocol (EAP) mechanism to authenticate against a network directory. 802.1X defines the use of EAP over Wireless (EAPoW) to allow an access point to forward authentication data without allowing any other type of network access. It is configured by selecting WPA2-Enterprise or WPA3-Enterprise as the security method on the access point.

With enterprise authentication, when a wireless client requests an association, the AP enables the channel for EAPoW traffic only. In this context, the wireless client is referred to as the supplicant device. The AP passes the credentials of the supplicant to an AAA (RADIUS or TACACS+) server on the wired network for validation. When the supplicant has been authenticated, the AAA server transmits a master key (MK) to the supplicant. The supplicant and authentication server then derive the same pairwise master key (PMK) from the MK. The AAA server transmits the PMK to the access point. The wireless station and access point use the PMK to derive session keys, using either the WPA2 four-way handshake or WPA3 SAE methods.



Using Cisco's Virtual Wireless LAN Controller to set security policies for a WLAN—This policy enforces use of WPA2 and the use of 802.1X (Enterprise) authentication.
(Image © and Courtesy of Cisco Systems, Inc. Unauthorized use not permitted.)

Guest Networks and Captive Portals

Configuring an access point for **open authentication** means that the client is not required to authenticate. With WPA2, this also means that data sent over the link is unencrypted. WPA3 uses the W-Fi Enhanced Open mechanism to encrypt open authentication traffic.

Open authentication would be used on a public AP or “hotspot” or on a private network to facilitate guest connections. In the latter scenario, separate SSIDs are created for employee and guest access. As well as allowing different authentication requirements, this keeps the traffic for each network separate. The administrator can also apply more restrictions to the guest network, such as only allowing Internet access, rather than access to LAN servers. Guest traffic can be fully segmented from employee traffic, improving security.

Open authentication may be combined with a secondary authentication mechanism managed via a browser. When the client associates with the open hotspot and launches the browser, the client is redirected to a **captive portal**. This will allow the client to authenticate to the hotspot provider's network (over HTTPS, so the login is secure). The portal may also be designed to enforce terms and conditions and/or take payment to access the Wi-Fi service.



Captive portal issues can arise when the redirect does not work. The captive portal should use HTTPS. Most modern browsers will block redirection to sites that do not use TLS. This means that the captive portal also needs to be installed with a digital certificate issued by a certification authority (CA) that is trusted by the client browser.

When using open wireless, users must ensure they send confidential web data only over HTTPS connections and only use messaging and file transfer services with TLS enabled. Another option is for the user to join a virtual private network (VPN). The user would associate with the open hotspot, then start the VPN connection. This creates an encrypted “tunnel” between the user’s computer and the VPN server. This allows the user to browse the web or connect to email services without anyone eavesdropping on the open Wi-Fi network being able to intercept those communications. The VPN could be provided by the user’s company, or they could use a third-party VPN service provider. Of course, if using a third party, the user needs to be able to trust them implicitly. The VPN must use certificate-based tunneling to set up the “inner” authentication method.

Bring Your Own Device Issues

Bring Your Own Device (BYOD) is a smartphone/tablet provisioning model that allows users to select a personal device to use to interact with corporate network services and cloud apps. Allowing user selection of devices introduces numerous compatibility, support, and security challenges:

- **Compatibility/support**—The wide range of devices, mobile OS versions, and vendor support for patches make the job of ensuring that each device can connect to corporate network apps and data resources complex.
- **Security**—This device variety also causes security issues, especially in terms of unpatched devices. Another issue is that the device is not fully under the administrative control of the IT department. An insider threat actor could install apps that might pose a risk to corporate data or misuse the device to exfiltrate data.

Some of the impact of these issues can be mitigated through the use of enterprise mobility management (EMM) suites and corporate workspaces. EMM (or mobile device management) is a type of network access control solution that registers devices as they connect to the network. It can then enforce security policies while the device is connected. These might restrict use of device functions or personal apps. A corporate workspace is an app that is segmented from the rest of the device and allows more centralized control over corporate data. Users must also agree to acceptable use policies, which might prohibit installing non-store apps and rooting/jailbreaking a device and require keeping the device up to date with patches. Users might have to permit some level of inspection of the device to protect corporate data.

Wireless Network Attacks

Wireless networks can open several avenues for a threat actor to gain unauthorized network access.


Rogue Access Points

A **rogue access point** is one that has been installed on the network without authorization, whether with malicious intent or not. A malicious user can set up such an AP with something as basic as a smartphone with tethering capabilities, and a non-malicious user could enable such an AP by accident. If connected to a LAN without security, an unauthorized AP creates a backdoor through which to attack the network.

Evil Twins

A rogue AP masquerading as a legitimate one is called an **evil twin**. An evil twin might advertise a similar network name (SSID) to the legitimate one. For example, an evil twin might be configured with the network name “compeny” where the legitimate network name is “company.” Alternatively, the evil twin might spoof the SSID and BSSID (MAC address) of an authorized access point, and then the attacker might use some DoS technique to overcome the legitimate AP. After a successful DoS attack, the users will be forced to disconnect from the network and then manually attempt to reconnect. At that point, with many users busy and trying to get back to work, some or all may associate with the evil twin AP and submit the network passphrase or their credentials for authentication.

However it is configured, when a user connects to an evil twin, it might be able to harvest authentication information and, if it is able to provide wider network or Internet access, execute an on-path attack to snoop on connections established with servers or websites.



Surveying Wi-Fi networks using Xirrus Wi-Fi Inspector (xirrus.com)—Note the presence of print devices configured with open authentication (no security) and a smart TV appliance (requiring authentication). (Screenshot used with permission from Xirrus.)

One solution to the risk of rogue access points is to use EAP-TLS security so that the authentication server and clients perform mutual authentication. There are also various scanners and monitoring systems that can detect rogue APs, referred to as a wireless intrusion detection system (WIDS) or wireless intrusion prevention system (WIPS).

Deauthentication Attacks

The use of an evil twin may be coupled with a **deauthentication attack**. This sends a stream of spoofed management frames to cause a client to deauthenticate from an AP. This might allow the attacker to interpose the evil twin, sniff information about the authentication process, or perform a denial of service (DoS) attack against the wireless infrastructure. The attacks can be mitigated if the wireless infrastructure supports Management Frame Protection (MFP/802.11w). Both the AP and clients must be configured to support MFP.

```

root@bt: ~ - Shell - Konsole
Session Edit View Bookmarks Settings Help
CH 5 | Elapsed: 16 mins | 2011-06-03 08:57 | WPA handshake: 00:26:44:
BSSID PWR RXQ Beacons #Data, #/s CH MB ENC CIPHER A
00:14:7C:BA:0B:36 0 92 9002 35813 126 5 54 WEP WEP 0
BSSID STATION PWR Rate Lost Packets Probes
00:14:7C:BA:0B:36 00:16:E0:03:2D:82 0 0 1 169169 101108

root@bt: ~ - Shell - Konsole <3>
Session Edit View Bookmarks Settings Help
Aircrack-ng 1.0 r1645
[00:00:12] Tested 123497 keys (got 29328 IVs)
KB depth byte(vote)
0 0/ 2 25(45056) 61(40704) 5B(38656) 5C(37376) 74(37120)
1 0/ 14 D1(40192) 3D(39680) 2C(36864) 4B(36608) 6C(36608)
2 1/ 7 50(40448) 52(39168) 54(37632) 0C(37376) F8(37120)
3 0/ 15 7C(40192) BA(38656) 46(38144) 10(37808) D1(36864)
4 40/ 43 EF(33792) 2C(33536) 2F(33536) 4C(33536) 98(33536)
KEY FOUND! [ 25:D1:50:7C:6D ]
Decrypted correctly: 100%
root@bt: #

root@bt: ~ - Shell - Konsole <4>
Session Edit View Bookmarks Settings Help
Wrote packet to: classroom-arp
root@bt: # aireplay-ng --interactive -r classroom-arp mon0
For information, no action required: Using gettimeofday() instead of /dev/rtc
No source MAC (-h) specified. Using the device MAC (00:16:E0:03:2D:82)

Size: 68, FromDS: 0, ToDS: 1 (WEP)
BSSID = 00:14:7C:BA:0B:36
Dest. MAC = FF:FF:FF:FF:FF:FF
Source MAC = 00:16:E0:03:2D:82
0x0000: 0041 0201 0014 7cba db36 0016 e003 2d82 .A...].6....
0x0010: ffff ffff ffff 0001 75a0 2a00 cbc1 5200 .....u*...R.
0x0020: c054 6b55 977a ac09 44fa 57ce 8a5c ed58 .TRU.z..D.W..\X
0x0030: 41d0 ae74 82a9 db4a bc65 8d71 c543 85dc A.t...J.e.q.C..
0x0040: cd00 810f .....
Use this packet ? y
Saving chosen packet in replay_src-0603-085243.cap
You should also start airodump-ng to capture replies.
Sent 17505 packets...(490 pps)
415 packets (got 60225 ARP requests and 190059 ACKs), sent 43575 packets
626 packets (got 60291 ARP requests and 190208 ACKs), sent 43625 packets
842 packets (got 60361 ARP requests and 190343 ACKs), sent 43675 packets
845 packets (got 60427 ARP requests and 190479 ACKs), sent 43726 packets
244 packets (got 60409 ARP requests and 190613 ACKs), sent 43776 packets
446 packets (got 60553 ARP requests and 190749 ACKs), sent 43826 packets
845 packets (got 60633 ARP requests and 191010 ACKs), sent 43875 packets
557 packets (got 60713 ARP requests and 191347 ACKs), sent 43925 packets
988 packets (got 60793 ARP requests and 191638 ACKs), sent 43976 packets
206 packets (got 60864 ARP requests and 191780 ACKs), sent 44026 packets
400 packets (got 60927 ARP requests and 191908 ACKs), sent 44076 packets
595 packets (got 60988 ARP requests and 192039 ACKs), sent 44125 packets
806 packets (got 61055 ARP requests and 192181 ACKs), sent 44176 packets
917 packets (got 61124 ARP requests and 192319 ACKs), sent 44226 packets
220 packets (got 61189 ARP requests and 192455 ACKs), sent 44276 packets
427 packets (got 61253 ARP requests and 192594 ACKs), sent 44326 packets
625 packets (got 61313 ARP requests and 192729 ACKs), sent 44376 packets
138 packets (got 61394 ARP requests and 193044 ACKs), sent 44426 packets
791 packets (got 61473 ARP requests and 193367 ACKs), sent 44476 packets
179 packets (got 61552 ARP requests and 193626 ACKs), sent 44526 packets
396 packets (got 61622 ARP requests and 193770 ACKs), sent 44577 packets
605 packets (got 61692 ARP requests and 193908 ACKs), sent 44627 packets
818 packets (got 61761 ARP requests and 194050 ACKs), sent 44676 packets
pps)
hell

```

Aireplay sniffs ARP packets to harvest IVs while Airodump saves them to a capture, which Aircrack can analyze to identify the correct encryption key.

Lesson 12.4

Wireless Troubleshooting



Exam Objectives Covered

5.4 Given a scenario, troubleshoot common performance issues.

5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues.

Some environments and performance demands can complicate wireless network deployment. A variety of tools and techniques are available to assess Wi-Fi performance and ensure a highly available network for all users.

As you study this lesson, answer the following questions:

- What are sources of wireless interference?
- How are signal strength and noise measured and evaluated?
- What are the steps for diagnosing signal degradation, coverage, channel overlap, roaming, and disassociation issues?

Wireless Performance Assessment

Wireless troubleshooting can be broadly divided into Physical layer issues with signal strength or interference and configuration issues. This lesson focuses on issues that affect signal strength and performance, but always check that the security and authentication parameters are correctly configured before assuming you have a Physical layer connectivity problem.

As with cabled networks, you should distinguish between bit rate and throughput when measuring and assessing wireless performance against the specifications and limitations of a particular Wi-Fi standard:


- Bit rate is the amount of total amount of data transferred per second established at the Physical and Data Link layers. The nominal link bit rate is determined by standards support (Wi-Fi 5 or Wi-Fi 6, for instance), use of bonded channels, and optimizations, such as MU-MIMO. If the sender and receiver are far apart or subject to interference, a lower rate will be negotiated to make the link more reliable.
- Throughput is the amount of data that can be transferred at the Network layer, discarding overhead from layers 1 and 2. Often the term “goodput” is used to describe data transfer achieved at the Application layer (accounting for overhead from header fields and packet loss/retransmissions).

Radio frequency (RF) attenuation or free space path loss is the degradation of a signal as the distance between a radio transmitter and receiver increases. The strength of the signal decreases per the inverse-square rule. For example, doubling the distance decreases the signal strength by a factor of four. Meanwhile, interference sources collectively overlay a competing background signal, referred to as noise. These factors impose distance limitations on how far a client can be from an access point. Loss of power/signal strength is measured in dB units. For example, if transmit power is 14 dBm (~25 mW), antenna gain is 3 dBi, and free space loss over 30 meters is 70 dB, the received signal strength is approximately 0.000005 mW:

$$(14 \text{ dBm} + 3 \text{ dBi}) - 70 \text{ dB} = -53 \text{ dBm} = \sim 0.000005 \text{ mW}$$

If noise is -80 dBm, the signal-to-noise ratio (SNR) would be 27 dB. Most sites would aim for a margin of 20+ dB, so there is enough margin for signal loss through typical indoor environmental obstacles. For example, you might budget for 5 dB loss per internal wall.

Signal strength and noise can be measured by using a **Wi-Fi analyzer**. This type of software can be installed to a laptop or smartphone. It will record statistics for the AP that the client is currently associated with and detect any other access points in the vicinity. There are also dedicated Wi-Fi tester hardware devices.



Surveying Wi-Fi networks using inSSIDer. The chart shows which channels are active and the signal strength of different networks in each channel. (Screenshot used with permission from MetaGeek.)

Insufficient Wireless Coverage Issues

Insufficient wireless coverage refers to spots within a building with no or weak Wi-Fi signal. If a sufficient signal strength cannot be obtained and sources of interference cannot be mitigated, the only solution is to install an additional device to cover the gap. If you cannot extend the distribution system (cabled network) to support an additional access point, you will need to configure a wireless bridge or use a range extender.

Antenna Placement

Incorrect antenna placement could cause or exacerbate attenuation and interference problems. Use a site survey and heat map to determine the optimum position for APs and (if available) the direction in which to point adjustable antennae. Also, using an incorrect antenna type may adversely affect the signal strength at any given point. A unidirectional antenna is only suitable for point-to-point connections, not for general client access. The internal antennae built into APs may also be optimized to transmit and receive in some directions more than others. For example, an AP designed for ceiling mounting may produce a stronger signal in a cone directed downward from its central axis, whereas the signal from a desktop AP is likely to radiate in a doughnut-like pattern. Consult the documentation for your specific model of AP or use site survey software to produce a heat map.



Remember that some client devices might support a standard such as 802.11n, but only have a single band 2.4 GHz radio. They will not be able to join a 5 GHz network.

Antenna Cable Attenuation

Another source of attenuation is where the antenna is connected at some distance from the access point via coax cabling. Signal loss along this cable is referred to as antenna cable attenuation. LMR/HDF/CFD 200 cable has attenuation of about 0.6 dB/m (decibels per meter), while 400 cable improves that to about 0.22 dB/m. Connector loss is usually calculated as 0.15 dB.



If a device has removable antennae, check that these are screwed in firmly. A loose or disconnected antenna may reduce the range of the device or prevent connectivity altogether.

Effective Isotropic Radiated Power/Power Settings

Wireless devices have a configurable transmit power. **Effective Isotropic Radiated Power (EIRP)** is calculated as the sum of transmit power, antenna cable/connector loss, and antenna gain. For example, if you are configuring a point-to-point link with a directional antenna, you might derive the following value for EIRP:

$$\begin{aligned} &15 \text{ dBm (Transmit Power)} - 1 \text{ dB (Cable Loss)} \\ &+ 6 \text{ dBi (Gain)} = 20 \text{ dBm (100 mW)} \end{aligned}$$

The EIRP for each radio is reported through the access point or controller management software. EIRP must not exceed regulatory limits. Power limits are different for the 2.4 GHz and 5 GHz bands and for point to multipoint versus point to point operation modes.

Increasing transmit power is not usually an effective solution for improving wireless coverage. While an AP might have an EIRP of around 23 dBm, smartphone devices are more likely to be around 10 to 14 dBm. If the client detects a strong signal, it will set a high data rate. However, because the EIRP of the client radio is much lower, it fails to transmit a strong signal back to the AP. Because it is trying to use a high data rate, this results in excessive packet errors.

As a general rule of thumb, AP power should be two-thirds of the weakest client power. For example, if the weakest client can output 14 dBm, the AP should transmit at 9 to 10 dBm.


Channel Overlap Issues

Channel overlap refers to interference issues resulting from multiple access points that are all in range of one another and are configured to use similar wavelengths. There are two main types of channel interference:

- **Co-channel interference (CCI)**—Each Wi-Fi channel is a collision domain. Where multiple devices use the same channel, there is contention. The wireless devices must use CSMA/CA to find opportunities to transmit. CCI can be measured as a channel utilization percentage. Channel utilization can be measured from the access point or by using a Wi-Fi analyzer. As a design goal, a channel should exhibit no more than 50% utilization.
- **Adjacent channel interference (ACI)**—This occurs when access points are configured to use different but overlapping channels, such as 1 and 3 in the 2.4 GHz band. ACI slows down the CSMA/CA process and raises noise levels.

One of the design goals for a multi-AP site is to create clean basic service areas so that clients can select an AP with the strongest signal easily and the WLAN operates with a minimum of co-channel interference. At least 25 MHz spacing should be allowed to avoid channel overlap. In practice, therefore, no more than three nearby APs using the 2.4 GHz band can have non-overlapping channels. This could be implemented, for example, by selecting channel 1 for AP1, channel 6 for AP2, and channel 11 for AP3. When you are using the 5 GHz band for 802.11a or Wi-Fi 4/5/6, more non-overlapping channels are available.

In a complex environment, it may be necessary to adjust the power level used by an AP on a given channel. Using the maximum available power on an AP can result in it interfering with other “cells” and to situations where a client can “hear” the AP but cannot “talk” to it because it lacks sufficient power.



Checking power levels on a wireless station using Intel's PROSet Wi-Fi configuration utility.
(Screenshot courtesy of Intel Corp.)

Interference Issues

If a device is within the supported range but the signal is weak or you cannot get a connection, there is likely to be interference. Apart from CCI and ACI, there are several other sources of interference to consider:

- **Reflection/bounce (multipath interference)**—Mirrors or shiny surfaces cause signals to reflect, meaning that a variable delay is introduced. This causes packets to be lost and consequently the data rate to drop.



The Wi-Fi 4/5/6 standards actually use bounce (multipath) as a means of optimizing throughput and range via MIMO.

- **Refraction**—Glass or water can cause radio waves to bend and take a different path to the receiver. This can also cause the data rate to drop.
- **Absorption**—This refers to the degree to which walls, windows, and people will reduce signal strength (some of the radio wave's energy is lost as heat when passing through construction materials or human bodies). An internal wall might "cost" 3 to 15 dB, depending on the material used (concrete being the most effective absorber). The 2.4 GHz frequency has better penetration than the 5 GHz one, given the same power output. To minimize absorption from office furniture (and people), use ceiling-mounted APs.
- **Electromagnetic interference (EMI)**—Interference from a powerful radio or electromagnetic source working in the same frequency band, such as a Bluetooth device, cordless phone, or microwave oven.

EMI can be detected by using a **spectrum analyzer**. Unlike a Wi-Fi analyzer, a spectrum analyzer must use a special radio receiver—Wi-Fi adapters filter out anything that isn't a Wi-Fi signal. They are usually supplied as handheld units with a directional antenna, so that the exact location of the interference can be pinpointed. A 6 dB change in the level of a particular source represents a halving or doubling of the distance between the analyzer and the source of the RF source.



Also consider that signal problems could be a result of someone trying to attack the network by jamming the legitimate AP and making clients connect to a rogue AP.

Roaming and Client Disassociation Issues

Roaming means that wireless clients can remain connected to the same network ESSID while moving around within an extended service area (ESA). An ESA is created by connecting APs via a wired network and configuring them with the same ESSID and security parameters. The access points are configured with different channels so that where the basic service area of each access point overlaps, there is no interference.

Roaming Misconfiguration Issues

In order to enable seamless roaming, the cells served by each AP need to overlap to some extent. This is one of the trickiest elements of site design to get right, as client behaviors and capabilities for roaming can vary widely. The main issues with roaming are the following:

- Sticky clients that do not identify signal issues and do not reassociate with a different AP that could provide a better connection.
- Flapping clients that switch repeatedly between access points.
- Clients that do not support roaming standards (802.11k, 802.11r, and 802.11v) and so experience service interruptions due to having to reauthenticate or associate too slowly with the new AP. 802.11r assists with reauthentication, support for 802.11k can mitigate sticky and flapping client issues, as it transmits information about the wireless topology to the client, and 802.11v can “push” a client toward a less congested access point.
- Inconsistent service areas for 2.4 GHz and 5 GHz. 2.4 GHz supports longer ranges than 5 GHz, and this can cause it to “attract” more clients. Typically, a 2.4 GHz BSS is configured with a lower transmit power than the equivalent 5 GHz BSS.

Issues with roaming can be identified by analyzing access point association times for client devices. A WLAN controller will be able to track client mobility, showing each access point and the time that the client associated with it.

Client Disassociation Issues

In the normal course of operations, an access point and client use management frames to control connections. The access point broadcasts a beacon frame to advertise service capabilities. Clients can choose to first authenticate and then associate to an access point when they move into range of the beacon. The client or access point can use **disassociation** and/or deauthentication frames to notify the other party that it has ended a connection. A legitimate client might disassociate but not deauthenticate because it is roaming between access points in an extended service area. A client might “flap” between two access points, causing numerous disassociations and reassociations. Investigate the access point or controller event log to identify the cause of disassociations.

If clients are disassociated unexpectedly and there is no roaming, interference, or standards support issue, you should consider the possibility of a malicious attack. A disassociation attack exploits the lack of encryption in management frame traffic to send spoofed frames. One type of disassociation attack injects management frames that spoof the MAC address of a single victim station in a disassociation notification, causing it to be disconnected from the network. Another variant of the attack broadcasts spoofed frames to disconnect all stations. Frames can be spoofed to send either disassociation or deauthentication notifications.

Disassociation/deauthentication attacks may be used to perform a denial of service attack against the wireless infrastructure or to exploit disconnected stations to try to force reconnection to a rogue or evil twin access point. Disassociation/deauthentication attacks might also be used in conjunction with a replay attack aimed at recovering the network key.

Overcapacity Issues

Overcapacity (or device saturation) occurs when too many client devices connect to the same AP. The maximum number of clients that an AP can support varies, depending on the Wi-Fi standard used and the type of network traffic generated. For example, web browsing will typically place a lighter load on the network than local client-server traffic or is likely at least to move any bottleneck further upstream to the WAN, rather than the wireless network. While individual circumstances must be considered, a maximum of 30 clients per AP is generally accepted as a rule of thumb. In designing the network, enough APs should be provided in appropriate locations to support the expected client density at this ratio. APs can usually be configured to enforce a maximum number of connections, so that additional clients will connect to the next nearest AP.

Even with a relatively low number of clients, the wireless network can suffer from bandwidth saturation. Since wireless is a broadcast medium, the available bandwidth is shared between all clients. Thus, if one client is a bandwidth hog, others may find it difficult to maintain a reliable connection. In an enterprise Wi-Fi solution, a controller will normally provide reporting tools to diagnose bandwidth issues and to report on which clients are consuming the most bandwidth. It could also report on wireless channel utilization and configure APs and clients to reassign channels dynamically to reduce overutilization. If a traffic shaper is deployed, it may work automatically to throttle bandwidth to overactive nodes.

Module 12

Summary

You should be able to install, configure, and troubleshoot appropriate wireless standards and technologies.

Guidelines for Deploying and Troubleshooting Wireless Networks

Follow these guidelines to deploy and troubleshoot wireless networks:

- Create a list of requirements for your network so that you can work toward meeting them. These requirements may include how many users need to connect, the physical area it will need to cover, external connections, and more.
- Consider the devices you will need and any compatibility requirements they have, in terms of Wi-Fi standards support, such as 802.11a, b, g or Wi-Fi 4 (n), 5 (ac), 6 (ax).
- Obtain a scale drawing of the building and a Wi-Fi analyzer to use to perform a site survey and generate heat maps of signal strength and channel utilization.
- Determine the range of the AP for the wireless technology you have chosen. This will help you to better determine how many APs you will need to ensure adequate coverage for the space.
- Balance the number of users who will have access to the AP, and ensure that the AP can cover all employees in the range of the AP. More employees in a given area means more APs.
- Tour the area in the range of the AP, and check to see if there are any devices that will interfere with the wireless network. This can include devices such as microwave ovens, Bluetooth-enabled devices, or an existing wireless network—whether from a community network, a neighboring building, or another floor of your company's building. These devices or networks can possibly interfere with your new implementation.
- Ensure that there are no obstacles in the path of the AP, such as doors, closed windows, walls, and furniture, that the wireless signal will need to pass through on its way to a client. If there are too many obstacles in the path, adjust the placement of your AP accordingly.
- Install the APs. The specific steps for installing the AP will vary by vendor, but the common steps may include the following:
 - Connecting the AP to the cabled network (distribution system) via a switch.
 - Setting the SSID/ESSID and an 802.11 beacon.
 - Configuring frequency bands and channel layout within each frequency band.
 - Adjusting transmit power to reduce channel overlap and using either transmit power or band steering to optimize frequency usage by clients.
 - Configuring the appropriate encryption and authentication schemes, such as WPA2/3 personal versus WPA2/3 enterprise. If appropriate, configure RADIUS or TACACS+ support for enterprise authentication.

- Test to ensure that the installation is appropriately sized, secure, and operational. Make sure these tests are done under real-world conditions so that you have an accurate test.
- Perform period site surveys to check RSSI at key locations and compare it to previous performance levels from previous site surveys.
- Document the steps and establish a baseline for future installations.

Module 13

Comparing Remote Access Methods

Module Introduction

Most local networks require some sort of external connection, whether to the global Internet or within an enterprise WAN. These long-distance communications are typically facilitated by service provider links. Supporting WAN and Internet access effectively is an essential competency to learn.

User services and network management often require the creation of various types of remote access, including virtual private networks (VPNs). While remote access makes networks more usable and accessible, it also broadens the attack surface. You must understand the implications of different remote access models and protocols so that you can support their secure use.

Module Objectives

In this module, you will do the following:

- Summarize WAN provider and Internet access types.
- Compare and contrast VPN topologies and protocols.
- Explain remote host access and management methods.

Lesson 13.1

WAN and Internet Connectivity

Understanding the various WAN connectivity devices and methods will help you support Internet connectivity and the configuration of enterprise WANs. You will need to understand the capabilities and limitations of WAN provider links to choose the one best suited for your network.

As you study this lesson, answer the following questions:

- How are private homes and offices connected to public telecommunications networks?
- What connection speeds should you expect from various types of Internet access methods?

Wide Area Networks and the OSI Model

Wide area network (WAN) technologies support data communications over greater distances than LANs. The term “enterprise WAN” is used to describe a WAN that is used and controlled by a single organization. However, even though an enterprise may control its WAN, it rarely owns all the infrastructure that supports it. Long-distance communications usually involve the use of public networks. Public networks are owned by telecommunications (telco) companies and provide WAN services to businesses and households. Organizations often choose to use public networks, as the cost is far less than implementing a private solution. Service providers often have rights of access to locations that are not available to other organizations, such as under roads.

As with a LAN, the WAN Physical layer describes the media type and interface specifications. Where the provider link is a copper cable, some type of modem is usually used, rather than a switch. A modem performs modulation of outgoing signals and demodulation of incoming data, working only at the Physical layer of the OSI model. Modulation means transforming an electromagnetic wave to represent information, such as using the amplitude (height) of the wave to represent bits. Legacy modems perform digital to analog modulation for transmission over voice lines. An analog (or dial-up) modem only supports low bandwidths (up to 56 Kbps). Digital modems perform a different type of modulation to transform digital signals received as Ethernet frames for transmission over the WAN media. Digital modem types include data service units (DSUs) for leased lines, digital subscriber line (DSL) modems, cable modems, and satellite modems.

At the Data Link layer, WANs often use simpler protocols than Ethernet LANs as the links are more likely to be point to point and do not need much complexity. That said, Ethernet is increasingly being deployed for end-to-end connectivity over WANs.

At the Network layer, the customer and provider site are addressed using the Internet Protocol (IP). A customer edge (CE) router connects to a provider edge (PE) via the underlying Link layer interface. The provider allocates public IPv4/IPv6 addresses or address ranges to the customer.

Internet Access Types

Establishing a WAN provider link means terminating the access provider's cabling at some point in your premises, and then attaching modem and routing equipment to that line. The service-related entry point at which the access provider's network terminates is called the demarcation point (or demarc for short) or minimum point of entry (MPOE). The demarc point represents the end of the telco's responsibility for maintaining that part of the cabling. Any cable problems arising from the other side of the demarc point are the responsibility of the customer.

Modems and routers or other access equipment that are provided or leased by the customer and installed at their site are referred to as customer premises equipment (CPE). Some providers may take on responsibility for faults that arise in CPE, depending on the contract and installation circumstances.

The demarc and CPE should be installed to a secure location within the premises, with access controls to restrict the area to authorized staff. This location is referred to as entrance facilities in TIA/EIA structured cabling standards.

T-Carrier

The T-carrier system enabled voice traffic to be digitized for transport around the core of the telecommunications network. It also enabled other types of digital data to be transported and could be provisioned directly to subscribers as a leased line. T-carrier is based on time division multiplexing (TDM). The protocol assigns each circuit (or channel) a time slot. Each 64 Kbps channel provides enough bandwidth for a digitized voice call.

A single 64 Kbps channel is known as a DS0 or narrowband link. For leased line data services, however, the foundation level of T-carrier is the DS1 or T1 digital signal circuit. This service comprises 24 channels multiplexed into a single 1.544 Mbps full-duplex digital connection that can be used for voice and data. The T1 lines themselves can be multiplexed to provide even more bandwidth.

At the Data Link layer, T1 leased lines typically use either High-Level Data Link Control (HDLC) or Point-to-Point Protocol (PPP).

Digital Subscriber Line

Digital subscriber line (DSL) is a technology for transferring data over voice-grade telephone lines, often referred to as the local loop. A DSL modem is installed as CPE, typically as a multifunction "wireless router," where the RJ11 WAN port connects to the provider's phone jack. DSL modems can also be supplied as separate appliances or plug-in cards for routers. A filter (splitter) must be installed on each phone point to prevent noise from affecting either voice calls or the DSL link.

There are various "flavors" of DSL, notably asymmetrical and symmetrical types:

- Asymmetrical DSL (ADSL) provides a fast downlink but a slow uplink. There are various iterations of ADSL, with the latest (ADSL2+) offering downlink rates up to about 24 Mbps and uplink rates of 1.25 Mbps or 2.5 Mbps.
- Symmetric versions of DSL offer the same uplink and downlink speeds. These are of more use to businesses and for branch office links, where more data is transferred upstream than with normal Internet use.

Cable Internet

A cable Internet connection is usually available along with Cable Access TV (CATV). These networks are sometimes described as hybrid fiber coax (HFC) because they combine a fiber optic core network with coax links to CPE, but are more simply just described as cable broadband.

Installation of a cable modem follows the same general principles as for a DSL modem. An Ethernet cable connects the cable modem to the customer's router, and a short segment of coax connects the WAN port to the provider network. More coax then links all the premises in a street with a Cable Modem Termination System (CMTS), which routes data traffic via the fiber backbone to the ISP's point of presence (PoP) and from there to the Internet. Cable based on the Data Over Cable Service Interface Specification (DOCSIS) supports downlink speeds of up to 38 Mbps (North America) or 50 Mbps (Europe) and uplinks of up to 27 Mbps. DOCSIS version 3 allows the use of multiplexed channels to achieve higher bandwidth.

Fiber to the Curb and Fiber to the Premises

The major obstacle to providing WAN access that can approach LAN performance is bandwidth in the last mile, where the copper wiring infrastructure is generally not good. The projects to update this wiring to use fiber optic links are referred to by the umbrella term Fiber to the X (FTTx).

Fiber to the Curb and VDSL


A fiber to the curb (FTTC) solution retains some sort of copper wiring to the customer premises while extending the fiber link from the point of presence to a communications cabinet servicing multiple subscribers. The service providers with their roots in telephone networks use very high-speed DSL (VDSL) to support FTTC. VDSL achieves higher bit rates than other DSL types at the expense of range. It allows for both symmetric and asymmetric modes. Over 300 m (1,000 feet), an asymmetric link supports 52 Mbps downstream and 6 Mbps upstream, while a symmetric link supports 26 Mbps in both directions. VDSL2 specifies a very short range (100 m/300 feet) rate of 100 Mbps (bidirectional).



The modem type must match the service. An ADSL-only modem cannot be used to access a VDSL service, for instance.

Fiber to the Premises and Optical Network Terminals

A fiber to the premises (FTTP) Internet connection means that the service provider's fiber optic cable is run all the way to the customer's building. This full fiber connection type is implemented as a passive optical network (PON). In a PON, a single fiber cable is run from an optical line terminal (OLT) to a splitter. The splitter directs each subscriber's traffic over a shorter length of fiber to an optical network terminal (ONT) installed at the customer's premises. The ONT converts the optical signal to an electrical one. The ONT is connected to the customer's router using an RJ45 Ethernet patch cord.



Optical network terminal—the PON port terminates the external fiber cable, and the LAN ports connect to local routers or computers over RJ45 patch cords. (Image by artush © 123RF.com.)

Lesson 13.2

Virtual Private Networks



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.5 Compare and contrast network access and management.

With today's mobile workforce, most networks have to support connections by remote employees, contractors, and customers to their network resources. These remote connections often make use of untrusted public networks, such as the Internet. To counter the security risks associated with public networks, organizations implement a virtual private network (VPN) over the public network to ensure secure communications. Consequently, understanding how to implement secure remote access VPN protocols will be a major part of your job as a network professional.

As you study this lesson, answer the following questions:

- Which protocols support VPNs?
- What are the ways a VPN can be implemented?
- What is the difference between a full tunnel and split tunnel?

Remote Access Considerations

Remote network access means that the user's device does not make a direct cabled or wireless connection to the network. The connection occurs over or through an intermediate network, usually a public WAN. Historically, remote network access might have used analog modems connecting over the telephone system. These days, most remote network access is implemented as a virtual private network (VPN), running over the Internet.

Given that, administering remote access involves essentially the same tasks as administering the local network. Only authorized users who have successfully authenticated should be allowed access to local network resources and communication channels. Additional complexity comes about because it can be more difficult to ensure the security of remote workstations and servers, and there is greater opportunity for remote logins to be exploited.

The creation of a remote access server (RAS) should be accompanied by documentation describing the uses of the service, security risks and countermeasures, and authorized users of the service. There should also be authorization to run the service from the network manager. The remote access policy should then implement the measures identified through compiling the documentation. Typical policy restrictions would be:

- Restricting access to defined users or groups.
- Restricting access to defined times of day or particular days of the week.

- Restricting privileges on the local network (ideally, remote users would only be permitted access to a clearly defined part of the network).
- Logging and auditing access logons and attempted logons.

In addition to this, a management plan should ensure that RASs and other hardware are kept up to date with the latest software or firmware updates. Administrative access to the devices should also be secured, using strong authentication.

Tunneling Protocols

Most modern remote network solutions use Internet access infrastructure to implement a **virtual private network (VPN)**. This requires a protocol that can create a secure tunnel for private communications through the Internet. Tunneling is where the hosts are on the same logical network but connected via different physical networks. The tunnel encapsulates the packet for the local network within a public network packet. Typically, the local network packet is encrypted. When the packet is delivered, the remote access server strips the public packet headers, extracts and decrypts the local packet, and forwards it over the local network.

Point-to-Point Protocol

The **Point-to-Point Protocol (PPP)** is an encapsulation protocol that works at the Data Link layer (layer 2). PPP is used to encapsulate IP packets for transmission over serial digital lines. PPP has no security mechanisms, so must be used with other protocols to provision a secure tunnel.

Generic Routing Encapsulation

Where PPP works at layer 2, **Generic Routing Encapsulation (GRE)** works at layer 3. A GRE packet can itself encapsulate an IP packet (or most other Network layer protocol types) as its payload. The “outer” GRE packet is assigned protocol number 47 and has its own IP source and header address fields. The GRE packet is then itself encapsulated in a layer 2 frame for transmission to the next hop router. Each intermediate router inspects only the outer GRE header to determine the forwarding destination. At the final destination, the receiving router de-encapsulates the GRE packet to extract the inner IP payload and forwards that inner packet to its destination. GRE does not have any mechanisms for authenticating users or devices and so is often used with other protocols in a VPN solution.

IP Security

Internet Protocol Security (IPSec) also operates at the Network layer of the OSI model to authenticate hosts and encrypt packets. IPSec is used with other protocols to provide connection security, and it is increasingly used as a standalone VPN protocol.

Transport Layer Security

Transport Layer Security (TLS) over TCP or datagram TLS (DTLS) over UDP can be used to encapsulate frames or IP packets. The main drawback is that as TLS already operates at the Session layer, the headers from the inner and outer packets add up to a significant overhead.

Internet Protocol Security

Internet Protocol Security (IPSec) can be used to secure IPv4 and/or IPv6 communications on local networks and as a remote access VPN protocol. IPSec operates at the Network layer of the OSI model (layer 3). This means that it can be implemented without having to configure specific application support and that it incurs less packet overhead.

There are two core protocols in IPSec, which can be applied singly or together, depending on the policy:

- **Authentication Header (AH)**—Performs a cryptographic hash on the whole packet, including the IP header, plus a shared secret key (known only to the communicating hosts), and adds this value in its header as an Integrity Check Value (ICV). The recipient performs the same function on the packet and key and should derive the same value to confirm that the packet has not been modified. The payload is not encrypted so this protocol does not provide confidentiality.
- **Encapsulating Security Payload (ESP)**—Can be used to encrypt the packet rather than simply calculating an ICV. ESP attaches three fields to the packet: a header, a trailer (providing padding for the cryptographic function), and an Integrity Check Value. Unlike AH, ESP excludes the IP header when calculating the ICV.



With ESP, algorithms for both confidentiality (symmetric cipher) and authentication/integrity (hash function) are usually applied together. It is possible to use one or the other, however.

IPSec can be used in two modes:

- **Transport mode**—This mode is used to secure communications between hosts on a private network (an end-to-end implementation). When ESP is applied in transport mode, the IP header for each packet is not encrypted, just the payload data. If AH is used in transport mode, it can provide integrity for the IP header.




IPSec datagram using AH and ESP in transport mode.

- **Tunnel mode**—This mode is used for communications between VPN gateways across an insecure network (creating a VPN). This is also referred to as a router implementation. With ESP, the whole IP packet (header and payload) is encrypted and encapsulated as a datagram with a new IP header. AH has no real use case in tunnel mode, as confidentiality will usually be required.



IPSec datagram using ESP in tunnel mode.



Configuring an IPsec tunnel in the OPNsense security appliance.
(Screenshot used with permission from OPNsense.)




The principles underlying IPsec are the same for IPv4 and IPv6, but the header formats are different. IPsec makes use of extension headers in IPv6 while in IPv4, ESP and AH are allocated new IP protocol numbers (50 and 51), and either modify the original IP header or encapsulate the original packet, depending on whether transport or tunnel mode is used.

Internet Key Exchange

Each host or router that uses IPsec must be assigned a policy. An IPsec policy sets the authentication mechanism and also the use of AH/ESP and transport or tunnel mode for a connection between two peers.

IPsec's encryption and hashing functions depend on a shared secret. The secret must be communicated to both peers, and the peers must perform mutual authentication to confirm one another's identity. The **Internet Key Exchange (IKE)** protocol implements an authentication method, selects which cryptographic ciphers are mutually supported by both peers, and performs key exchange. The set of properties is referred to as a security association (SA).



*Configuring IKE in the OPNsense security appliance.
(Screenshot used with permission from OPNsense.)*

IKE negotiations take place over two phases:

1. Phase I establishes the identity of the two peers and performs key agreement using the Diffie-Hellman algorithm to create a secure channel. Two methods of authenticating peers are commonly used:
 - **Digital certificates**—Are issued to each peer by a mutually trusted certificate authority to identify one another.
 - **Pre-shared key (group authentication)**—Is when the same passphrase is configured on both peers.
2. Phase II uses the secure channel created in Phase I to establish which ciphers and key sizes will be used with AH and/or ESP in the IPsec session.

There are two versions of IKE. Version 1 was designed for site-to-site and host-to-host topologies and requires a supporting protocol to implement remote access VPNs. IKEv2 has some additional features that have made the protocol popular for use as a stand-alone remote access client-to-site VPN solution. The main changes are the following:

- Supports EAP authentication methods, allowing, for example, user authentication against a RADIUS server.
- Provides a simple setup mode that reduces bandwidth without compromising security.
- Allows network address translation (NAT) traversal and MOBIKE multihoming. NAT traversal makes it easier to configure a tunnel allowed by a home router/firewall. Multihoming means that a smartphone client with Wi-Fi and cellular interfaces can keep the IPsec connection alive when switching between them.

Client-to-Site VPNs

A VPN can be implemented in several topologies. In a client-to-site or remote access topology, the VPN client connects over the public network to a VPN gateway (a VPN-enabled router) positioned on the edge of the local network (typically the VPN access server will be in a screened subnet). Client-to-site is the “telecommuter” model, allowing homeworkers and employees working in the field to connect to the corporate network.


Client-to-site VPNs can be configured using a number of protocols. An SSL/TLS VPN solution uses certificates to establish the secure tunnel. One example is Microsoft’s Secure Socket Tunneling Protocol (SSTP). Cisco’s Layer 2 Tunneling Protocol (L2TP) is also widely used, in conjunction with IPsec. All these solutions require client software to operate. Most VPN solutions use EAP and AAA/RADIUS architecture to authenticate client devices and users.



Microsoft’s Point-to-Point Tunneling Protocol (PPTP) was once very widely used but has too many security flaws to be deployed safely.


When a client connected to a remote access VPN tries to access other sites on the Internet, there are two ways to manage the connection:

- **Split tunnel**—The client accesses the Internet directly using its ISP-manged IP configuration, routers, and DNS servers.



Split tunnel VPN traffic flow. (Images © 123RF.com.)

- **Full tunnel**—Internet access is mediated by the corporate network, which will alter the client’s IP address and DNS servers and may use a proxy.



Full tunnel VPN traffic flow. (Images © 123RF.com.)

Full tunnel offers better security, but the network address translations and DNS operations required may cause problems with some websites, especially cloud services. It also means more data is channeled over the link, and the connection can exhibit higher latency.


Clientless VPNs

Microsoft's Remote Desktop Protocol (RDP) can be used to access a physical machine on a one-to-one basis. Alternatively, the site can operate a remote desktop gateway that facilitates access to virtual desktops or individual apps running on the network servers (docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/welcome-to-rds). Similar services are provided by Citrix's products (citrix.com/products).

Traditionally, remote desktop products and client-to-site VPNs require a client app that implements the protocols and authentication methods supported by the remote desktop/VPN gateway. The canvas element introduced in HTML5 allows a browser to draw and update a desktop with relatively little lag. It can also handle audio. This allows ordinary browser software to connect to a remote desktop or to a VPN portal that publishes a number of web applications. This is referred to as an **HTML5 VPN or clientless VPN** (guacamole.apache.org). This solution also uses a protocol called WebSockets, which enables bidirectional messages to be sent between the server and client without requiring the overhead of separate HTTP requests.

Site-to-Site VPNs

A VPN can also be deployed in a site-to-site model to connect two or more private networks. Where remote access VPN connections are typically initiated by the client, a site-to-site VPN is configured to operate automatically. The gateways exchange security information using whichever protocol the VPN is based on. This establishes a trust relationship between the gateways and sets up a secure connection through which to tunnel data. Hosts at each site do not need to be configured with any information about the VPN. The routing infrastructure at each site determines whether to deliver traffic locally or send it over the VPN tunnel. This is also referred to as compulsory tunneling. Compulsory tunnels can be in place permanently (static), or they can be put in place based on the data or client type (dynamic).



Site-to-site VPN. (Images © 123rf.com.)



VPNs are not always established over the public Internet. A WAN service provider can implement VPNs via its network. The provider can use VLAN-like technology to isolate a customer's data from other traffic. This is a common model for site-to-site VPNs.



While VPNs are being covered here as part of remote access, they can be just as usefully deployed on local networks as a type of network segmentation. For example, the department for product development might need to provide secure communications with SCADA workstations in an industrial internet of things (IIoT) segment.

Lesson 13.3

Remote Management



Exam Objectives Covered

- 1.4 Explain common networking ports, protocols, services, and traffic types.
- 3.5 Compare and contrast network access and management.

A remote management tool allows you to configure servers and devices over the network. Having to perform configuration and troubleshooting activity at a local console would be incredibly time consuming. Efficient network administration depends upon remote access tools. It is imperative to configure these tools securely, however.

As you study this lesson, answer the following questions:

- What protocols support secure remote access?
- What is the difference between in-band and out-of-band management?
- What is the role of APIs in remote management, and what security issues do they raise?

Remote Host Access

A remote access VPN refers to extending local network access over an intermediate public network so that a remote computer is effectively joined to the local network. Remote access can also refer to remote host access, where a user operates a computer or configures a network appliance without having to use a local terminal. This type of remote host access can be implemented within a local network or over a public network. It can be used for a variety of purposes:

- Remote configuration of network appliances. Most of these appliances are headless (they do not have a video monitor or input devices), and remote connections are the only practical configuration option. This type of connection is typically implemented using Secure Shell (SSH).
- Remote desktop connections either allow an administrator to configure a server or a user to operate a computer remotely. Where remote desktop protocols provide GUI access, other protocols can be used for terminal-only access.
- Remote desktop gateways allow user access to networked apps. A gateway can also be used to connect a user to a virtual desktop, where a client OS and applications software are provisioned as a virtual appliance. Alternatively, a remote desktop gateway is a means of implementing a clientless VPN.

Secure Shell


The name “terminal” comes from the early days of computing where configuration was performed by a teletype (TTY) device. The TTY is the terminal or endpoint for communication between the computer and the user. The TTY handles text input and output between the user and the shell, or command environment. Where the terminal accepts input and displays output, the shell performs the actual processing.

A terminal emulator is any kind of software that replicates this TTY input/output function. A given terminal emulator application might support connections to multiple types of shell. A remote terminal emulator allows you to connect to the shell of a different host over the network.

Secure Shell (SSH) is the principal means of obtaining secure remote access to UNIX and Linux servers and to most types of network appliances (switches, routers, and firewalls). As well as terminal emulation, SSH can be used as the secure file transfer protocol (SFTP). There are numerous commercial and open source SSH servers and terminal emulation clients available for all the major NOS platforms (UNIX, Linux, Windows, and macOS). The most widely used is OpenSSH (openssh.com). An SSH server listens on TCP port 22 by default.

SSH Host Key

An SSH server is identified by a public/private key pair, referred to as the host key. A mapping of host names to public keys can be kept manually by each SSH client, or there are various enterprise software products designed for SSH key management.



Confirming the SSH server's host key using the SSH client. (Screenshot courtesy of Microsoft.)



The host key must be changed if any compromise of the host is suspected. If an attacker has obtained the private key of a server or appliance, they can masquerade as that server or appliance and perform a spoofing attack, usually with a view to obtaining other network credentials. You might also change the key to use a longer bit strength.

SSH Client Authentication

The server's host key is used to set up a secure channel to use for the client to submit authentication credentials. SSH allows various methods for the client to authenticate to the SSH server. Each of these methods can be enabled or disabled as required on the server:

- **Username/password**—The client submits credentials that are verified by the SSH server either against a local user database or using a network authentication server.
- **Public key authentication**—Each remote user's public key is added to a list of keys authorized for each local account on the SSH server.
- **Kerberos**—The client submits the Kerberos credentials (a Ticket Granting Ticket) obtained when the user logged onto the workstation to the server using the Generic Security Services Application Program Interface (GSSAPI). The SSH server contacts the Ticket Granting Service (in a Windows environment, this will be a domain controller) to validate the credential.



Managing valid client public keys is a critical security task. Many recent attacks on web servers have exploited poor key management. If a user's private key is compromised, delete the public key from the appliance, then regenerate the key pair on the user's (remediated) client device and copy the public key to the SSH server. Always delete public keys if the user's access permissions have been revoked.


Secure Shell Commands

SSH features a rich command set, fully documented at the OpenSSH website (openssh.com/manual.html). Some of the most important commands are the following:

- `ssh`—Start the SSH Daemon (server). Parameters such as the host's certificate file, port to listen on, and logging options can be set via switches or in a configuration file.
- `ssh-keygen`—Create a key pair to use to access servers. The private key must be stored securely on your local computer. The public key must be copied to the server. You can use the `ssh-copy-id` command to do this, or you can copy the file manually.
- `ssh-agent`—Configure a service to use to store the keys used to access multiple hosts. The agent stores the private key for each public key securely and reduces the number of times use of a private key has to be confirmed with a passphrase. This provides a single sign-on (SSO) mechanism for multiple SSH servers. The `ssh-add` command is used to add a key to the agent.
- `ssh Host`—Use the SSH client to connect to the server running at *Host*. *Host* can be an FQDN or IP address. You can also create a client configuration file.
- `ssh Username@Host`—Use the SSH client to connect to the server running at *Host* with a different *Username*.
- `ssh Host "Command or Script"`—Use the SSH client to execute a command or script on the remote server running at *Host* without starting a shell.
- `scp Username@Host:RemoteFile/Local/Destination`—A file transfer client with remote copy/rcp-like command interface.
- `sftp`—A file transfer client with FTP-like command interface.

Telnet

Telnet is both a protocol and a terminal emulation software tool that transmits shell commands and output between a client and the remote host. In order to support Telnet access, the remote computer must run a service known as the Telnet Daemon. The Telnet Daemon listens on TCP port 23 by default.



```

mail.classroom.local - PuTTY
220 mail.classroom.local ESMT
helo localhost
250 Hello.
mail from:<administrator@web.local>
250 OK
rcpt to:<administrator@classroom.local>
250 OK
data
354 OK, send.
from: Tech Support <administrator@web.local>
to: Hostmaster <administrator@classroom.local>
subject: Virus infection
mime-version: 1.0;
content-type: text/html;

<html>
<body>
<p>Viruses have been detected on your hosted server. Visit the <a href="http://www.notagoodidea.net">Hosting Services Portal</a> and enter your password to scan and remove them.</p>
</body>
</html>
.
250 Queued (199.078 seconds)
  
```

PuTTY Telnet client. (Screenshot courtesy of PuTTY.)

A Telnet interface can be password protected, but the password and other communications are not encrypted and therefore could be vulnerable to packet sniffing and replay. Historically, Telnet provided a simple means to configure switch and router equipment, but only secure access methods should be used for these tasks now. Ensure that the Telnet service is uninstalled or disabled, and block access to port 23.



If use of Telnet to manage legacy systems is unavoidable, these legacy systems must be deployed to a secure segment.

Remote Desktop Protocol

Telnet and SSH provide terminal emulation for command line shells. This is sufficient for most administrative tasks, but where users want to connect to a desktop, they usually prefer to work with an interface that can be controlled with a mouse. A graphical user interface (GUI) remote administration tool sends screen and audio data from the remote host to the client and transfers mouse and keyboard input from the client to the remote host. **Remote Desktop Protocol (RDP)** is Microsoft's protocol for operating remote GUI connections to a Windows machine. RDP uses TCP port 3389.

The administrator can specify permissions to connect to the server via RDP and can configure encryption on the connection. RDP should also be configured to require Network Level Authentication (NLA). NLA protects the RDP server against denial of service attacks. Without NLA, the system configures a desktop before the user logs on. A malicious user can create multiple pending connections to try to crash the system. NLA authenticates the user before committing any resources to the session.

If Remote Desktop is used to connect to a server that has been compromised by malware, the credentials of the user account used to make the connection become highly vulnerable. RDP Restricted Admin (RDPRA) Mode and Remote Credential Guard are means of mitigating this risk. You can read more about these technologies at docs.microsoft.com/en-us/windows/security/identity-protection/remote-credential-guard.



*To protect privileged account credentials, it is important not to initiate remote connections to servers and appliances from a low-trust workstation. A **secure administrative workstation (SAW)** is a computer with a very low attack surface running the minimum possible apps used solely for remote management.*

There are several popular alternatives to Remote Desktop. Most support remote access to platforms other than Windows (macOS and iOS, Linux, Chrome OS, and Android, for instance). Examples include TeamViewer (teamviewer.com/en) and Virtual Network Computing (VNC), which is implemented by several different providers (notably realvnc.com/en).

RDP is mainly used for the remote administration of a Windows server or client, but another function is to publish software applications on a server, rather than installing them locally on each client (application virtualization). A site can operate a remote desktop gateway that facilitates access to virtual desktops or individual apps running on the network servers (docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/welcome-to-rds). Similar services are provided by Citrix's products (citrix.com/products).

Console Connections and Out-of-Bound Management

Some network appliances, such as unmanaged switches, do not offer any configuration options or interface. You just have to plug them in, and they operate automatically. Managed switches and appliances, such as routers, firewalls, switches, and access points, support more complex functions and can be configured and monitored over several interfaces. The functions of a managed appliance can be accessed via one of the device's management interfaces. An appliance may support the following interfaces:

- **Console port**—This requires connecting a laptop running terminal emulator software to the switch or router via a separate physical interface using a special console (or rollover) cable. The terminal emulator can then be used to start a command line interface (CLI).
- **AUX port**—This port is designed to connect to an analog modem and provide remote access over a dial-up link. Once the AUX port is enabled and configured, the modem can be connected to it by using an RS-232 serial cable, a specially wired RJ45 rollover cable and terminal adapter (RJ45 to DB9), or a management cable (RJ45 to DB9). Configure the modem with appropriate serial link settings (refer to the vendor guide), connect it to an appropriate telephone line, and allocate an extension number. A remote host can connect to the appliance CLI by using a terminal emulation program such as HyperTerminal or PuTTY.
- **Management port**—This means configuring a virtual network interface and IP address on the device to use for management functions and connecting to it via one of the normal Ethernet ports. The port must be enabled for this function (some appliances come with a dedicated management port). Using Telnet (insecure) or Secure Shell (SSH) to connect to a CLI remotely over the management interface in this way is referred to as a virtual terminal.



USB and RJ45 type console ports plus AUX and other management interfaces on a router.
(Image © 123RF.com.)

Management methods can be described as either in-band or **out-of-band (OOB)**. An in-band management link is one that shares traffic with other communications on the “production” network. The console port is a physically out-of-band management method; the link is limited to the attached device. When you are using a browser-based management interface or a virtual terminal, the link can be made out-of-band by connecting the port used for management access to physically separate network infrastructure. Obviously, this is costly to implement, but out-of-band management is more secure and means that access to the device is preserved when there are problems affecting the production network.

With an in-band connection, better security can be implemented by using a VLAN to isolate management traffic. This makes it harder for potential eavesdroppers to view or modify traffic passing over the management interface. This sort of virtual OOB does still mean that access could be compromised by a system-wide network failure, however.




Use a secure connection protocol (HTTPS rather than HTTP, or SSH rather than Telnet) for the management interface. This applies to OOB too, but it is critical for in-band management.

Jump Boxes

One of the challenges of managing hosts exposed to the Internet, such as in a screened subnet or cloud network, is providing administrative access to the servers and appliances located within it. On the one hand, a link is necessary; on the other, the administrative interface could be compromised and exploited as a pivot point into the rest of the network. Consequently, management of hosts permitted to access administrative interfaces on hosts in the secure zone must be tightly controlled. Configuring and auditing this type of control when there are many different servers operating in the zone is complex.

One solution to this complexity is to add a single administration server, or **jump box/host/server**, to the secure zone. The jump box only runs the necessary administrative port and protocol, such as SSH or RDP. Administrators connect to the jump box and then use the jump host to connect to the admin interface on the application server. The application server’s admin interface has a single entry in its ACL (the jump server) and denies connection attempts from any other hosts.



Securing management traffic using a jump server. (Images © 123RF.com.)

API Connection Methods

Networks are increasingly making use of automation and orchestration to configure physical and virtual appliances via scripts instead of manually. An appliance that can be automated via scripting exposes an application programming interface (API). The API is the means by which external entities interact with the appliance, calling it with expected parameters and receiving the expected output.

For example, Google's firewall service is an example of a virtual firewall (cloud.google.com/firewalls). The service can be deployed and configured using either the CSP's web console, or programmatically via a command line interface (CLI) or application programming interface (API). To add a firewall rule via the API, a developer would create a JavaScript Object Notation (JSON) formatted request with the appropriate rule syntax and post it to the firewall API endpoint, such as `https://compute.googleapis.com/compute/v1/projects/515web/global/firewalls`.

Most APIs are exposed to remote access connections. A failure of credential management is likely to be exploited by malicious actors. You must enforce strong authentication policies to mitigate risks:

- Do not use the root user account for any day-to-day logon activity or automation. Configure specific accounts for automation that are allocated with least privileges only.
- Principals—user accounts, security groups, roles, and services—are enabled for programmatic access by assigning a secret key to the account. Only the secret key (not the ordinary account credential) can be used for programmatic access. When a secret key is generated for an account, it must immediately be transferred to the host and kept securely on that host.
- Only use secure protocols, such as HTTPS, for API communications. Configure mutual authentication and access controls so that API requests can only be issued from authorized clients.

Module 13

Summary

You should be able to explain WAN provider links and compare and contrast remote access methods and security implications.

Guidelines for Supporting WAN Links and Remote Access Methods

Follow these guidelines to support WAN links and remote access methods:

- Evaluate available types of WAN provider, including leased line, DSL, cable, and FTTx, to select a service that matches reliability and bandwidth requirements.
- Provision layer 1/2 WAN connectivity at the demarc/entrance facilities by connecting an appropriate type of digital modem to the service provider smartjack or network terminal. If the digital modem is not provisioned as a WAN interface card, connect it to a router to establish connectivity at layer 3.
- Develop a remote access policy to ensure only authorized users can connect and ensure that the network is not compromised by remote clients with weak security configurations.
- Support client-to-site VPNs and/or remote desktop services by selecting a protocol supported by client devices and installing the remote access server to the network edge by using a secure firewall configuration to prevent compromise.
- Support remote access to critical network infrastructure using secure protocols (SSH/RDP) over out-of-band links. Provision secure programmatic/API access where appropriate.

Module 14

Summarizing Cloud Concepts

Module Introduction

As the Internet becomes more robust and capable of matching the performance of local networks, many services are being moved from on-premises servers to cloud providers. Even where services are kept on-site, the different requirements and design principles of datacenters are essential competencies for network technicians at all levels.

This module completes the Network+ course by summarizing the software-driven virtualization, automation, and orchestration functionality that underpins cloud services.

Module Objectives

In this module, you will do the following:

- Explain datacenter and storage network architecture.
- Summarize cloud concepts.
- Summarize the use of software, coding, and zero trust in modern network environments.

Lesson 14.1

Datacenter and Storage Networks



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.*
- 1.5 Compare and contrast transmission media and transceivers.*
- 1.6 Compare and contrast network topologies, architectures, and types.*

Datacenters and storage networks play critical parts in both on-premises and cloud networks. Understanding the different topology and automation requirements of these networks will be critical for pursuing a successful career in networking.

As you study this lesson, answer the following questions:

- What is the difference between north/south and east/west traffic patterns?
- What is a spine and leaf topology?
- What components and cabling are used to create a storage network, and how are these different from Ethernet?

Datacenter Network Design

A datacenter is a site that is dedicated to provisioning server resources. The datacenter hosts network services (such as authentication, addressing, and name resolution), application servers, and storage area networks (SANs). Most datacenters are housed in purpose-built facilities, but some of the concepts also apply to server rooms.

A datacenter has dedicated networking, power, climate control, and physical access control features all designed to provide a highly available environment for running critical applications. Unlike an office network, a datacenter contains no client PCs, other than hardened secure administrative workstations (SAWs) used solely to manage servers.

Historically, datacenters were designed to use the same three-tiered architecture as an enterprise campus network, with core, distribution, and access layer switches. The design of modern applications as services with virtualization and on-demand instances has changed the nature of datacenter traffic flows. These changes are reflected in different topology designs in the datacenter.

Traffic that goes to and from a datacenter is referred to as **north-south**. This traffic represents clients outside the datacenter making requests and receiving responses. Corporate network traffic flows are also typically north-south. A client device is located on a workgroup switch connected to a router, while the server is connected to a separate switch or VLAN. Traffic from the client to the server passes “north” from the client’s switch to the router and then back “south” to the server’s switch.


In datacenters that support cloud and other Internet services, most traffic is actually between servers within the datacenter. This is referred to as **east-west traffic**. Consider a client uploading a photograph as part of a social media post. The image file might be checked by an analysis server for policy violations (indecent or copyright images, for instance), a search/indexing service would be updated with the image metadata, the image would be replicated to servers that provision content delivery networks (CDNs), the image would be copied to backup servers, and so on. A single request to the cloud tends to cascade into multiple requests and transfers *within* the cloud. Consequently, datacenters need to use a topology that optimizes secure server-to-server communications.

The preponderance of east-west traffic complicates security design. If each of these cascading transactions were to pass through a firewall or other security appliance, it would create a severe bottleneck. These requirements are driving the creation of virtualized security appliances that can monitor traffic as it passes between servers (blogs.cisco.com/security/trends-in-data-center-security-part-1-traffic-trends). At the same time, security implementations are moving toward zero trust architectures. Zero trust implies a highly segmented network where each request from one server to another must be authenticated and authorized.

Spine and Leaf Topology

The **spine and leaf topology** provides better support for east-west traffic and the use of SDN and overlay networks within datacenters. A spine and leaf topology has two layers:

- The spine layer comprises a backbone of top-tier distribution switches. Note that while this is described as a backbone, the spine switches are not linked to one another.
- The leaf layer contains access switches. Each access switch is connected to every spine switch in a full mesh topology. The access switches never have direct connections to one another.



Spine and leaf topology diagram. (Image © 123RF.com.)

The spine and leaf topology has a number of advantages:

- Each server is only ever a single hop from the backbone, making network latency more predictable.
- There are multiple redundant paths between a leaf switch and the backbone, allowing for load balancing and failover.
- As there are no direct connections between spine switches in the backbone or between leaf switches, the network is loop free and does not need to run spanning tree. Instead, each leaf switch runs a protocol called Equal Cost Multipathing (ECMP) to distribute traffic between the links to the spine switches.
- Servers are connected to multiple leaf switches for multipath redundancy, using a first hop gateway protocol to determine the active path.
- Scalability is improved because adding spine and/or leaf nodes does not change the topology. This means adding capacity for service needs that change unpredictably, such as storage, is easy.


The leaf layer access switches are implemented as **top-of-rack (ToR)** switch models. These are switch models designed to provide high-speed connectivity to a rack of server appliances and support higher bandwidths than ordinary workgroup switches. For example, where a workgroup switch might have 1 Gbps access ports and a 10 Gbps uplink port, top-of-rack switches have 10 Gbps access ports and 40/100 Gbps uplink ports.



A ToR switch doesn't have to be placed at the top of the server rack. This is a common practice, however, as it ensures cleaner cable management and better accessibility.

Storage Area Networks

Most datacenters (and server environments generally) require shared access to large amounts of storage. A storage device installed inside a server is referred to as direct-attached. This type of storage is typically only used for OS and software images. Variable data is hosted on a **storage area network (SAN)**. A SAN provisions access to storage devices using block input/output (I/O). Each read or write operation addresses the actual location of data on the media, just like direct-attached storage. The difference is that these read/write requests are communicated from the server to the storage device over the network. This does require an extremely fast and reliable network that is dedicated only to the storage function. A SAN is isolated from the main network. It is only accessed by servers, not by client PCs and laptops. SAN clients are servers running databases or applications that require access to shared storage.



Storage area network. (Images © 123RF.com.)

A SAN can integrate different types of storage technology—RAID arrays and tape libraries, for instance. It can contain a mixture of high-speed and low-cost devices, allowing for tiered storage to support different types of file access requirements without having to overprovision high-cost, fast drives.

Fibre Channel

A SAN fabric can be implemented using a variety of technologies.

Fibre Channel

Fibre Channel is defined in the T11 ANSI standard. The British spelling “fibre” is deliberately chosen to distinguish the standard from fiber optic cabling, which it often uses but on which it does not rely. A SAN based on a Fibre Channel (FC) Switched Fabric (FC-SW) involves three main types of components:

- **Initiator**—This is a client device of the SAN, such as a file or database server installed with a fibre channel host bus adapter (HBA).
- **Target**—This is the network port for a storage device. Typical devices include single drives, RAID drive arrays, tape drives, and tape libraries. Space on the storage devices is divided into logical volumes, each identified by a 64-bit logical unit number (LUN). The initiator will use SCSI, Serial Attached SCSI (SAS), SATA, or Nonvolatile Memory Express (NVMe) commands to operate the storage devices in the network, depending on which interface they support. Most devices have multiple ports for load balancing and fault tolerance.

The initiators and targets are identified by 64-bit WorldWide Names (WWN), similar to network adapter MAC addresses. Collectively, initiators and targets are referred to as nodes. Nodes can be allocated their own WWN, referred to as a WWNN (WorldWide Node Name). Also, each port on a node can have its own WorldWide Port Name (WWPN).

- **FC switch**—This provides the interconnections between initiators and targets (a fabric). The switch topology and interconnections would be designed to provide multiple paths between initiators and targets, allowing for fault tolerance and load balancing. High performance FC switches are often referred to as directors.

Fibre Channel can use rates from 1GFC (1 Gbps) up to 128GFC. Fibre Channel uses dedicated SFP+ and QSFP optical and twinax modular transceivers. Fibre Channel transceivers and Ethernet transceivers are not interchangeable.



As most SANs now depend on the use of fast, highly parallel SSDs, the NVMe specification is typically preferred over the older SCSI interface. Using NVMe in a networked environment is referred to as NVMe over Fabrics (NVMe-oF). NVMe over FC can be referred to as either FC-NVMe or NVMe/FC.

Converged Ethernet

The reason for using Fibre Channel fabric over standard Ethernet is that a SAN requires quality of service (QoS) mechanisms to ensure flow control and guaranteed delivery. A new iteration of Ethernet, referred to as lossless Ethernet, Data Center Ethernet, or Converged Enhanced Ethernet, has been developed as an alternative SAN fabric.

Fibre Channel over Ethernet (FCoE) is a means of delivering Fibre Channel packets over lossless Ethernet components. FCoE requires special 10/40/100G adapters that combine the function of NIC and HBA, referred to as converged network adapters (CNAs).

A more modern option is NVMe over Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE). RDMA is a way of offloading storage transfers from the CPU and OS to improve performance, compared to NVMe/FC.

TCP/IP

Internet Small Computer Systems Interface (iSCSI) is an IP tunneling protocol that enables the transfer of SCSI data over an IP-based network. iSCSI works with ordinary Ethernet network adapters and switches. iSCSI can be used to link SANs but is also seen as an alternative to Fibre Channel or Converged Ethernet, as it works with regular Ethernet adapters and switches.

Another option is NMEe over TCP (NVMe/TCP), which uses the reliability mechanisms built into TCP to substitute for the lossless mechanisms of FC or CE. While there is greater packet header and latency compared to NVMe/FC or RoCE, the use of standard Ethernet products can simplify procurement and support procedures.



A SAN should not be implemented on the same cabling as a production data network, even if technologies such as iSCSI and NVMe/TCP make that technically possible. The performance of the SAN will be heavily impacted. As a best practice, implement a dedicated network infrastructure (cabling, switches, and NICs) that is restricted to only SAN traffic.

Lesson 14.2

Cloud Concepts



Exam Objectives Covered

- 1.2 Compare and contrast networking appliances, applications, and functions.
- 1.3 Summarize cloud concepts and connectivity options.

Cloud services allow companies to outsource computing power and network/application infrastructure. Cloud computing encompasses different implementations and services. If you plan to use a cloud service, you need to know what the choices are and the advantages and disadvantages. Having a solid grasp of these choices will enable you to better manage and implement these technologies in your environment.

As you study this lesson, answer the following questions:

- How does the cloud provision scalable and elastic services?
- What are the types of cloud deployment model, and how do they impact security considerations?
- What is a cloud service model, and which are the main types?

Cloud Scalability and Elasticity

From the consumer point of view, cloud computing is a service that provides on-demand resources—server instances, file storage, databases, or applications—over a network, typically the Internet. The service is a cloud because the end user is not aware of or responsible for any details of the procurement, implementation, or management of the infrastructure that underpins those resources. The end user is interested in and pays for only the services provided by the cloud.

From the provider point of view, implementing a cloud is like provisioning any other type of large-scale datacenter. Cloud computing almost always uses one or more methods of virtualization to ensure that resources are quickly and automatically provisioned to the client who requires them.

Among other benefits, the cloud provides scalability and elasticity:

- **Scalability** means that the resources or costs involved in supplying the service to more users are linear. For example, if the number of users doubles in a scalable system, the costs to maintain the same level of service would also double (or less than double). If it takes quadruple the resources to maintain the service, the system is less scalable. Scalability can be achieved by adding nodes (horizontal/scaling out) or by adding resources to each node (vertical/scaling up).
- **Elasticity** refers to the system's ability to handle changes to demand in real time. A system with high elasticity will not experience loss of service or performance if demand suddenly doubles (or triples, or quadruples). Conversely, it may be important for the system to be able to reduce costs or deprovision resources when demand is low.

In order to meet scalability and elasticity requirements, cloud providers must be able to provision and deprovision resources automatically. This is achieved through resource pooling and virtualization. Resource pooling means that the hardware making up the cloud provider's datacenter is not dedicated or reserved to a single customer account. The layers of virtualization used in the cloud architecture allow the provider to provision more CPU, memory, disk, or network resources using management software, rather than (for instance) having to go to the datacenter floor, unplug a server, add a memory module, and reboot.

Cloud Deployment Models

A **cloud deployment model** classifies how the service is owned and provisioned. These cloud deployment models can be broadly categorized as follows:

- **Public (or multitenant)**—A service offered over the Internet by cloud service providers (CSPs) to cloud consumers, often referred to as tenants. With this model, businesses can offer subscriptions or pay-as-you-go financing, while at the same time providing lower-tier services free of charge. As a shared resource, there are risks regarding performance and security. Multicloud architectures are where the consumer organization uses services from more than one CSP.
- **Hosted Private (or single tenant)**—Hosted by a third party for the exclusive use of one organization. This is more secure and can guarantee a better level of performance, but it is correspondingly more expensive.
- **Private**—Cloud infrastructure that is completely private to and owned by the organization. In this case, there is likely to be one business unit dedicated to managing the cloud while other business units make use of it. With private cloud computing, organizations can exercise greater control over the privacy and security of their services. This type of delivery method is geared more toward banking and governmental services that require strict access control in their operations.
- **Hybrid**—A cloud computing solution that implements a mixed public/private solution. For example, a travel organization may run a sales website for most of the year using a private cloud but “break out” the solution to a public cloud at times when much higher utilization is forecast. As another example, a hybrid deployment may be used to provide some functions via a public cloud, but keep sensitive or regulated infrastructure, applications, and data on-premises.

Flexibility is a key advantage of cloud computing, but the implications for data risk must be well understood when you are moving data between private and public storage environments.

Cloud Service Models

As well as the deployment model—public, private, or hybrid—cloud services are often differentiated on the level of complexity and preconfiguration provided. These [cloud service models](#) are referred to as Something/Anything/Everything as a Service (XaaS). Some of the most common XaaS models are infrastructure, software, and platforms.

Infrastructure as a Service

Infrastructure as a Service (IaaS) is a means of provisioning IT resources such as servers, load balancers, and storage area network (SAN) components quickly. Rather than purchase these components and the Internet links they require, you rent them on an as-needed basis from the service provider's datacenter. Examples include Amazon Elastic Compute Cloud (aws.amazon.com/ec2), Microsoft Azure Virtual Machines (azure.microsoft.com/services/virtual-machines), and OpenStack (openstack.org).

Software as a Service

Software as a Service (SaaS) is a different model of provisioning software applications. Rather than purchasing software licenses for a given number of seats, a business would access software hosted on a supplier's servers on a pay-as-you-go or lease arrangement (on-demand). Virtual infrastructure allows developers to provision on-demand applications much more quickly than previously. The applications can be developed and tested in the cloud without the need to test and deploy on client computers. Examples include Microsoft Office 365 (support.office.com), Salesforce (salesforce.com), and Google Workspace (workspace.google.com).

Platform as a Service

Platform as a Service (PaaS) provides resources somewhere between SaaS and IaaS. A typical PaaS solution would deploy servers and storage network infrastructure (as per IaaS) but also provide a multi-tier web application/database platform on top. This platform could be based on Oracle or MS SQL or PHP and MySQL. Examples include Oracle Database (cloud.oracle.com/paas), Microsoft Azure SQL Database (azure.microsoft.com/services/sql-database), and Google App Engine (cloud.google.com/appengine).

As distinct from SaaS though, this platform would not be configured to actually do anything. Your own developers would have to create the software (the CRM or e-commerce application) that runs using the platform. The service provider would be responsible for the integrity and availability of the platform components, but you would be responsible for the security of the application you created on the platform.


The screenshot shows the AWS Management Console interface. At the top, there's a navigation bar with the AWS logo and various menu items like 'Services', 'Resource Groups', and 'Support'. The main content area is divided into three columns. The left column contains a sidebar with navigation links for 'EC2 Dashboard', 'Events', 'Tags', 'Reports', 'Limits', 'INSTANCES', 'IMAGES', 'ELASTIC BLOCK STORE', and 'NETWORK &'. The middle column, titled 'Resources', shows a summary of EC2 resources in the 'US East (N. Virginia)' region, including counts for Running Instances, Elastic IPs, Dedicated Hosts, Snapshots, Volumes, Load Balancers, Key Pairs, Security Groups, and Placement Groups. Below this, there's a 'Create Instance' section with a 'Launch Instance' button. The right column, titled 'Account Attributes', contains links to 'Supported Platforms', 'Resource ID length management', 'Additional Information', 'Getting Started Guide', 'Documentation', 'All EC2 Resources', 'Forums', 'Pricing', 'Contact Us', 'AWS Marketplace', and 'Find free software trial products in the AWS Marketplace from the EC2 Launch Wizard. Or try these'.

*Dashboard for Amazon Web Services Elastic Compute Cloud (EC2) IaaS/PaaS.
(Screenshot courtesy of Amazon.)*

Content Delivery Networks

A **content delivery network (CDN)** is a type of caching solution for high-traffic websites and application services. The purpose of a CDN is to provide a mix of high availability with performance and scalability. CDNs improve web performance by placing servers and media files near the consumers. Content can mean web apps, Software as a Service apps, websites, scripts, downloads, and other files. The result is quicker data access than would result from a single central storage location that might be very far from the consumer.

CDNs provide high availability by distributing content to multiple datacenters that are geographically dispersed. This also increases scalability. CDN content is delivered via multiple Internet backbones and service providers, further enhancing availability. The CDN servers replicate content to remain current with each other.



Content delivery network with datacenters around the world. (Images © 123RF.com.)

Lesson 14.3

Cloud Networking



Exam Objectives Covered

1.3 Summarize cloud concepts and connectivity options.

Cloud networking is the collection of technologies and practices that enable communications between the virtual instances created within the cloud and communications between the cloud service and external customers and networks. This is a hugely complex area, but it is important that you be able to summarize some of the connectivity options so that you can support the way your organization uses cloud services.

As you study this lesson, answer the following questions:

- What are workloads, and how are they supported by virtualization?
- What is a virtual private cloud, and how does it support familiar subnetting and routing concepts?
- What are the functions of cloud gateways and connectivity options?
- What options are there for filtering traffic within the cloud and between the cloud and external networks?

Cloud Instances

When using an IaaS or PaaS model, the customer will use the cloud to build a solution, such as creating a popular video streaming platform. Each solution will comprise one or more workloads. For example, each time an end-customer requests a video, a workload is created to stream the video to them. Each workload requires compute (CPU and memory), storage, and network resources. These are allocated using some type of virtualization:

- A virtual machine (VM) is an instance of a computer or network appliance running an OS and applications software. The VM can be allocated with a number of CPUs, an amount of system RAM, local storage, and network links. A VM can be managed just like a normal computer by connecting to it via RDP or SSH.
- A container is a lightweight computing instance designed to run a single application service or a single workload task. Containers don't have emulated hardware components, but they do still use compute resources and can connect to storage devices and networks.

A VM could be an instance of a general server running Windows or Linux. It is also possible to provision virtual appliances. With a virtual appliance, the vendor either develops a software product that emulates the functions of an existing dedicated hardware appliance (router, firewall, load balancer, or malware/intrusion detection, for instance) or creates software that implements that kind of functionality in a new product. These virtual appliances might be developed against a standard architecture, such as ETSI's **Network Functions Virtualization (NFV)**. NFV divides the provisioning of these appliances into three domains:

- **Virtual network function (VNF)**—Specifies and deploys instances of each virtual appliance. VNFs are designed to run as VMs on standard CPU platforms.
- **NFV infrastructure**—Controls the allocation of compute (CPU and memory) plus storage and networking resources to each VNF.
- **Management and orchestration (MANO)**—Positions VNFs within workflows to perform the forwarding and filtering tasks they are designed for.

The advantage of using a standard NFV architecture is that it makes the configuration and operation of the appliance more open to automation and orchestration via scripting.

Virtual Private Clouds

Within the cloud, the CSP establishes a virtualization layer that abstracts the underlying physical network of the datacenter hosting the cloud. This allows the CSP to operate a public cloud where the networking performed by each tenant account is isolated from the others. In terms of tenant-configured cloud networking, there are various contexts:

- Networks by which the tenant operates and manages the cloud systems.
- Virtual networks established to connect VM and container instances within the cloud.
- Virtual networks by which cloud services are published to guests or customers on the Internet.

To establish “local” networks within the cloud to deploy instances to, each tenant can create one or more **virtual private clouds (VPCs)** attached to their account. By default, a VPC is isolated from other CSP accounts and from other VPCs operating in the same account. This means that tenant A cannot view traffic passing over tenant B's VPC. The instances assigned to each VPC are isolated from other VPCs. Any communications between them must be created by configuring routing. Within each VPC, the cloud consumer can assign an IPv4 CIDR block and configure one or more subnets within that block. Optionally, an IPv6 CIDR block can be assigned also.



These notes focus on features of networking in AWS. Other vendors support similar functionality, though sometimes with different terminology. For example, in Microsoft Azure, VPCs are referred to as virtual networks.

Cloud Gateways

As with on-premises networking, a **cloud gateway** refers to the route that instances within a VPC subnet use to establish communications with other subnets in the same VPC, subnets in other VPCs, or over the Internet.

Each subnet within a VPC can either be private or public. To configure a public subnet, first an Internet gateway (virtual router) must be attached to the VPC configuration. Secondly, the Internet gateway must be configured as the default route for each public subnet. If a default route is not configured, the subnet remains private, even if an Internet gateway is attached to the VPC. Each instance in the subnet must also be configured with a public IP in its cloud profile. The Internet gateway performs 1:1 network address translation (NAT) to route Internet communications to and from the instance.



The instance network interface is not configured with this public IP address. The instance's network interface is configured with an IP address for the subnet. The public address is used by the virtualization management layer only. Public IP addresses can be assigned from your own pool or from a CSP-managed service, such as Amazon's Elastic IP (docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html).

There are other ways to provision external connectivity for a subnet if it is not appropriate to make it public:

- **NAT gateway**—This feature allows an instance to connect out to the Internet or to other AWS services but does not allow connections initiated from the Internet.
- **VPN**—There are various options for establishing connections to and between VPCs using virtual private networks (VPNs) at the software layer or using CSP-managed features.



Note that both an Internet gateway and a NAT gateway use NAT, but in different ways. An Internet gateway is a two-way gateway and requires the VM to be associated with a public IP address instance. A NAT gateway is a one-way (outbound only) gateway and does not require the VM to be associated with a public IP.

Cloud Connectivity Options

Cloud connectivity is the mechanism by which clients connect to whatever infrastructure, platform, or software that the tenant has configured in the cloud. There are several connectivity scenarios:

- Permitting secure access to the cloud for individual hosts or users.
- Connecting on-premises networks with cloud infrastructure.
- Connecting cloud infrastructure established in different geographical regions.
- Creating multicloud infrastructure using different cloud providers.

Internet/Virtual Private Network

A virtual private network (VPN) solution means that the tenant configures a VPN gateway for the VPC. Customers can establish a connection to this VPN gateway using either a client-to-site model or a site-to-site model. A site-to-site model would be used to connect cloud instances to an on-premises network or to another provider's cloud. Within the cloud, a virtual customer gateway is configured to represent the public IP address and security properties of the on-premises site.

VPN-based methods have the advantages of being cost-effective and straightforward to set up wherever there is Internet connectivity, which is ideal for organizations that have a fragmented or distributed network structure. However, any connection running over the public Internet can suffer from poor performance due to latency and bandwidth throttling, so this would not normally be a solution for a mission-critical or high-volume application.

Direct Connect/Colocation

Colocation within a datacenter offers a higher bandwidth solution by providing a direct connect or private link. The customer establishes infrastructure within a datacenter supported by the cloud provider or provisions a **direct connect** link from their enterprise network to the datacenter, possibly using private connections configured within a service provider's network. The datacenter installs a cross-connect cable or VLAN between the customer and the cloud provider, establishing a low-latency, high-bandwidth secure link. This solution is preferred for organizations which have a more centralized operation where the connection to the cloud can be from the main HQ and the company's own enterprise network is used to allow branch locations access.

Transit Gateways

Connectivity can also be configured between VPCs in the same account or with VPCs belonging to different accounts, and between VPCs and on-premises networks. Configuring additional VPCs rather than subnets within a VPC allows for a greater degree of segmentation between instances. A complex network might split segments between different VPCs across different cloud accounts for performance or compliance reasons.

Traditionally, VPCs can be interconnected using peering relationships and connected with on-premises networks using VPN gateways. These one-to-one VPC peering relationships can quickly become difficult to manage, especially if each VPC must interconnect in a mesh-like structure. A transit gateway is a simpler means of managing these interconnections. Essentially, a transit gateway is a virtual router that handles routing between the subnets in each attached VPC and any attached VPN gateways (aws.amazon.com/transit-gateway).



Amazon's white paper sets out options for configuring multi-VPC infrastructure in more detail (d1.awsstatic.com/whitepapers/building-a-scalable-and-secure-multi-vpc-aws-network-infrastructure.pdf).

Cloud Firewall Security

When traffic is routed between subnets, it can be subject to security rules that allow or block connections. These rules can be enforced by a cloud provider's virtual firewall solution, or the traffic could be routed or switched through a virtual firewall instance, or other security appliance. Firewalls work with multiple accounts, VPCs, subnets within VPCs, and instances within subnets to enforce the segmentation required by the architectural design. Segmentation may be needed for many different reasons, including separating workloads for performance and load balancing, keeping data processing within an isolated segment for compliance with laws and regulations, and compartmentalizing data access and processing for different departments or functional requirements.

Filtering decisions can be made based on packet headers and payload contents at various layers, identified in terms of the OSI model:

- **Network layer (layer 3)**—The firewall accepts or denies connections on the basis of IP addresses or address ranges and TCP/UDP port numbers (the latter are actually contained in layer 4 headers, but this functionality is still always described as basic layer 3 packet filtering).
- **Transport layer (layer 4)**—The firewall can store connection states and use rules to allow established or related traffic. Because the firewall must maintain a state table of existing connections, this requires more processing power (CPU and memory).
- **Application layer (layer 7)**—The firewall can parse application protocol headers and payloads (such as HTTP packets) and make filtering decisions based on their contents. This requires even greater processing capacity (or load balancing), or the firewall will become a bottleneck and increase network latency.

A cloud firewall can be implemented in several ways to suit different purposes:

- As software running on an instance. This sort of host-based firewall is identical to ones that you would configure for an on-premises host. It could be a stateful packet filtering firewall or a web application firewall (WAF) with a ruleset tuned to preventing malicious attacks. The drawback is that the software consumes instance resources and so is not very efficient. Also, managing the rulesets across many instances can be challenging.
- As a service at the virtualization layer to filter traffic between VPC subnets and instances. This equates to the concept of an on-premises network firewall.

Default cloud application-aware firewalls incur transaction costs, typically calculated on time deployed and traffic volume. These costs might be a reason to choose a third-party solution instead of the CSP's firewall service.

Security Groups and Security Lists

In AWS, basic packet filtering rules managing traffic that each instance will accept can be managed through **security groups** (docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html). A security group provides stateful inbound and outbound filtering at layer 4. The stateful filtering property means that it will allow established and related traffic if a new connection has been accepted.

The default security group allows any outbound traffic and any inbound traffic from instances also bound to the default security group. A custom security group sets the ports and endpoints that are allowed for inbound and outbound traffic. There are no deny rules for security groups; any traffic that does not match an allow rule is dropped. Consequently, a custom group with no rules will drop all network traffic. Multiple instances can be assigned to the same security group, and instances within the same subnet can be assigned to different security groups. You can assign multiple security groups to the same instance. You can also assign security groups to VPC endpoint interfaces.

Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

Assign a security group: ☐ Create a new security group ☒ Select an existing security group

Inbound rules for sg-0c731c5ae81f65a41 (Selected security groups: sg-0c731c5ae81f65a41)

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	10.0.0.0/24	
HTTPS	TCP	443	0.0.0.0/0	HTTPS
HTTPS	TCP	443	::/0	HTTPS

[Cancel](#) [Previous](#) [Review and Launch](#)

[Feedback](#) [English \(US\)](#) [Privacy Policy](#) [Terms of Use](#)

© 2008 - 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Adding a custom security group when launching a new instance in AWS EC2. This policy allows SSH access from a single IP address (redacted) and access to HTTPS from any IP address. (Screenshot courtesy of Amazon.com.)

Most cloud providers support similar filtering functionality, though they may be implemented differently. For example, in Azure, network security groups can be applied to network interfaces or to subnets (docs.microsoft.com/en-us/azure/virtual-network/security-overview). In Oracle Cloud Infrastructure (OCI), a **security list** is a set of rules that applies to an entire subnet. An OCI security group is similar to the AWS concept, as it can be applied to selected network interfaces (docs.oracle.com/en-us/iaas/Content/Network/Concepts/securityrules.htm#comparison).

Lesson 14.4

Modern Network Environments



Exam Objectives Covered

1.8 Summarize evolving use cases for modern network environments.

Twentieth-century networking was dominated by the on-premises model. On-premises networking means that companies install clients and servers at a single site, with extended access provisioned via branch office and remote access VPNs. Over the last decade, increased Internet and mobile bandwidth and reliability plus improvements in virtualization have driven the adoption of a new, modern network model. A modern network means that services are located in the cloud and can be accessed from diverse locations. A modern network is also designed for security and scalability. These requirements mean a greater role for scripted automation.

As you study this lesson, answer the following questions:

- What are the uses of infrastructure as code and software-defined networking, and how do they relate to one another?
- What is the function of an overlay network, and what is the role of VXLAN in implementing it?
- What is Secure Access Service Edge, and how does it relate to Security Service Edge and zero trust architecture?

Infrastructure as Code

The use of cloud technologies encourages the use of scripted approaches to provisioning, rather than installing operating systems and apps and making configuration changes or installing patches manually. An approach to infrastructure management where automation and orchestration fully replace manual configuration is referred to as **infrastructure as code (IaC)**.

One of the goals of IaC is to eliminate snowflake systems. A snowflake is a configuration or build that is different from any other. The lack of consistency—or **drift**—in the platform environment leads to security issues, such as patches that have not been installed, and stability issues, such as scripts that fail to run because of some small configuration difference.

IaC is often deployed to provision immutable architecture. Immutable architecture means that instances are never updated in place. If a change or update is required, a new instance is deployed to replace the old one. By rejecting manual configuration and ad hoc patching, IaC ensures idempotence. Idempotence means that making the same call with the same parameters will always produce the same result.



Note that IaC is not simply a matter of using scripts to perform repetitive tasks. Running scripts that have been written ad hoc is just as likely to cause environment drift as manual configuration. IaC means using carefully developed and tested scripts and orchestration playbooks to generate consistent builds.

Automation and Playbooks

When human engineers perform a task manually, they follow a standard operating procedure (SOP). **Automation** using scripted playbooks means that each configuration or build task is performed by a block of code. The script will take standard arguments as data. This means that there is less scope for uncertainty over configuration choices leading to errors. It also ensures that systems are in compliance with network or security policies.

The aim of an automation **playbook** is to script as many stages of the SOP as possible, leaving clearly defined interaction points for human intervention or verification. These interaction points should try to present all the contextual information and guidance needed for the technician to make a quick, informed decision about how to remediate a failed task, or one that needs manual input.

There are two types of automation tool:

- Imperative tools take the precise series of steps required to achieve the desired configuration as input.
- Declarative tools take the desired configuration as input and leave the detail of how that configuration should be achieved to the implementation platform.

Orchestration

Where automation focuses on making a single, discrete task easily repeatable, **orchestration** performs a sequence of automated tasks. For example, you might orchestrate adding a new VM to a load-balanced cluster. This end-to-end process might include provisioning the VM, configuring it with an app and network settings, adding the new VM to the load-balanced cluster, and reconfiguring the load-balancing weight distribution given the new cluster configuration. In doing this, the orchestrated steps would have to run numerous automated scripts or API service calls.

For orchestration to work properly, automated steps must occur in the right sequence, taking dependencies into account; it must provide the right security credentials at every step along the way; and it must have the rights and permissions to perform the defined tasks. Orchestration can automate processes that are complex, requiring dozens or hundreds of manual steps.

Automation and orchestration platforms connect to and provide administration, management, and orchestration for many popular cloud platforms and services. One of the advantages of using a third-party orchestration platform is protection from vendor lock in. If you wish to migrate from one cloud provider to another, or wish to move to a multicloud environment, automated workflows can often be adapted for use on new platforms. Industry leaders in this space include Chef (chef.io), Puppet (puppet.com), Ansible (ansible.com), and Kubernetes (kubernetes.io).

Uses for Infrastructure as Code

Infrastructure as code has many different applications, but many use cases center around the deployment of VMs and containers to cloud infrastructure.

Non-persistence and Templates

Automation works best in an environment designed for non-persistence.

Non-persistence means that any given instance is completely static in terms of processing function. Data is kept separate from the instance OS/app, so that the instance can be swapped out for an “as new” copy without suffering any configuration problems. When provisioning a new or replacement instance, the automation system may use one of two types of mastering instructions:

- **Master image**—This is the “gold” copy of a VM or container instance, with the OS, applications, and patches all installed and configured. This is faster than using a template, but keeping the image up to date can involve more work than updating a template.
- **Automated build from a template**—This is the build instructions (imperative automation) or desired configuration (declarative automation) for an instance. Rather than storing a master image, the software builds and provisions an instance according to the template instructions or desired configuration.

Upgrades

An upgrade is a move from an old OS or software version to a newer one. While an upgrade might have benefits for usability, performance, and security, the upgrade process can be highly complex. The changes introduced by an upgrade can have unforeseen impacts. An upgrade project must be treated as a major change and should be supported by test and rollback plans. The project must identify dependencies between systems and how they will be impacted by the upgrade. Automation can assist with this by speeding up deployment of systems into a test network, and performing scripted test suites to check for known or anticipated compatibility issues. It can also be used to deploy the upgraded systems on the production network.

Dynamic Inventories

Instances of VMs and containers launched into a cloud environment need to be tracked as inventory just like switches, routers, and servers in an on-premises network. Additional complexity comes from the fact that cloud instances are ephemeral. Rather than fixed asset IDs, they need to be identified by tags. Tags can be assigned in the cloud management system when the instance is launched. It is imperative to devise and enforce a tagging system that properly identifies ownership and roles for all instances.

As with on-premises virtualization, it is important to manage instances to avoid sprawl. Sprawl is where undocumented instances are launched and left unmanaged. As well as restricting rights to launch instances, you should configure logging and monitoring to track usage. This process is supported by dynamic inventory features of automation suites. A dynamic inventory queries the cloud API to return a list of instances and their properties for storage in a database.

Reusable Tasks

In any coding environment, it is helpful to minimize the number of lines of code or the number of scripts. The more code there is to manage, the more likely it is to develop bugs and vulnerabilities. A reusable task or module is a block of code that can perform a function in multiple different contexts. For example, given different inputs, the same block of code could launch a Windows VM into subnet A and a Linux VM into subnet B. Writing separate scripts for these tasks would be less scalable and more likely to lead to inconsistencies. Rather than separate scripts, you develop function libraries that can be reused for multiple tasks.

Source Control

Source control is the overall process of managing code for a software development project. When using infrastructure as code, it is important to use the correct version of a script to perform a task. Tasks performed by different versions can lead to configuration drift and noncompliance. Also, software development is typically a collaborative process, and there needs to be procedures and tools to allow multiple developers to work on the same project.

Version Control

Within the overall process of source control, **version control** is an ID system for each iteration of a software product or automation script. Most version control numbers represent both the version, as made known to the customer or end user, and internal build numbers for use in the development process. Version control supports the change management process for software development projects.

Central Repository

Software development environments use a repository server to maintain source code. One example is the Global Information Tracker commonly known as Git (git-scm.com). When a developer commits new or changed code to the repository, the new source code is tagged with an updated version number and the old version archived. This allows changes to be rolled back if a problem is discovered.

Branching

As scripts are developed and updated, there will be times when new features or changed functionality needs to be created and tested. To facilitate this, changes can be made in a branch copy of source code stored separately to the main or production version. When the branch code is ready, the developer issues a pull request, and it is tested and validated for merging back into the main branch.

Conflict Identification

Even with a branching strategy, there can still be instances where two (or more) competing changes to code need to be integrated back into the main branch. Conflict identification highlights these clashes and provides developers with tools to resolve them.

Software-Defined Networking

Cloud services require the rapid provisioning and deprovisioning of server instances and networks using automation and orchestration, plus the use of overlay networks to establish logical point to point links quickly and reliably. This means that these components must be fully accessible to scripting—representing the ideal of infrastructure as code. **Software-defined networking (SDN)** is a model for how these processes can be used to provision and deprovision networks. Some of the properties of SDN are the following:


- **Central policy management**—There is a single “source of truth” for how the network should operate. These business and security rules are automatically converted into device configuration states. There is central policy management but distributed policy enforcement. Also, status reporting ensures that “single pane of glass” monitoring and oversight is available to administrators.
- **Transport agnostic**—The overlay network can make use of any available forwarding fabric, whether that is Ethernet, Wi-Fi, 4G/5G cellular, leased line, or satellite. The underlying network fabric is automatically configured to establish the logical network link.
- **Zero-touch provisioning**—When new nodes are deployed to the network, they use automation to achieve the desired configuration, rather than needing to be manually configured by a technician. Similarly, if network policies change, nodes are reconfigured automatically.
- **Application aware**—Forwarding nodes can identify types of traffic, such as voice, video, or IoT. They can reserve capacity for these applications to ensure sufficient bandwidth, low latency, and lossless transfers.

SDN Architecture

In the SDN model defined by IETF (datatracker.ietf.org/doc/html/rfc7426), network functions are divided into three layers. The top and bottom layers are application and infrastructure:

- **Application layer**—Applies the business logic to make decisions about how traffic should be prioritized and secured and where it should be switched. This layer defines policies such as segmentation, ACLs, and traffic prioritization.
- **Infrastructure layer**—Uses the devices (physical or virtual) that handle the actual forwarding (switching and routing) of traffic and imposition of ACLs and other policy configurations for security.

The principal innovation of SDN is to insert a control layer between the Application and Infrastructure layers. The functions of the control plane are implemented by a virtual device called the SDN controller. Each layer exposes an application programming interface (API) that can be automated by scripts that call functions in the layer above or below. The interface between SDN applications and the SDN controller is described as the service interface or as the “northbound” API, while that between the SDN controller and infrastructure devices is the “southbound” API.



*Layers and components in a typical software-defined networking architecture.
(Images © 123RF.com.)*

Management Plane

In IETF's SDN model, there are separate forwarding (data) and operational planes at the infrastructure level. The operational plane implements device state, such as CPU and memory utilization. A management plane sits at the same level as the control plane to interface with the operational plane. This is used to implement monitoring of traffic conditions and network status.


Overlay Networks

An **overlay network** is used to implement logical links between nodes or networks. The overlay network abstracts the complexity of the underlying physical topology. A virtual private network (VPN) is an example of an overlay network. Other types of overlay network use encapsulation protocols and software-defined networking (SDN) to create a logical tunnel between nodes or networks that might be located in different physical topologies. An overlay network also allows for the segmentation of the same physical network. For example, a cloud provider can use an overlay network to isolate each tenant's traffic from other tenants.

When used inside the datacenter, overlay networks are typically implemented using virtual extensible LANs (VXLANs).

Virtual Extensible LANs

A **virtual extensible LAN (VXLAN)** uses layer 2 encapsulation to create an overlay network that runs on a layer 3 IP underlay network. Each overlay network is allocated a 24-bit VXLAN network identifier (VNI). A VXLAN Tunnel Endpoint (VTEP) server, switch, or router encapsulates the layer 2 frames tagged with the VNI in UDP packets. The UDP packets are routed over the IP network to another VTEP. The receiving VTEP decapsulates the packets to extract and process the frames.



VXLAN overlay network with VNI 101 allows the VMs with IP addresses 172..1 and ..2 to establish layer 2 adjacency, even though they are located on different hypervisors in physically separate racks. The hypervisors are configured with VTEP IDs and IP addresses. The VXLAN header allows the encapsulated packet from .1 to .2 to be tunneled through the underlying IP network, which uses a different 10. addressing scheme. (Image © 123RF.com.)



Traditional VLANs use a 12-bit field that allows for 4096 values, some of which are reserved. The 24-bit VXLAN format allows for millions of IDs.

As an alternative to manually configuring connections between VNIs and VTEPs, an **Ethernet VPN (EVPN)** uses Border Gateway Protocol (BGP) to advertise VXLAN networks and nodes as routes. This implements a control plane for automated configuration and management of the overlay network.

Datacenter Interconnect


An overlay network could also span two geographically separate locations, such as connecting instances or logical networks hosted in two different datacenters. **Data Center Interconnect (DCI)** refers to ways of creating links between datacenters and hosts/networks in different datacenters.

Datacenter services typically use cluster technologies and other applications that depend upon layer 2 adjacency. This means that the clustered servers or VMs must be part of the same broadcast domain and subnet, even if they are in different datacenters. Simply “stretching” the layer 2 boundaries over physical fiber links between datacenters with hosts in the different locations configured as part of the same VLAN can generate complex broadcast and spanning tree issues. This problem can be mitigated using VXLAN and Ethernet VPN (EVPN) to implement datacenter interconnects. EVPN allows servers to discover adjacent MAC addresses and forward data using an overlay network to tunnel traffic between them.

Software-Defined WAN

The hub and branch office design with on-premises datacenters has a number of performance and reliability drawbacks. Shifting services to one or more dedicated datacenters in the cloud mitigates some of these issues. Service availability and integrity is separated from site accessibility considerations. In this model, access to the datacenter from the corporate network, branch offices, and remote/teleworker locations can be facilitated through a **software-defined WAN (SD-WAN)**. SD-WAN replaces hub-and-spoke-type designs with more efficient, but still secure, connectivity to corporate clouds.

In a branch office topology, access to the datacenter or the cloud would be routed and authorized via the hub office. An SD-WAN is a type of overlay network that provisions a corporate WAN across multiple locations and can facilitate secure access to the cloud directly from a branch office or other remote location. It uses automation and orchestration to provision links dynamically based on application requirements and network congestion, using IPSec to ensure that traffic is tunneled through the underlying transport networks securely. An SD-WAN solution should also apply microsegmentation and zero trust security policies to ensure that all requests and responses are authenticated and authorized.



Components in an SD-WAN solution. (Image © 123RF.com.)


The SD-WAN is managed by a controller and management software located in a corporate datacenter or public cloud. Each site has a SD-WAN capable router, gateway, or VPN app. The SDN controller orchestrates connections to networks and clouds enrolled in the SD-WAN. It uses any available IP underlay network, such as broadband Internet, 4G/5G cellular, or private Multiprotocol Label Switching (MPLS) VPNs to provision the fastest or most reliable available transport to networks and clouds enrolled in the SD-WAN. The controller also ensures that each access request is authenticated and authorized.

Zero Trust Architecture

Zero trust architecture (ZTA) is a security model that assumes that all devices, users, and services are not inherently trusted, regardless of whether inside or outside a network's perimeter. Instead, the zero trust model requires all users and devices to be authenticated and authorized before accessing network resources. The zero trust model includes several fundamental concepts that provide a comprehensive security solution:

- **Policy-based authentication** recognizes that user identities are not static and that identity verification must be continuous and based on a user's current context and the resources they are attempting to access.
- **Threat scope reduction and least privilege access** means that access to network resources is granted on a need-to-know basis, and access is limited to only those resources required to complete a specific task. These concepts reduce the network's attack surface and limit the damage that a successful attack can cause.
- **Policy-driven authorization** describes how least privilege access control policies are used to enforce permissions and restrictions based on user identity, device posture, and network context.

In a zero trust architecture, the control and data planes are implemented separately and have different functions.



In a zero trust architecture, the control and data planes are implemented separately and have different functions.

The control plane manages policies that dictate how users and devices are authorized to access network resources. It is implemented through a centralized policy decision point. The policy decision point is responsible for defining policies that limit access to resources on a least privilege basis, monitoring network activity for suspicious behavior, and updating policies to reflect changing network conditions and security threats. The policy decision point is comprised of two subsystems:

- The policy engine is configured with subject and host identities and credentials, access control policies, up-to-date threat intelligence, behavioral analytics, and other results of host and network security scanning and monitoring. This comprehensive state data allows it to define an algorithm and metrics for making dynamic authentication and authorization decisions on a per-request basis.
- The policy administrator is responsible for managing the process of issuing access tokens and establishing or tearing down sessions, based on the decisions made by the policy engine. The policy administrator implements an interface between the control plane and the data plane.

Where systems in the control plane define policies and make decisions, systems in the data plane establish sessions for secure information transfers. In the data plane, a subject (user or service) uses a system (such as a client host PC, laptop, or smartphone) to make requests for a given resource. A resource is typically an enterprise app running on a server or cloud. Each request is mediated by a policy enforcement point. The enforcement point might be implemented as a software agent running on the client host that communicates with an app gateway. The policy enforcement point interfaces with the policy administrator to set up a secure data pathway if access is approved, or tear down a session if access is denied or revoked.



The processes implementing the policy enforcement point are the only ones permitted to interface with the policy administrator. It is critical to establish a root of trust for these processes so that policy decisions cannot be tampered with.

The data pathway established between the policy enforcement point and the resource is referred to as an implicit trust zone. For example, the outcome of a successful access request might be an IPSec tunnel established between a digitally signed agent process running on the client, a trusted web application gateway, and the resource server. Because the data is protected by IPSec transport encryption, no tampering by anyone with access to the underlying network infrastructure (switches, access points, routers, and firewalls) is possible.

The goal of zero trust design is to make this implicit trust zone as small as possible, and as transient as possible. Trusted sessions might only be established for individual transactions. This granular or microsegmented approach is in contrast with perimeter-based models, where trust is assumed once a user has authenticated and joined the network. In zero trust, place in the network is not a sufficient reason to trust a subject request. Similarly, even if a user is nominally authenticated, behavioral analytics might cause a request to be blocked or a session to be terminated.

Secure Access Service Edge

The shift of service provisioning from on-premises servers to the cloud has prompted the development of a new **Secure Access Service Edge (SASE)** paradigm. SASE combines the SD-WAN technologies used to implement secure tunnels from remote sites to enterprise services (the “A” for access part) with a Security Service Edge (SSE). In an SASE topology, when a user initiates an SD-WAN connection, the endpoint is not to a cloud service directly, but to a mediating SSE service.

The **Security Service Edge (SSE)** is a set of technologies that mediate access to cloud services and web applications. Zero trust architecture is one example of such an SSE technology. Another is a **cloud access security broker (CASB)**. Some of the functions of a CASB are the following:

- Enable single sign-on authentication and enforce access controls and authorizations from the enterprise network or remote employees to the cloud provider.
- Scan for malware and rogue or noncompliant device access.
- Monitor and audit user and resource activity.
- Mitigate data exfiltration by preventing access to unauthorized cloud services from managed devices.

In general, CASBs are implemented in one of three ways:

- **Forward proxy**—This is a security appliance or host positioned at the client network edge that forwards user traffic to the cloud network if the contents of that traffic comply with policy. This requires configuration of users’ devices or installation of an agent. In this mode, the proxy can inspect all traffic in real time, even if that traffic is not bound for sanctioned cloud applications. The problem with this mode is that users may be able to evade the proxy and connect directly. Proxies are also associated with poor performance as without a load balancing solution, they become a bottleneck and potentially a single point of failure.
- **Reverse proxy**—This is positioned at the cloud network edge and directs traffic to cloud services if the contents of that traffic comply with policy. This does not require configuration of the users’ devices. This approach is only possible if the cloud application has proxy support.
- **Application programming interface (API)**—Rather than placing a CASB appliance or host inline with cloud consumers and the cloud services, an API-based CASB uses brokers’ connections between the cloud service and the cloud consumer. For example, if a user account has been disabled or an authorization has been revoked on the local network, the CASB would communicate this to the cloud service and use its API to disable access there too. This depends on the API supporting the range of functions that the CASB and access and authorization policies demand. CASB solutions are quite likely to use both proxy and API modes for different security management purposes.

Another SSE technology is a secure web gateway (SWG). An on-premises SWG is a proxy-based firewall, content filter, and intrusion detection/prevention system that mediates user access to Internet sites and services.

Module 14

Summary

You should be able to explain datacenter network architecture, summarize cloud concepts and connectivity options, and identify requirements for software-defined network environments.

Guidelines for Supporting Cloud and Datacenter Architecture

Follow these guidelines to support the deployment of cloud and datacenter architecture and technologies:

- Identify apps and services that can utilize the elasticity and scalability benefits of cloud provision and determine an appropriate deployment model (such as public, private, and on-/off-premises) and service model (such as IaaS, SaaS, PaaS, or DaaS).
- Consider developing cloud-based apps using the infrastructure as code model to gain the most benefit from automation, orchestration, and software-defined networking.
- When implementing a private cloud/datacenter:
 - Consider a spine and leaf topology with aggregation and top-of-rack switch models to create a network fabric that best supports east-west traffic flows and use of overlay networks.
 - Identify virtualization and SAN products that can support the goals of elasticity and scalability and benefit from SDN and network function virtualization.
- When using a public cloud vendor, create a cloud responsibility matrix and perform regular risk assessments and security audits.
- Develop a WAN access strategy that provisions secure and high-performing links between corporate data networks, branch offices, remote teleworkers, and on-/off-premises datacenters and clouds, making use of technologies such as SD-WAN.

Appendix A

Mapping Course Content to CompTIA Certification

Achieving CompTIA Network+ certification requires candidates to pass Exam N10-009. This table describes where the exam objectives for Exam N10-009 are covered in this course.

1.0 Networking Concepts

1.1 Explain concepts related to the Open Systems Interconnection (OSI) reference model

Covered in

Layer 1 - Physical

Module 1, Lesson 2

Layer 2 - Data link

Module 1, Lesson 2

Layer 3 - Network

Module 1, Lesson 2

Layer 4 - Transport

Module 1, Lesson 2

Layer 5 - Session

Module 1, Lesson 2

Layer 6 - Presentation

Module 1, Lesson 2

Layer 7 - Application

Module 1, Lesson 2

1.2 Compare and contrast networking appliances, applications, and functions

Covered in

Physical and virtual appliances

Module 3, Lesson 2
Module 5, Lesson 1
Module 5, Lesson 4
Module 7, Lesson 2
Module 7, Lesson 4
Module 10, Lesson 5
Module 11, Lesson 1
Module 14, Lesson 1

Router

Module 5, Lesson 1

Switch

Module 3, Lesson 2

Firewall

Module 5, Lesson 4

Intrusion detection system (IDS)/intrusion prevention system (IPS)

Module 11, Lesson 1

Load balancer

Module 7, Lesson 4

Proxy

Module 10, Lesson 5

Network attached storage (NAS)

Module 7, Lesson 2

Storage area network (SAN)

Module 14, Lesson 1

Wireless

Module 12, Lesson 2

Access point (AP)

Module 12, Lesson 2

Controller

Module 12, Lesson 2

1.2 Compare and contrast networking appliances, applications, and functions	Covered in
Applications	Module 1, Lesson 1
Content delivery network (CDN)	Module 14, Lesson 2
Functions	Module 1, Lesson 1
Virtual private network (VPN)	Module 13, Lesson 2
Quality of service (QoS)	Module 8, Lesson 6
Time to live (TTL)	Module 5, Lesson 1
1.3 Summarize cloud concepts and connectivity options	Covered in
Network Functions Virtualization (NFV)	Module 14, Lesson 3
Virtual private cloud (VPC)	Module 14, Lesson 3
Network security groups	Module 14, Lesson 3
Network security lists	Module 14, Lesson 3
Cloud gateways	Module 14, Lesson 3
Internet gateway	Module 14, Lesson 3
Network address translation (NAT) gateway	Module 14, Lesson 3
Cloud connectivity options	Module 14, Lesson 3
VPN	Module 14, Lesson 3
Direct Connect	Module 14, Lesson 3
Deployment models	Module 14, Lesson 2
Public	Module 14, Lesson 2
Private	Module 14, Lesson 2
Hybrid	Module 14, Lesson 2
Service models	Module 14, Lesson 2
Software as a Service (SaaS)	Module 14, Lesson 2
Infrastructure as a Service (IaaS)	Module 14, Lesson 2
Platform as a Service (PaaS)	Module 14, Lesson 2
Scalability	Module 14, Lesson 2
Elasticity	Module 14, Lesson 2
Multitenancy	Module 14, Lesson 2
1.4 Explain common networking ports, protocols, services, and traffic types	Covered in
Protocols/Ports	
File Transfer Protocol (FTP) 20/21	Module 7, Lesson 2
Secure File Transfer Protocol (SFTP) 22	Module 7, Lesson 2
Secure Shell (SSH) 22	Module 13, Lesson 3
Telnet 23	Module 13, Lesson 3
Simple Mail Transfer Protocol (SMTP) 25	Module 7, Lesson 3
Domain Name System (DNS) 53	Module 6, Lesson 5

1.4 Explain common networking ports, protocols, services, and traffic types	Covered in
Dynamic Host Configuration Protocol (DHCP) 67/68	Module 6, Lesson 2
Trivial File Transfer Protocol (TFTP) 69	Module 7, Lesson 2
Hypertext Transfer Protocol (HTTP) 80	Module 7, Lesson 2
Network Time Protocol (NTP) 123	Module 7, Lesson 1
Simple Network Management Protocol (SNMP) 161/162	Module 8, Lesson 3
Lightweight Directory Access Protocol (LDAP) 389	Module 10, Lesson 2
Hypertext Transfer Protocol Secure (HTTPS) 443	Module 7, Lesson 2
Server Message Block (SMB) 445	Module 7, Lesson 2
Syslog 514	Module 8, Lesson 4
Simple Mail Transfer Protocol Secure (SMTPS) 587	Module 7, Lesson 3
Lightweight Directory Access Protocol over SSL (LDAPS) 636	Module 10, Lesson 2
Structured Query Language (SQL) Server 1433	Module 7, Lesson 2
Remote Desktop Protocol (RDP) 3389	Module 13, Lesson 3
Session Initiation Protocol (SIP) 5060/5061	Module 7, Lesson 3
Internet Protocol (IP) types	Module 4, Lesson 1
Internet Control Message Protocol (ICMP)	Module 4, Lesson 1
Transmission Control Protocol (TCP)	Module 6, Lesson 1
User Datagram Protocol (UDP)	Module 6, Lesson 1
Generic Routing Encapsulation (GRE)	Module 13, Lesson 2
Internet Protocol Security (IPSec)	Module 13, Lesson 2
Authentication Header (AH)	Module 13, Lesson 2
Encapsulating Security Payload (ESP)	Module 13, Lesson 2
Internet Key Exchange (IKE)	Module 13, Lesson 2
Traffic types	Module 4, Lesson 1
Unicast	Module 4, Lesson 1
Multicast	Module 4, Lesson 1
Anycast	Module 4, Lesson 1
Broadcast	Module 4, Lesson 1

1.5 Compare and contrast transmission media and transceivers	Covered in
Wireless	Module 12, Lesson 1
802.11 standards	Module 12, Lesson 1
Cellular	Module 12, Lesson 1
Satellite	Module 12, Lesson 1
Wired	Module 2, Lesson 1
	Module 2, Lesson 2
	Module 2, Lesson 4
802.3 standards	Module 2, Lesson 1
Single-mode vs. multimode fiber	Module 2, Lesson 4

1.5 Compare and contrast transmission media and transceivers	Covered in
Direct Attach Copper (DAC) cable	Module 2, Lesson 2
Twinaxial cable	Module 2, Lesson 2
Coaxial cable	Module 2, Lesson 2
Cable speeds	Module 2, Lesson 1
Plenum vs. non-plenum cable	Module 2, Lesson 2
Transceivers	Module 3, Lesson 1 Module 14, Lesson 1
Protocol	Module 3, Lesson 1
Ethernet	Module 3, Lesson 1
Fibre Channel (FC)	Module 14, Lesson 1
Form factors	Module 3, Lesson 1
Small form-factor pluggable (SFP)	Module 3, Lesson 1
Quad small form-factor pluggable (QSFP)	Module 3, Lesson 1
Connector types	Module 2, Lesson 2 Module 2, Lesson 4
Subscriber connector (SC)	Module 2, Lesson 4
Local connector (LC)	Module 2, Lesson 4
Straight tip (ST)	Module 2, Lesson 4
Multi-fiber push-on (MPO)	Module 2, Lesson 4
Registered Jack (RJ)11	Module 2, Lesson 2
RJ45	Module 2, Lesson 2
F-type	Module 2, Lesson 2
Bayonet Neill-Concelman (BNC)	Module 2, Lesson 2
1.6 Compare and contrast network topologies, architectures, and types	Covered in
Mesh	Module 1, Lesson 1
Hybrid	Module 5, Lesson 5
Star/hub and spoke	Module 1, Lesson 1
Spine and leaf	Module 14, Lesson 1
Point to point	Module 1, Lesson 1
Three-tier hierarchical model	Module 5, Lesson 5
Core	Module 5, Lesson 5
Distribution	Module 5, Lesson 5
Access	Module 5, Lesson 5
Collapsed core	Module 5, Lesson 5
Traffic flows	Module 14, Lesson 1
North-south	Module 14, Lesson 1
East-west	Module 14, Lesson 1

1.7 Given a scenario, use appropriate IPv4 network addressing	Covered in
Public vs. private	Module 4, Lesson 3
Automatic Private IP Addressing (APIPA)	Module 6, Lesson 3
RFC1918	Module 4, Lesson 3
Loopback/localhost	Module 4, Lesson 3
Subnetting	Module 4, Lesson 2
Variable length subnet mask (VLSM)	Module 4, Lesson 3
Classless Inter-Domain Routing (CIDR)	Module 4, Lesson 3
IPv4 address classes	Module 4, Lesson 3
Class A	Module 4, Lesson 3
Class B	Module 4, Lesson 3
Class C	Module 4, Lesson 3
Class D	Module 4, Lesson 3
Class E	Module 4, Lesson 3

1.8 Summarize evolving use cases for modern network environments	Covered in
Software-defined network (SDN) and software-defined wide area network (SD-WAN)	Module 14, Lesson 4
Application aware	Module 14, Lesson 4
Zero-touch provisioning	Module 14, Lesson 4
Transport agnostic	Module 14, Lesson 4
Central policy management	Module 14, Lesson 4
Virtual Extensible Local Area Network (VXLAN)	Module 14, Lesson 4
Data Center Interconnect (DCI)	Module 14, Lesson 4
Layer 2 encapsulation	Module 14, Lesson 4
Zero trust architecture (ZTA)	Module 14, Lesson 4
Policy-based authentication	Module 14, Lesson 4
Authorization	Module 14, Lesson 4
Least privilege access	Module 14, Lesson 4
Secure Access Service Edge (SASE)/ Security Service Edge (SSE)	Module 14, Lesson 4
Infrastructure as code (IaC)	Module 14, Lesson 4
Automation	Module 14, Lesson 4
Playbooks/templates/reusable tasks	Module 14, Lesson 4
Configuration drift/compliance	Module 14, Lesson 4
Upgrades	Module 14, Lesson 4
Dynamic inventories	Module 14, Lesson 4

1.8 Summarize evolving use cases for modern network environments	Covered in
Source control	Module 14, Lesson 4
Version control	Module 14, Lesson 4
Central repository	Module 14, Lesson 4
Conflict identification	Module 14, Lesson 4
Branching	Module 14, Lesson 4
IPv6 addressing	Module 4, Lesson 5
Mitigating address exhaustion	Module 4, Lesson 5
Compatibility requirements	Module 4, Lesson 5
Tunneling	Module 4, Lesson 5
Dual stack	Module 4, Lesson 5
NAT64	Module 4, Lesson 5

2.0 Network Implementation	
2.1 Explain characteristics of routing technologies	Covered in
Static routing	Module 5, Lesson 1
Dynamic routing	Module 5, Lesson 2
Border Gateway Protocol (BGP)	Module 5, Lesson 2
Enhanced Interior Gateway Routing Protocol (EIGRP)	Module 5, Lesson 2
Open Shortest Path First (OSPF)	Module 5, Lesson 2
Route selection	Module 5, Lesson 2
Administrative distance	Module 5, Lesson 2
Prefix length	Module 5, Lesson 2
Metric	Module 5, Lesson 2
Address translation	Module 5, Lesson 3
NAT	Module 5, Lesson 3
Port Address Translation (PAT)	Module 5, Lesson 3
First Hop Redundancy Protocol (FHRP)	Module 7, Lesson 4
Virtual IP (VIP)	Module 7, Lesson 4
Subinterfaces	Module 5, Lesson 6

2.2 Given a scenario, configure switching technologies and features	Covered in
Virtual Local Area Network (VLAN)	Module 5, Lesson 6
VLAN database	Module 5, Lesson 6
Switch Virtual Interface (SVI)	Module 5, Lesson 6
Interface configuration	Module 3, Lesson 2
	Module 3, Lesson 3
	Module 5, Lesson 6
Native VLAN	Module 5, Lesson 6
Voice VLAN	Module 5, Lesson 6

2.2 Given a scenario, configure switching technologies and features	Covered in
802.1Q tagging	Module 5, Lesson 6
Link aggregation	Module 3, Lesson 3
Speed	Module 3, Lesson 2
Duplex	Module 3, Lesson 2
Spanning tree	Module 3, Lesson 3
Maximum transmission unit (MTU)	Module 3, Lesson 3
Jumbo frames	Module 3, Lesson 3

2.3 Given a scenario, select and configure wireless devices and technologies	Covered in
Channels	Module 12, Lesson 1
Channel width	Module 12, Lesson 1
Non-overlapping channels	Module 12, Lesson 1
Regulatory impacts (802.11h)	Module 12, Lesson 1
Frequency options	Module 12, Lesson 1
2.4 GHz	Module 12, Lesson 1
5 GHz	Module 12, Lesson 1
6 GHz	Module 12, Lesson 1
Band steering	Module 12, Lesson 1
Service Set Identifier (SSID)	Module 12, Lesson 2
Basic Service Set Identifier (BSSID)	Module 12, Lesson 2
Extended Service Set Identifier (ESSID)	Module 12, Lesson 2
Network types	Module 12, Lesson 2
Mesh networks	Module 12, Lesson 2
Ad hoc	Module 12, Lesson 2
Point to point	Module 12, Lesson 2
Infrastructure	Module 12, Lesson 2
Encryption	Module 12, Lesson 3
Wi-Fi Protected Access 2 (WPA2)	Module 12, Lesson 3
WPA3	Module 12, Lesson 3
Guest networks	Module 12, Lesson 3
Captive portals	Module 12, Lesson 3
Authentication	Module 12, Lesson 3
Pre-shared key (PSK) vs. Enterprise	Module 12, Lesson 3
Antennas	Module 12, Lesson 2
Omnidirectional vs. directional	Module 12, Lesson 2
Autonomous vs. lightweight access point	Module 12, Lesson 2

2.4 Explain important factors of physical installations		Covered in
Important installation implications		Module 2, Lesson 3 Module 2, Lesson 4 Module 2, Lesson 5
Locations		Module 2, Lesson 3
Intermediate distribution frame (IDF)		Module 2, Lesson 3
Main distribution frame (MDF)		Module 2, Lesson 3
Rack size		Module 2, Lesson 5
Port-side exhaust/intake		Module 2, Lesson 5
Cabling		Module 2, Lesson 3
Patch panel		Module 2, Lesson 3
Fiber distribution panel		Module 2, Lesson 4
Lockable		Module 2, Lesson 5
Power		Module 2, Lesson 5
Uninterruptible power supply (UPS)		Module 2, Lesson 5
Power distribution unit (PDU)		Module 2, Lesson 5
Power load		Module 2, Lesson 5
Voltage		Module 2, Lesson 5
Environmental factors		Module 2, Lesson 5
Humidity		Module 2, Lesson 5
Fire suppression		Module 2, Lesson 5
Temperature		Module 2, Lesson 5
3.0 Network Operations		
3.1 Explain the purpose of organizational processes and procedures		Covered in
Documentation		Module 8, Lesson 1
Physical vs. logical diagrams		Module 8, Lesson 1
Rack diagrams		Module 8, Lesson 1
Cable maps and diagrams		Module 8, Lesson 1
Network diagrams		Module 8, Lesson 1
Layer 1		Module 8, Lesson 1
Layer 2		Module 8, Lesson 1
Layer 3		Module 8, Lesson 1
Asset inventory		Module 8, Lesson 1
Hardware		Module 8, Lesson 1
Software		Module 8, Lesson 1
Licensing		Module 8, Lesson 1
Warranty support		Module 8, Lesson 1
IP address management (IPAM)		Module 8, Lesson 1
Service-level agreement (SLA)		Module 8, Lesson 1
Wireless survey/heat map		Module 12, Lesson 2

3.1 Explain the purpose of organizational processes and procedures	Covered in
Lifecycle management	Module 8, Lesson 1
End-of-life (EOL)	Module 8, Lesson 1
End-of-support (EOS)	Module 8, Lesson 1
Software management	Module 8, Lesson 1
Patches and bug fixes	Module 8, Lesson 1
Operating system (OS)	Module 8, Lesson 1
Firmware	Module 8, Lesson 1
Decommissioning	Module 8, Lesson 1
Change management	Module 8, Lesson 1
Request process tracking/service request	Module 8, Lesson 1
Configuration management	Module 8, Lesson 1
Production configuration	Module 8, Lesson 1
Backup configuration	Module 8, Lesson 1
Baseline/golden configuration	Module 8, Lesson 1

3.2 Given a scenario, use network monitoring technologies	Covered in
Methods	Module 8, Lesson 3
SNMP	Module 8, Lesson 3
Traps	Module 8, Lesson 3
Management information base (MIB)	Module 8, Lesson 3
Versions (v2c, v3)	Module 8, Lesson 3
Community strings	Module 8, Lesson 3
Authentication	Module 8, Lesson 3
Flow data	Module 8, Lesson 6
Packet capture	Module 8, Lesson 5
Baseline metrics	Module 8, Lesson 2
Anomaly alerting/notification	Module 8, Lesson 4
Log aggregation	Module 8, Lesson 4
Syslog collector	Module 8, Lesson 4
Security information and event management (SIEM)	Module 8, Lesson 4
Application programming interface (API) integration	Module 8, Lesson 4
Port mirroring	Module 8, Lesson 5
Solutions	Module 8, Lesson 2
Network discovery	Module 8, Lesson 2
Ad hoc	Module 8, Lesson 2
Scheduled	Module 8, Lesson 2
Traffic analysis	Module 8, Lesson 6
Performance monitoring	Module 8, Lesson 2
Availability monitoring	Module 8, Lesson 2
Configuration monitoring	Module 8, Lesson 2

3.3 Explain disaster recovery (DR) concepts	Covered in
DR metrics	Module 7, Lesson 4
Recovery point objective (RPO)	Module 7, Lesson 4
Recovery time objective (RTO)	Module 7, Lesson 4
Mean time to repair (MTTR)	Module 7, Lesson 4
Mean time between failures (MTBF)	Module 7, Lesson 4
DR sites	Module 7, Lesson 4
Cold site	Module 7, Lesson 4
Warm site	Module 7, Lesson 4
Hot site	Module 7, Lesson 4
High-availability approaches	Module 7, Lesson 4
Active-active	Module 7, Lesson 4
Active-passive	Module 7, Lesson 4
Testing	Module 7, Lesson 4
Tabletop exercises	Module 7, Lesson 4
Validation tests	Module 7, Lesson 4

3.4 Given a scenario, implement IPv4 and IPv6 network services	Covered in
Dynamic addressing	Module 6, Lesson 2 Module 6, Lesson 3 Module 6, Lesson 4
DHCP	Module 6, Lesson 2 Module 6, Lesson 3
Reservations	Module 6, Lesson 2 Module 6, Lesson 3
Scope	Module 6, Lesson 2 Module 6, Lesson 3
Lease time	Module 6, Lesson 2 Module 6, Lesson 3
Options	Module 6, Lesson 2 Module 6, Lesson 3
Relay/IP helper	Module 6, Lesson 4
Exclusions	Module 6, Lesson 2 Module 6, Lesson 3
Stateless address autoconfiguration (SLAAC)	Module 6, Lesson 3
Name resolution	Module 6, Lesson 5

3.4 Given a scenario, implement IPv4 and IPv6 network services	Covered in
DNS	Module 6, Lesson 5
Domain Name System Security Extensions (DNSSEC)	Module 6, Lesson 5
DNS over HTTPS (DoH) and DNS over TLS (DoT)	Module 6, Lesson 5
Record types	Module 6, Lesson 5
Address (A)	Module 6, Lesson 5
AAAA	Module 6, Lesson 5
Canonical name (CNAME)	Module 6, Lesson 5
Mail exchange (MX)	Module 6, Lesson 5
Text	Module 6, Lesson 5
Name server (NS)	Module 6, Lesson 5
Pointer	Module 6, Lesson 5
Zone types	Module 6, Lesson 5
Forward	Module 6, Lesson 5
Reverse	Module 6, Lesson 5
Authoritative vs. non-authoritative	Module 6, Lesson 5
Primary vs. secondary	Module 6, Lesson 5
Recursive	Module 6, Lesson 5
Hosts file	Module 6, Lesson 6
Time protocols	Module 7, Lesson 1
NTP	Module 7, Lesson 1
Precision Time Protocol (PTP)	Module 7, Lesson 1
Network Time Security (NTS)	Module 7, Lesson 1

3.5 Compare and contrast network access and management methods	Covered in
Site-to-site VPN	Module 13, Lesson 2
Client-to-site VPN	Module 13, Lesson 2
Clientless	Module 13, Lesson 2
Split tunnel vs. full tunnel	Module 13, Lesson 2
Connection methods	Module 13, Lesson 3
SSH	Module 13, Lesson 3
Graphical user interface (GUI)	Module 13, Lesson 3
API	Module 13, Lesson 3
Console	Module 13, Lesson 3
Jump box/host	Module 13, Lesson 3
In-band vs. out-of-band management	Module 13, Lesson 3

4.0 Network Security	
4.1 Explain the importance of basic network security concepts	Covered in
Logical security	Module 9, Lesson 1
Encryption	Module 9, Lesson 1
Data in transit	Module 9, Lesson 1
Data at rest	Module 9, Lesson 1
Certificates	Module 10, Lesson 1
Public key infrastructure (PKI)	Module 10, Lesson 1
Self-signed	Module 10, Lesson 1
Identity and access management (IAM)	Module 10, Lesson 1
Authentication	Module 10, Lesson 1
Multifactor authentication (MFA)	Module 10, Lesson 1
Single sign-on (SSO)	Module 10, Lesson 1
Remote Authentication Dial-In User Service (RADIUS)	Module 10, Lesson 1
LDAP	Module 10, Lesson 2
Security Assertion Markup Language (SAML)	Module 10, Lesson 1
Terminal Access Controller Access Control System Plus (TACACS+)	Module 10, Lesson 1
Time-based authentication	Module 10, Lesson 1
Authorization	Module 10, Lesson 2
Least privilege	Module 10, Lesson 2
Role-based access control	Module 10, Lesson 2
Geofencing	Module 11, Lesson 3
Physical security	Module 11, Lesson 3
Camera	Module 11, Lesson 3
Locks	Module 11, Lesson 3
Deception technologies	Module 9, Lesson 1
Honeypot	Module 9, Lesson 1
Honeynet	Module 9, Lesson 1
Common security terminology	Module 9, Lesson 1
Risk	Module 9, Lesson 1
Vulnerability	Module 9, Lesson 1
Exploit	Module 9, Lesson 1
Threat	Module 9, Lesson 1
Confidentiality, integrity, and availability (CIA) triad	Module 9, Lesson 1
Audits and regulatory compliance	Module 9, Lesson 1
Data locality	Module 9, Lesson 1
Payment Card Industry Data Security Standard (PCI DSS)	Module 9, Lesson 1
General Data Protection Regulation (GDPR)	Module 9, Lesson 1

4.1 Explain the importance of basic network security concepts	Covered in
Network segmentation enforcement	Module 11, Lesson 1
Internet of things (IoT) and industrial internet of things (IIoT)	Module 11, Lesson 2
Supervisory control and data acquisition (SCADA), industrial control system (ICS), operational technology (OT)	Module 11, Lesson 2
Guest	Module 12, Lesson 3
Bring your own device (BYOD)	Module 12, Lesson 3
4.2 Summarize various types of attacks and their impact to the network	Covered in
Denial of service (DoS)/distributed denial of service (DDoS)	Module 9, Lesson 2
VLAN hopping	Module 9, Lesson 3
Media access control (MAC) flooding	Module 9, Lesson 3
Address Resolution Protocol (ARP) poisoning	Module 9, Lesson 3
ARP spoofing	Module 9, Lesson 3
DNS poisoning	Module 9, Lesson 4
DNS spoofing	Module 9, Lesson 4
Rogue devices and services	Module 9, Lesson 4
DHCP	Module 9, Lesson 4
AP	Module 12, Lesson 3
Evil twin	Module 12, Lesson 3
On-path attack	Module 9, Lesson 3
Social engineering	Module 9, Lesson 5
Phishing	Module 9, Lesson 5
Dumpster diving	Module 9, Lesson 5
Shoulder surfing	Module 9, Lesson 5
Tailgating	Module 9, Lesson 5
Malware	Module 9, Lesson 2
4.3 Given a scenario, apply network security features, defense techniques, and solutions	Covered in
Device hardening	Module 10, Lesson 3
Disable unused ports and services	Module 10, Lesson 3
Change default passwords	Module 10, Lesson 3
Network access control (NAC)	Module 10, Lesson 4
Port security	Module 10, Lesson 4
802.1X	Module 10, Lesson 4
MAC filtering	Module 10, Lesson 4
Key management	Module 10, Lesson 1
Security rules	Module 10, Lesson 5

4.3 Given a scenario, apply network security features, defense techniques, and solutions	Covered in
Access control list (ACL)	Module 10, Lesson 5
Uniform resource locator (URL) filtering	Module 10, Lesson 5
Content filtering	Module 10, Lesson 5
Zones	Module 11, Lesson 1
Trusted vs. untrusted	Module 11, Lesson 1
Screened subnet	Module 11, Lesson 1

5.0 Network Troubleshooting	
5.1 Explain the troubleshooting methodology	Covered in
Identify the problem	Module 1, Lesson 4
Gather information	Module 1, Lesson 4
Question users	Module 1, Lesson 4
Identify symptoms	Module 1, Lesson 4
Determine if anything has changed	Module 1, Lesson 4
Duplicate the problem, if possible	Module 1, Lesson 4
Approach multiple problems individually	Module 1, Lesson 4
Establish a theory of probable cause	Module 1, Lesson 4
Question the obvious	Module 1, Lesson 4
Consider multiple approaches	Module 1, Lesson 4
Top-to-bottom/bottom-to-top OSI model	Module 1, Lesson 4
Divide and conquer	Module 1, Lesson 4
Test the theory to determine the cause	Module 1, Lesson 4
If theory is confirmed, determine next steps to resolve problem	Module 1, Lesson 4
If theory is not confirmed, establish a new theory or escalate	Module 1, Lesson 4
Establish a plan of action to resolve the problem and identify potential effects	Module 1, Lesson 4
Implement the solution or escalate as necessary	Module 1, Lesson 4
Verify full system functionality and implement preventive measures if applicable	Module 1, Lesson 4
Document findings, actions, outcomes, and lessons learned throughout the process	Module 1, Lesson 4

5.2 Given a scenario, troubleshoot common cabling and physical interface issues	Covered in
Cable issues	Module 2, Lesson 6
Incorrect cable	Module 2, Lesson 6
Single mode vs. multimode	Module 2, Lesson 6
Category 5/6/7/8	Module 2, Lesson 6
Shielded twisted pair (STP) vs. unshielded twisted pair (UTP)	Module 2, Lesson 6

5.2 Given a scenario, troubleshoot common cabling and physical interface issues	Covered in
Signal degradation	Module 2, Lesson 6
Crosstalk	Module 2, Lesson 6
Interference	Module 2, Lesson 6
Attenuation	Module 2, Lesson 6
Improper termination	Module 2, Lesson 6
Transmitter (TX)/Receiver (RX) transposed	Module 2, Lesson 6
Interface issues	Module 3, Lesson 4
Increasing interface counters	Module 3, Lesson 4
Cyclic redundancy check (CRC)	Module 3, Lesson 4
Runts	Module 3, Lesson 4
Giants	Module 3, Lesson 4
Drops	Module 3, Lesson 4
Port status	Module 3, Lesson 4
Error disabled	Module 3, Lesson 4
Administratively down	Module 3, Lesson 4
Suspended	Module 3, Lesson 4
Hardware issues	Module 3, Lesson 1
Power over Ethernet (PoE)	Module 3, Lesson 4
Power budget exceeded	Module 3, Lesson 4
Incorrect standard	Module 3, Lesson 4
Transceivers	Module 3, Lesson 1
Mismatch	Module 3, Lesson 1
Signal strength	Module 3, Lesson 1

5.3 Given a scenario, troubleshoot common issues with network services	Covered in
Switching issues	Module 3, Lesson 3 Module 3, Lesson 4 Module 5, Lesson 7 Module 10, Lesson 5
STP	Module 3, Lesson 3 Module 3, Lesson 4
Network loops	Module 3, Lesson 4
Root bridge selection	Module 3, Lesson 3
Port roles	Module 3, Lesson 3
Port states	Module 3, Lesson 3
Incorrect VLAN assignment	Module 5, Lesson 7
ACLs	Module 10, Lesson 5
Route selection	Module 5, Lesson 7
Routing table	Module 5, Lesson 7
Default routes	Module 5, Lesson 7
Address pool exhaustion	Module 6, Lesson 4

5.3 Given a scenario, troubleshoot common issues with network services	Covered in
Incorrect default gateway	Module 4, Lesson 6
Incorrect IP address	Module 4, Lesson 6
Duplicate IP address	Module 4, Lesson 6
Incorrect subnet mask	Module 4, Lesson 6
5.4 Given a scenario, troubleshoot common performance issues	Covered in
Congestion/contention	Module 8, Lesson 6
Bottlenecking	Module 8, Lesson 6
Bandwidth	Module 8, Lesson 6
Throughput capacity	Module 8, Lesson 6
Latency	Module 8, Lesson 6
Packet loss	Module 8, Lesson 6
Jitter	Module 8, Lesson 6
Wireless	Module 12, Lesson 4
Interference	Module 12, Lesson 4
Channel overlap	Module 12, Lesson 4
Signal degradation or loss	Module 12, Lesson 4
Insufficient wireless coverage	Module 12, Lesson 4
Client disassociation issues	Module 12, Lesson 4
Roaming misconfiguration	Module 12, Lesson 4
5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues	Covered in
Software tools	Module 4, Lesson 4
	Module 5, Lesson 1
	Module 6, Lesson 1
	Module 6, Lesson 6
	Module 8, Lesson 2
	Module 8, Lesson 5
	Module 8, Lesson 6
Protocol analyzer	Module 8, Lesson 5
Command line	Module 4, Lesson 4
ping	Module 4, Lesson 4
tracert/traceroute	Module 5, Lesson 1
nslookup	Module 6, Lesson 6
tcpdump	Module 8, Lesson 5
dig	Module 6, Lesson 6
netstat	Module 6, Lesson 1
ip/ifconfig/ipconfig	Module 4, Lesson 4
arp	Module 4, Lesson 4

5.5 Given a scenario, use the appropriate tool or protocol to solve networking issues	Covered in
Nmap	Module 8, Lesson 2
Link Layer Discovery Protocol (LLDP)/Cisco Discovery Protocol (CDP)	Module 8, Lesson 2
Speed tester	Module 8, Lesson 6
Hardware tools	Module 2, Lesson 6
	Module 8, Lesson 5
	Module 12, Lesson 4
Toner	Module 2, Lesson 6
Cable tester	Module 2, Lesson 6
Taps	Module 8, Lesson 5
Wi-Fi analyzer	Module 12, Lesson 4
Visual fault locator	Module 2, Lesson 6
Basic networking device commands	Module 3, Lesson 4
	Module 5, Lesson 1
	Module 5, Lesson 6
show mac-address-table	Module 3, Lesson 4
show route	Module 5, Lesson 1
show interface	Module 3, Lesson 4
show config	Module 3, Lesson 4
show arp	Module 5, Lesson 1
show vlan	Module 5, Lesson 6
show power	Module 3, Lesson 4

Glossary

802.11 standards Specifications developed by IEEE for wireless networking over microwave radio transmission in the 2.4 GHz, 5 GHz, and 6 GHz frequency bands. The Wi-Fi standards brand has six main iterations: a, b, g, Wi-Fi 4 (n), Wi-Fi 5 (ac), and Wi-Fi 6 (ax). These specify different modulation techniques, supported distances, and data rates, plus special features, such as channel bonding, MIMO, and MU-MIMO.

802.11h Amendment to Wi-Fi standards that defines a Dynamic Frequency Selection (DFS) mechanism to avoid interference with radar and cellular communications in the 5 GHz frequency band.

802.1p IEEE standard defining a 3-bit (0 to 7) class of service priority field within the 802.1Q format.

802.1Q Trunking protocols enable switches to exchange data about VLAN configurations. The 802.1Q protocol is often used to tag frames destined for different VLANs across trunk links.

802.1X A standard for encapsulating EAP communications over a LAN (EAPoL) or WLAN (EAPoW) to implement port-based authentication.

access control list (ACL) The collection of access control entries (ACEs) that determines which subjects (user accounts, host IP addresses, and so on) are allowed or denied access to the object and the privileges given (read-only, read/write, and so on).

access point (AP) A device that provides a connection between wireless devices and can connect to wired networks, implementing an infrastructure mode WLAN.

access/edge layer Lowest tier in a hierarchical network topology acting as the attachment point for end systems.

active-active High availability cluster configuration where all nodes are utilized continually.

active-passive High availability cluster configuration where one or more nodes are only utilized during failover.

ad hoc network A type of wireless network where connected devices communicate directly with each other instead of over an established medium.

Address Resolution Protocol (ARP) Broadcast mechanism by which the hardware MAC address of an interface is matched to an IP address on a local network segment.

addressing (network) Unique identifier for a network node, such as a MAC address, IPv4 address, or IPv6 address.

adjacent channel interference (ACI) Troubleshooting issue where access points within range of one another are configured to use different but overlapping channels, causing increased noise.

administrative distance (AD) Metric determining the trustworthiness of routes derived from different routing protocols.

administratively down Switch or router port that has been purposefully disabled via the management interface.

advanced persistent threat (APT) Threat actors with the ability to craft novel exploits and techniques to obtain, maintain, and diversify unauthorized access to network systems over a long period.

angled physical contact (APC) Fiber optic connector finishing type that uses an angled polish for the ferrule.

antenna type Specially arranged metal wires that can send and receive radio signals, typically implemented as either an omnidirectional or a unidirectional type.

anycast IP delivery mechanism whereby a packet is addressed to a single host from a group sharing the same address.

Application layer OSI model layer providing support to applications requiring network services (file transfer, printing, email, databases, and so on).

application programming interface (API) Methods exposed by a script, program, or web application that allow other scripts or apps to interact with it.

arp command Utility to display and modify contents of host's cache of IP to MAC address mappings, as resolved by Address Resolution Protocol (ARP) replies.

ARP spoofing A network-based attack where an attacker with access to the target local network segment redirects an IP address to the MAC address of a computer that is not the intended recipient. This can be used to perform a variety of attacks, including DoS, spoofing, and on-path.

attenuation Attenuation, or degradation of a signal as it travels over media, determines the maximum distance for a particular media type at a given bit rate.

Authentication Header (AH) IPSec protocol that provides authentication for the origin of transmitted data as well as integrity and protection against replay attacks.

authoritative name server DNS server designated by a name server record for the domain that holds a complete copy of zone records.

Automatic Private IP Addressing (APIPA) Mechanism for Windows hosts configured to obtain an address automatically that cannot contact a DHCP server to revert to using an address from the range 169.254.x.y. This is also called a link local address.

automation Using scripts and APIs to provision and deprovision systems without manual intervention.

autonomous AP Access point whose firmware contains enough processing logic to be able to function autonomously and handle clients without the use of a wireless controller.

autonomous system (AS) Group of network prefixes under the

administrative control of a single organization used to establish routing boundaries.

availability monitoring Processes and tools that facilitate reporting and alerting when a host or app cannot be contacted over the network.

backup configuration Configuration settings that will be applied if an appliance, instance, or app is restored from backup media.

band steering Feature of Wi-Fi that allows an access point to try to ensure that clients use a particular frequency band, such as 5 GHz rather than 2.4 GHz.

bandwidth Generally used to refer to the amount of data that can be transferred through a connection over a given period. Bandwidth more properly means the range of frequencies supported by transmission media, measured in Hertz.

bandwidth speed tester Hosted utility used to measure actual speed obtained by an Internet link to a representative server or to measure the response times of websites from different locations on the Internet.

baseline metrics Values for resource utilization that assess the performance or stability of a service based on historical information or vendor guidance.

basic service set ID (BSSID) MAC address of an access point supporting a basic service area.

Bayonet Neill-Concelman (BNC) Twist and lock connector for coaxial cable.

bidirectional wavelength division multiplexing (BWDM) System that allows bidirectional data transfer over a single fiber strand by using separate wavelengths for transmit and receive streams.

bit rate Amount of data that can be transferred over a network connection in a given amount of time, typically measured in bits or bytes per second (or some more suitable multiple thereof). Transfer rate is also described variously as data rate, bit rate, connection speed,

transmission speed, or bandwidth. Transfer rates are often quoted as the peak, maximum, theoretical value; sustained, actual throughput is often considerably less.

Border Gateway Protocol (BGP) Path vector exterior gateway routing protocol used principally by ISPs to establish routing between autonomous systems.

botnet A group of hosts or devices that have been infected by a control program called a bot, which enables attackers to exploit the hosts to mount attacks.

bottleneck Troubleshooting issue where performance for a whole network or system is constrained by the performance of a single link, device, or subsystem.

bridge Intermediate system that isolates collision domains to separate segments while joining segments within the same broadcast domain.

bring your own device (BYOD) Security framework and tools to facilitate use of personally owned devices to access corporate networks and data.

broadcast Packet or frame addressed to all hosts on a local network segment, subnet, or broadcast domain. Routers do not ordinarily forward broadcast traffic. The broadcast address of IP is one where the host bits are all set to 1; at the MAC layer it is the address ff:ff:ff:ff:ff:ff.

broadcast domain Network segment in which all nodes receive the same broadcast frames at layer 2.

broadcast storm Traffic that is recirculated and amplified by loops in a switching topology, causing network slowdowns and crashing switches.

brute force attack A type of password attack where an attacker uses an application to exhaustively try every possible alphanumeric combination to crack encrypted passwords.

bugfix Update to software code that addresses a single discrete error and is typically applied in a development or test environment rather than a production one.

business continuity (BC) A collection of processes that enable an organization to maintain normal business operations in the face of some adverse event.

business impact analysis (BIA) Systematic activity that identifies organizational risks and determines their effect on ongoing, mission-critical operations.

cable map Physical plan showing cable routes through building spaces between communications closets and work areas.

cable stripper Tool for stripping cable jacket or wire insulation.

cable tester Two-part tool used to test successful termination of copper cable by attaching to each end of a cable and energizing each wire conductor in turn with an LED to indicate an end-to-end connection.

canonical notation Format for representing IPv6 addresses using hex double-bytes with colon delimitation and zero compression.

captive portal Webpage or website to which a client is redirected before being granted full network access.

carrier sense multiple access with collision avoidance (CSMA/CA) Mechanism used by 802.11 Wi-Fi standards to cope with contention over the shared access media.

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) In a contention-based system, each network device competes with the other connected devices for use of the transmission media. Contention-based systems require a set of protocols that reduce the possibility of data collisions, since if the devices compete and simultaneously send data packets, neither packet will reach its intended destination. The Carrier Sense Multiple Access (CSMA) protocols allow contention-based networks to successfully communicate by detecting activity on the network media (Carrier Sense) and reacting to this (for example, if the medium is busy). CSMA/CD (Collision Detection) recognizes a signal collision on the basis of electrical

fluctuations produced when signals combine.

cat cable standards ANSI/TIA/EIA cable category designations, with higher numbers representing better support for higher data rates.

cellular radio Standards for implementing data access over cellular networks are implemented as successive generations. For 2G (up to about 48 Kb/s) and 3G (up to about 42 Mb/s), there are competing GSM and CDMA provider networks. Standards for 4G (up to about 90 Mb/s) and 5G (up to about 300 Mb/s) are developed under converged LTE standards.

certificate authority (CA) A server that guarantees subject identities by issuing signed digital certificate wrappers for their public keys.

change management The process through which changes to the configuration of information systems are implemented as part of the organization's overall configuration management efforts.

channel In Wi-Fi, subdivision of frequency bands into smaller channels to allow multiple networks to operate at the same location without interfering with one another.

channel bonding Capability to aggregate one or more adjacent wireless channels to increase bandwidth.

cipher suite Lists of cryptographic algorithms that a server and client can use to negotiate a secure connection.

Cisco Discovery Protocol (CDP) Proprietary protocol used by Cisco network appliances to discover layer 2 adjacent devices or neighbors.

classful addressing Legacy form of IP addressing where the network ID is determined automatically from the first octet of the address. Netmasks that align to whole octet boundaries are still sometimes referred to as class A, B, or C.

classless interdomain routing (CIDR) Using network prefixes to aggregate routes to multiple network blocks ("supernetting"). This replaced the old

method of assigning class-based IP addresses based on the network size.

client-server Administration paradigm where some host machines are designated as providing server and services, and other machines are designated as client devices that only consume server services.

cloud access security broker (CASB) Enterprise management software designed to mediate access to cloud services by users across all types of devices.

cloud deployment model Classifying the ownership and management of a cloud as public, private, community, or hybrid.

cloud direct connection A dedicated connection between the on-premises network and a cloud service provider.

cloud gateway In cloud infrastructure, a virtual router that facilitates routing between subnets and public networks. External connectivity can be provisioned using various types of NAT and VPN.

cloud service model Classifying the provision of cloud services and the limit of the cloud service provider's responsibility as software, platform, infrastructure, and so on.

clustering A load balancing technique where a group of servers is configured as a unit and work together to provide network services.

coarse wavelength division multiplexing (CWDM) Technology for multiplexing up to 16 signal channels on a single fiber using different wavelengths.

coaxial cable Media type using two separate conductors that share a common axis categorized using the Radio Grade (RG) specifications.

co-channel interference (CCI) Troubleshooting issue where access points within range of one another are configured to use the same channel, causing increased contention.

cold site A predetermined alternate location where a network can be rebuilt after a disaster.

collapsed core Two-tier hierarchical network topology where access layer switches connect directly to a full mesh core layer.

collision domain Network segment where nodes are attached to the same shared access media, such as a bus network or Ethernet hub.

colocation Deploying private servers, network appliances, and interconnects to a hosted datacenter facility shared with other customers.

command and control (C2) Infrastructure of hosts and services with which attackers direct, distribute, and control malware over botnets.

community string In Simple Network Management Protocol (SNMP), a password-like value that permits a management system to access an agent.

confidentiality, integrity, and availability (CIA) Three principles of security control and management. Also known as the information security triad. Also referred to in reverse order as the AIC triad.

configuration baseline Settings for services and policy configuration for a network appliance or for a server operating in a particular application role (web server, mail server, file/print server, and so on).

configuration drift Risk that systems and networks will deviate from a baseline or golden configuration over time.

configuration management A process through which an organization's information systems components are kept in a controlled state that meets the organization's requirements, including those for security and compliance.

configuration monitoring Processes and tools that facilitate reporting and alerting when a host or app's configuration deviates from a baseline or golden configuration.

content delivery network (CDN) Distributing and replicating the components of any service (such as web apps, media and storage) across all the key service areas needing access to the content.

content filtering Security measure performed on email and Internet traffic to identify and block suspicious, malicious, and/or inappropriate content in accordance with an organization's policies.

convergence Process whereby routers agree on routes through the network to establish the same network topology in their routing tables (steady state). The time taken to reach steady state is a measure of a routing protocol's convergence performance.

core layer Highest tier in a hierarchical network topology providing interconnections between blocks.

crimper Tool to join a Registered Jack (RJ) form factor connector to the ends of twisted-pair patch cable.

crosstalk Phenomenon whereby one wire causes interference in another as a result of their close proximity.

cyclical redundancy check (CRC) Calculation of a checksum based on the contents of a frame used to detect errors.

data at rest Information that is primarily stored on specific media, rather than moving from one medium to another.

Data Center Interconnect (DCI) Technologies such as VXLAN and EVPN that establish links between hosts in two or more separate datacenter facilities.

data in transit Information that is being transmitted between two hosts, such as over a private network or the Internet.

Data Link layer OSI model layer responsible for transferring data between nodes.

data remnant Leftover information on a storage medium even after basic attempts have been made to remove that data.

data sovereignty In data protection, the principle that countries and states may impose individual requirements on data collected or stored within their jurisdiction.

datacenter Facility dedicated to the provisioning of reliable power, environmental controls, and network fabric to server computers.

deauthentication/disassociation attack Spoofing frames to disconnect a wireless station to try to obtain authentication data to crack.

decibel loss (dB) Loss of signal strength between a transmitter and receiver due to attenuation and interference measured in decibels.

decommissioning In asset management, the policies and procedures that govern the removal of devices and software from production networks, and their subsequent disposal through sale, donation, or as waste.

default gateway IP configuration parameter that identifies the address of a router on the local subnet that the host can use to contact other networks.

default route Entry in the routing table to represent the forwarding path that will be used if no other entries are matched.

default VLAN Default VLAN ID (1) for all unconfigured switch ports.

defense in depth Security strategy that positions the layers of diverse security control categories and functions as opposed to relying on perimeter controls.

demarcation point Location that represents the end of the access provider's network (and therefore their responsibility for maintaining it). The demarc point is usually at the minimum point of entry (MPOE). If routing equipment cannot be installed at this location, demarc extension cabling may need to be laid.

denial of service attack (DoS) Any type of physical, application, or network attack that affects the availability of a managed resource.

dense wavelength division multiplexing (DWDM) Technology for multiplexing 40 or 80 signal channels on a single fiber using different wavelengths.

DHCP relay Configuration of a router to forward DHCP traffic where the client and server are in different subnets.

DHCP snooping Switchport protection mechanism that blocks DHCP offers from unauthorized sources.

dictionary attack A type of password attack that compares encrypted passwords against a predetermined list of possible password values.

DiffServ Header field used to indicate a priority value for a layer 3 (IP) packet to facilitate quality of service (QoS) or class of service (CoS) scheduling.

dig command Utility to query a DNS server and return information about a particular domain name or resource record.

digital certificate Identification and authentication information presented in the X.509 format and issued by a certificate authority (CA) as a guarantee that a key pair (as identified by the public key embedded in the certificate) is valid for a particular subject (user or host).

direct attach copper (DAC) Factory-terminated twinax patch cords used for 10+ Gbps Ethernet connections, typically between rack-mounted appliances.

directly connected route Entry in the routing table representing a subnet in which the router has an active interface.

disassociation Management frame handling process by which a station is disconnected from an access point.

disaster recovery plan (DRP) A documented and resourced plan showing actions and responsibilities to be used in response to critical incidents.

discretionary access control (DAC) An access control model where each resource is protected by an access control list (ACL) managed by the resource's owner (or owners).

distance vector Algorithm used by routing protocols that select a forwarding path based on the next hop router with the lowest hop count to the destination network.

distributed denial-of-service (DDoS) An attack that involves the use of infected Internet-connected computers and devices to disrupt the normal

flow of traffic of a server or service by overwhelming the target with traffic.

distribution system (DS) Connecting access points to a switched network via cabling to facilitate roaming within an extended service area (ESA). A wireless distribution system uses access points configured in repeater mode to facilitate roaming.

distribution/aggregation layer Intermediate tier in a hierarchical network topology providing interconnections between the access layer and the core.

DNS caching Data store on DNS clients and servers holding results of recent queries.

DNS over HTTPS (DoH) Protocol that mitigates risks from snooping and modification when a client queries a DNS server by encapsulating DNS traffic within an HTTP-Secure (HTTPS) session.

DNS over TLS (DoT) Protocol that mitigates risks from snooping and modification when a client queries a DNS server by encapsulating DNS traffic within a Transport Layer Security (TLS) session.

DNS spoofing An attack where a threat actor injects false resource records into a client or server cache to redirect a domain name to an IP address of the attacker's choosing.

domain name system (DNS) Service that maps fully qualified domain name labels to IP addresses on most TCP/IP networks, including the Internet.

Domain Name System Security Extensions (DNSSEC) Security protocol that provides authentication of DNS data and upholds DNS data integrity.

dotted decimal notation Format for expressing IPv4 addresses using four decimal values from 0 to 255 for each octet.

dual stack Host operating multiple protocols simultaneously on the same interface. Most hosts are capable of dual stack IPv4 and IPv6 operation, for instance.

dumpster diving The social engineering technique of discovering things about an organization (or person) based on what it throws away.

Dynamic Host Configuration Protocol (DHCP) Protocol used to automatically assign IP addressing information to hosts that have not been configured manually.

dynamic route Entry in the routing table that has been learned from another router via a dynamic routing protocol.

east-west traffic Design paradigm accounting for the fact that datacenter traffic between servers is greater than that passing in and out (north-south).

effective isotropic radiated power (EIRP) Signal strength from a transmitter, measured as the sum of transmit power, antenna cable/connector loss, and antenna gain.

elasticity Property by which a computing environment can add or remove resources in response to increasing and decreasing demands in workload.

electromagnetic interference (EMI) Noise that occurs when a magnetic field around one electrical circuit or device interferes with the signal being carried on an adjacent circuit.

Encapsulating Security Payload (ESP) IPSec sub-protocol that enables encryption and authentication of the header and payload of a data packet.

encapsulation A method by which protocols build data packets by adding headers and trailers to existing data.

encryption Scrambling the characters used in a message so that the message can be seen but not understood or modified unless it can be deciphered. Encryption provides for a secure means of transmitting data and authenticating users. It is also used to store data securely. Encryption uses different types of algorithm/cipher and one or more keys. The size of the key is one factor in determining the strength of the encryption product.

end of life (EOL) Product life cycle phase where mainstream vendor support is no longer available.

end of service life (EOSL) Product life cycle phase where support is no longer available from the vendor.

Enhanced Interior Gateway Routing Protocol (EIGRP) Advanced distance vector dynamic routing protocol using bandwidth and delay metrics to establish optimum forwarding paths.

enterprise authentication A wireless network authentication mode where the access point acts as pass-through for credentials that are verified by an AAA server.

enumeration An attack that aims to list resources on the network, host, or system as a whole to identify potential targets for further attack.

escalation In the context of support procedures, incident response, and breach-reporting, escalation is the process of involving expert and senior staff to assist in problem management.

Ethernet Standards developed as the IEEE 802.3 series describing media types, access methods, data rates, and distance limitations at OSI layers 1 and 2 using xBASE-y designations.

Ethernet header Fields in a frame used to identify source and destination MAC addresses, protocol type, and error detection.

Ethernet virtual private network (EVPN) Using Border Gateway Protocol (BGP) to advertise virtual extensible LAN (VXLAN) networks as routes.

evil twin A wireless access point that deceives users into believing that it is a legitimate network access point.

explicit deny Firewall ACL rule configured manually to block any traffic not matched by previous rules.

exploit A specific method by which malware code infects a target host, often via some vulnerability in a software process.

Extended Service Set ID (ESSID) Network name configured on multiple

access points to form an extended service area.

extended unique identifier (EUI) IEEE's preferred term for a network interface's unique identifier. An EUI-48 corresponds to a MAC address while an EUI-64 is one that uses a 64-bit address space.

Extensible Authentication Protocol (EAP) Framework for negotiating authentication methods that enable systems to use hardware-based identifiers, such as fingerprint scanners or smart card readers, for authentication and to establish secure tunnels through which to submit credentials.

fiber distribution panel Type of distribution frame with pre-wired connectors used with fiber optic cabling.

fiber optic cable Network cable type that uses light signals as the basis for data transmission. Infrared light pulses are transmitted down the glass core of the fiber. The cladding that surrounds this core reflects light back to ensure transmission efficiency. At the receiving end of the cable, light-sensitive diodes re-convert the light pulse into an electrical signal. Fiber optic cable is immune to eavesdropping and EMI, has low attenuation, supports rates of 10 Gb/s+, and is light and compact.

Fibre Channel (FC) High-speed network communications protocol used to implement SANs.

File Transfer Protocol (FTP) Application protocol used to transfer files between network hosts. Variants include S(ecure) FTP, FTP with SSL (FTPS and FTPES), and T(rivial)FTP. FTP utilizes ports 20 and 21.

firewall Software or hardware device that protects a network segment or individual host by filtering packets to an access control list.

first hop redundancy protocol (FHRP) Provisioning failover routers to serve as the default gateway for a subnet.

fragmentation Mechanism for splitting a layer 3 datagram between multiple frames to fit the maximum transmission unit (MTU) of the underlying Data Link network.

frame Common term for the protocol data unit for layer 2.

frequency band Portion of the microwave radio-frequency spectrum in which wireless products operate, such as 2.4 GHz band or 5 GHz band.

F-type connector Screw down connector used with coaxial cable.

full tunnel VPN configuration where all traffic is routed via the VPN gateway.

full-duplex Network link that allows interfaces to send and receive simultaneously.

fully qualified domain name (FQDN) Unique label specified in a DNS hierarchy to identify a particular host within a subdomain within a top-level domain.

General Data Protection Regulation (GDPR) Provisions and requirements protecting the personal data of European Union (EU) citizens. Transfers of personal data outside the EU Single Market are restricted unless protected by like-for-like regulations, such as the US's Privacy Shield requirements.

Generic Routing Encapsulation (GRE) Tunneling protocol allowing the transmission of encapsulated frames or packets from different types of network protocol over an IP network.

geofencing Security control that can enforce a virtual boundary based on real-world geography.

giant Ethernet frame that is larger than the receiving interface will accept.

Global Positioning System (GPS) A means of determining a receiver's position on Earth based on information received from orbital satellites.

half-duplex Network link where simultaneously sending and receiving is not possible.

hardening A process of making a host or app configuration secure by reducing its attack surface, through running only necessary services, installing monitoring software to protect against malware and intrusions, and establishing a maintenance schedule to ensure the system is patched to be secure against software exploits.

hashing A function that converts an arbitrary-length string input to a fixed-length string output. A cryptographic hash function does this in a way that reduces the chance of collisions, where two different inputs produce the same output.

heat map In a Wi-Fi site survey, a diagram showing signal strength and channel utilization at different locations.

heating, ventilation, air conditioning (HVAC) Control systems that maintain an optimum heating, cooling, and humidity level working environment for different parts of the building.

high availability (HA) A metric that defines how closely systems approach the goal of providing data availability 100% of the time while maintaining a high level of system performance.

honeypot A host (honeypot), network (honeynet), file (honeyfile), or credential/token (honeytokens) set up with the purpose of luring attackers away from assets of actual value and/or discovering attack strategies and weaknesses in the security configuration.

hop One link in the path from a host to a router or from router to router. Each time a packet passes through a router, its hop count (or TTL) is decreased by one.

host name Label applied to a host computer that is unique on the local network.

hosts (file) List of static name to IP address mappings maintained on a host computer that will typically take precedence over name resolution queries.

hot site A fully configured alternate processing site that can be brought online either instantly or very quickly after a disaster.

HTML5 VPN Using features of HTML5 to implement remote desktop/VPN connections via browser software (clientless).

hub Layer 1 (Physical) network device used to implement a star network topology on legacy Ethernet networks, working as a multiport repeater.

hub-and-spoke Wide area network topology with the same layout as a star topology.

hybrid cloud A cloud deployment that uses both private and public elements.

hybrid topology A network that uses a combination of physical or logical topologies. In practice, most networks use hybrid topologies. For example, modern types of Ethernet are physically wired as stars but logically operate as buses.

HyperText Transfer Protocol/HTTP Secure (HTTP) Application protocol used to provide web content to browsers. HTTP uses port 80. HTTPS(ecure) provides for encrypted transfers, using TLS and port 443.

identity and access management (IAM) A security process that provides identification, authentication, and authorization mechanisms for users, computers, and other entities to work with organizational assets such as networks, operating systems, and applications.

ifconfig command Deprecated Linux command tool used to gather information about the IP configuration of the network adapter or to configure the network adapter.

implicit deny Firewall ACL rule configured by default to block any traffic not matched by previous rules.

industrial control system (ICS) Network managing embedded devices (computer systems that are designed to perform a specific, dedicated function).

infrastructure as a service (IaaS) A cloud service model that provisions virtual machines and network infrastructure.

infrastructure as code (IaC) Provisioning architecture in which deployment of resources is performed by scripted automation and orchestration.

instant secure erase (ISE) Media sanitization command built into HDDs and SSDs that are self-encrypting that works by erasing the encryption key, leaving remnants unrecoverable.

insulation displacement connector (IDC) Block used to terminate twisted pair cabling at a wall plate or patch panel available in different formats, such as 110, BIX, and Krone.

interface statistics Metrics recorded by a host or switch that enable monitoring of link state, resets, speed, duplex setting, utilization, and error rates.

intermediate distribution frame (IDF) Passive wiring panel providing a central termination point for cabling. An IDF is an optional layer of distribution frame hierarchy that cross-connects “vertical” backbone cabling to an MDF to “horizontal” wiring to wall ports on each floor of a building or each building of a campus network.

Internet Control Message Protocol (ICMP) IP-level protocol for reporting errors and status information supporting the function of troubleshooting utilities such as ping.

Internet Group Management Protocol (IGMP) Layer 3 protocol that allows hosts to join and leave groups configured to receive multicast communications.

Internet Key Exchange (IKE) Framework for creating a Security Association (SA) used with IPSec. An SA establishes that two hosts trust one another (authenticate) and agree on secure protocols and cipher suites to use to exchange data.

Internet Message Access Protocol (IMAP) Application protocol providing a means for a client to access and manage email messages stored in a mailbox on a remote server. IMAP4 utilizes TCP port number 143, while the secure version IMAPS uses TCP/993.

internet of things (IoT) Devices that can report state and configuration data and be remotely managed over IP networks.

Internet Protocol header (IP header) Fields in a datagram used to identify source and destination IP addresses, protocol type, and other layer 3 properties.

Internet Protocol Security (IPSec) Network protocol suite used to secure

data through authentication and encryption as the data travels across the network or the Internet.

Internet Service Provider (ISP)

Provides Internet connectivity and web services to its customers.

intrusion detection system (IDS)

A security appliance or software that analyzes data from a packet sniffer to identify traffic that violates policies or rules.

intrusion prevention system (IPS)

A security appliance or software that combines detection capabilities with functions that can actively block attacks.

IP address management (IPAM)

Software consolidating management of multiple DHCP and DNS services to provide oversight into IP address allocation across an enterprise network.

ip command Linux command tool used to gather information about the IP configuration of the network adapter or to configure the network adapter.

IP helper Command set in a router OS to support DHCP relay and other broadcast forwarding functionality.

IP protocol type Identifier for a protocol working over the Internet Protocol, such as TCP, UDP, ICMP, GRE, EIGRP, or OSPF.

IP scanner Utility that can probe a network to detect which IP addresses are in use by hosts.

ipconfig command Command tool used to gather information about the IP configuration of a Windows host.

iperf Utility used to measure the bandwidth achievable over a network link.

iterative lookup DNS query type whereby a server responds with information from its own data store only.

jitter Variation in the time it takes for a signal to reach the recipient. Jitter manifests itself as an inconsistent rate of packet delivery. If packet loss or delay is excessive, then noticeable audio or video problems (artifacts) are experienced by users.

jumbo frame Ethernet frame with a payload larger than 1,500 bytes (up to 9,216 bytes).

jump server A hardened server that provides access to other hosts.

Kerberos A single sign-on authentication and authorization service that is based on a time-sensitive, ticket-granting system.

latency Time taken for a signal to reach the recipient, measured in milliseconds. Latency is a particular problem for two-way applications, such as VoIP (telephone) and online conferencing.

layer 3 switch Switch appliance capable of IP routing between virtual LAN (VLAN) subnets using hardware-optimized path selection and forwarding.

least privilege A basic principle of security stating that something should be allocated the minimum necessary rights, privileges, or information to perform its role.

lifecycle roadmap Method to track the life cycle phases of one or more hardware, service, or software systems in your organization.

lightweight AP Access point that requires a wireless controller in order to function.

Lightweight Directory Access Protocol (LDAP) Protocol used to access network directory databases, which store information about authorized users and their privileges, as well as other organizational information.

link aggregation Combining the bandwidth of two or more switch ports into a single channel link.

Link Aggregation Control Protocol (LACP) IEEE protocol governing the use of bonded Ethernet ports (NIC teaming).

Link Layer Discovery Protocol (LLDP) Standards-based protocol used by network appliances to discover layer 2 adjacent devices or neighbors.

link local IP addressing scheme used within the scope of a single broadcast domain only.

link state Algorithm used by routing protocols that build a complete network

topology to use to select optimum forwarding paths.

load balancer A type of switch, router, or software that distributes client requests between different resources, such as communications links or similarly configured servers. This provides fault tolerance and improves throughput.

local area network (LAN) Network scope restricted to a single geographic location and owned/managed by a single organization.

local connector (LC) Small form factor push-pull fiber optic connector; available in simplex and duplex versions.

logging level Threshold for storing or forwarding an event message based on its severity index or value.

Long Term Evolution (LTE) Packet data communications specification providing an upgrade path for 2G and 3G cellular networks. LTE services use a SIM card to identify the subscriber and network provider. LTE Advanced is designed to provide 4G standard network access.

loopback address IP address by which a host can address itself over any available interface.

main distribution frame (MDF) Passive wiring panel providing a central termination point for cabling. A MDF distributes backbone or “vertical” wiring through a building and connections to external access provider networks.

malware Software that serves a malicious purpose, typically installed without the user’s consent (or knowledge).

management information base (MiB) Database that stores Simple Network Management Protocol (SNMP) properties and values of a network device and its components.

maximum tolerable downtime (MTD) The longest period that a process can be inoperable without causing irrevocable business failure.

maximum transmission unit (MTU) Maximum size in bytes of a frame’s payload. If the payload cannot be encapsulated within a single frame

at the Data Link layer, it must be fragmented.

mean time between failures (MTBF) A metric for a device or component that predicts the expected time between failures.

mean time to failure (MTTF) Metric indicating average time a device or component is expected to be in operation.

mean time to repair/replace/recover (MTTR) A metric representing average time taken for a device or component to be repaired, replaced, or otherwise recover from a failure.

media access control (MAC) Hardware address that uniquely identifies each network interface at layer 2 (Data Link). A MAC address is 48 bits long with the first half representing the manufacturer’s organizationally unique identifier (OUI).

Media Access Control address table (MAC) Data store on a switch that keeps track of the MAC addresses associated with each port. As the switch uses a type of memory called content addressable memory (CAM), this is sometimes called the CAM table.

Media Access Control filtering Applying an access control list to a switch or access point so that only clients with approved MAC addresses can connect to it.

Media Access Control flooding (MAC flooding) Network attack where a switch’s cache table is inundated with frames from random source MAC addresses so that it starts flooding unicast traffic, facilitating snooping attacks.

media converter Layer 1 (Physical) network device that translates signals received over one media type for transmission over a different media type.

medium dependent interface/medium dependent interface crossover (MDI/MDI-X) System that distinguishes transmit and receive pins on different interface types. The interface on an end system is MDI while that on an intermediate system is MDI-X.

memorandum of understanding

(MoU) Usually a preliminary or exploratory agreement to express an intent to work together that is not legally binding and does not involve the exchange of money.

mesh topology A topology often used in WANs where each device has (in theory) a point-to-point connection with every other device (fully connected); in practice, only the more important devices are directly interconnected (partial mesh).

microsegmentation (switching)

Function of an Ethernet switch whereby collision domains are reduced to the scope of a single port only.

missing route Troubleshooting issue where a routing table does not contain a required entry due either to manual misconfiguration or failure of a dynamic routing protocol update.

mission essential function (MEF)

Business or organizational activity that is too critical to be deferred for anything more than a few hours, if at all.

multicast A packet addressed to a selection of hosts (in IP, those belonging to a multicast group).

multifactor authentication (MFA) An authentication scheme that requires the user to present at least two different factors as credentials; for example, something you know, something you have, something you are, something you do, and somewhere you are. Specifying two factors is known as "2FA."

multi-fiber push-on (MPO) Fiber optic cable type that terminates multiple strands to a single compact connector, supporting parallel links.

multimode fiber (MMF) Fiber optic cable type using LED or vertical cavity surface emitting laser optics and graded using optical multimode types for core size and bandwidth.

multiple input multiple output

(MIMO) Use of multiple reception and transmission antennas to boost wireless bandwidth via spatial multiplexing and to boost range and signal reliability via spatial diversity.

multiuser MIMO (MU-MIMO) Use of spatial multiplexing to allow a wireless access point to support multiple client stations simultaneously.

NAT64 IPv6 transition mechanism that uses Network Address Translation (NAT) to convert destination IPv4 addresses to IPv6 format at routing boundaries.

native VLAN VLAN ID used for any untagged frames received on a trunk port. The same ID should be used on both ends of the trunk and the ID should not be left as the default VLAN ID (1).

Neighbor Discovery Protocol (ND)

IPv6 protocol used to identify link local nodes.

NetFlow Cisco-developed means of reporting network flow information to a structured database. NetFlow allows better understanding of IP traffic flows as used by different network applications and hosts.

netstat command Cross-platform command tool to show network information on a machine running TCP/IP, notably active connections, and the routing table.

network access control (NAC) A general term for the collected protocols, policies, and hardware that authenticate and authorize access to a network at the device level.

Network Address Translation (NAT)

Routing mechanism that conceals internal addressing schemes from the public Internet by translating between a single public address on the external side of a router and private, non-routable addresses internally.

network attached storage (NAS)

Storage device enclosure with network port and an embedded OS that supports typical network file access protocols (FTP and SMB for instance).

network discovery Processes and tools that facilitate identification of hosts present on a network or subnet.

Network Functions Virtualization

(NFV) Provisioning virtual network appliances, such as switches, routers, and firewalls, via VMs and containers.

network interface card (NIC) Adapter card that provides one or more Ethernet ports for connecting hosts to a network so that they can exchange data over a link.

Network layer OSI model layer responsible for logical network addressing and forwarding.

network loop Troubleshooting issue where layer 2 frames are forwarded between switches or bridges in an endless loop.

network mask Number of bits applied to an IP address to mask the network ID portion from the host/interface ID portion. This can be expressed as a bit prefix in slash notation or as a dotted decimal subnet mask.

network security group Rules that filter communication between cloud networks and from cloud networks to the Internet.

network security list In Oracle Cloud Infrastructure, traffic filtering rules that apply to a subnet, rather than just network interfaces.

network separation Enforcing a security zone by separating a segment of the network from access by the rest of the network. This could be accomplished using firewalls or VPNs or VLANs. A physically separate network or host (with no cabling or wireless links to other networks) is referred to as air-gapped.

Network Time Protocol (NTP) Application protocol allowing machines to synchronize to the same time clock that runs over UDP port 123.

Network Time Security (NTS) Method of securing NTP queries and responses using Transport Layer Security (TLS). NTS typically uses TCP port 3443.

NIC teaming Two or more NIC aggregated into a single channel link for fault tolerance and increased throughput. Also known as NIC bonding.

Nmap An IP and port scanner used for topology, host, service, and OS discovery and enumeration.

nondisclosure agreement (NDA) An agreement that stipulates that

entities will not share confidential information, knowledge, or materials with unauthorized third parties.

north-south traffic Network data flows that go into and out of an organization's network or datacenter.

nslookup command Cross-platform command tool for querying DNS resource records.

on-path attack An attack where the threat actor makes an independent connection between two victims and is able to read and possibly modify traffic.

open authentication Wireless network authentication mode where guest (unauthenticated) access is permitted.

Open Shortest Path First (OSPF) Dynamic routing protocol that uses a link state algorithm and a hierarchical topology.

Open Systems Interconnection reference model (OSI) Assigns network and hardware components and functions at seven discrete layers: Physical, Data Link, Network, Transport, Session, Presentation, and Application.

operational technology (OT) A communications network designed to implement an industrial control system rather than data networking.

optical link budget Assessment of allowable signal loss over a fiber optic link.

optical multimode (OM) Classification system for multimode fiber designating core size and modal bandwidth.

option (DCHP) DHCP configuration that assigns additional parameters, such as DNS server addresses. In DHCPv4, an option is used to identify the default gateway address.

orchestration Automation of multiple coordinated steps in a deployment process.

out of band management (OOB) Accessing the administrative interface of a network appliance using a separate network from the usual data network. This could use a separate VLAN or a different kind of link, such as a dial-up modem.

overlay network Network protocols that use encapsulation to provision virtual tunnels and networks without requiring reconfiguration of the underlying transport network.

packet loss Network PDUs that do not reach their destination due to transmission errors, congestion, or security policies. A packet drop or discard is where a switch or router does not forward a packet due to congestion or because the packet does not match the requirements of an ACL.

packet sniffer A monitor that records (or “sniffs”) data from frames as they pass over network media, using methods such as a mirror port or TAP device.

patch A small unit of supplemental code meant to address either a security problem or a functionality flaw in a software package or operating system.

patch panel Type of distribution frame used with twisted pair cabling with IDCs to terminate fixed cabling on one side and modular jacks to make cross-connections to other equipment on the other.

Payment Card Industry Data Security Standard (PCI DSS) The information security standard for organizations that process credit or bank card payments.

peer-to-peer Administration paradigm whereby any computer device may be configured to operate as both server and client.

performance metric Measurement of a value affecting system performance, such as CPU or memory utilization.

personally identifiable information (PII) Data that can be used to identify or contact an individual (or, in the case of identity theft, to impersonate them).

phishing A email-based social engineering attack in which the attacker sends email from a supposedly reputable source, such as a bank, to try to elicit private information from the victim.

Physical layer (PHY) Lowest layer of the OSI model providing for the

transmission and receipt of data bits from node to node. This includes the network medium and mechanical and electrical specifications for using the media.

ping command Cross-platform command tool for testing IP packet transmission.

platform as a service (PaaS) A cloud service model that provisions application and database services as a platform for development of apps.

playbook A checklist of actions to perform to complete a standard procedure or detect and respond to a specific type of incident.

plenum Cable for use in building voids designed to be fire-resistant and to produce a minimal amount of smoke if burned.

Point to Point Protocol (PPP) Dial-up protocol working at layer 2 (Data Link) used to connect devices remotely to networks.

point-to-point A point-to-point topology is one where two nodes have a dedicated connection to one another.

polarization Orientation of the wave propagating from an antenna.

port (TCP/UDP) In TCP and UDP applications, a unique number assigned to a particular application protocol. Server ports are typically assigned well-known or registered numbers, while client ports use dynamic or ephemeral numbering.

port address translation (PAT) Maps private host IP addresses onto a single public IP address. Each host is tracked by assigning it a random high TCP port for communications.

port mirroring Copying ingress and/or egress communications from one or more switch ports to another port. This is used to monitor communications passing over the switch.

port role In Spanning Tree Protocol (STP), each port is assigned a role (root, designated, blocked, or disabled) depending on its position in the topology.

port scanner Utility that can probe a host to enumerate the status of TCP and UDP ports.

port security Preventing a device attached to a switch port from communicating on the network unless it matches a given MAC address or other protection profile.

port state In Spanning Tree Protocol (STP), topology changes cause ports to transition through different states (blocking, listening, learning, forwarding, and disabled).

port tagging On a switch with VLANs configured, a port with an end station host connected operates in untagged mode (access port). A tagged port will normally be part of a trunk link.

port-side exhaust/intake Feature of switches that allows fans to switch between expelling hot air and drawing in cool air from the side with ports.

posture assessment Audit process and tools for verifying compliance with a compliance framework or configuration baseline.

power budget When configuring Power over Ethernet, the maximum amount of power available across all switchports.

power distribution unit (PDU) An advanced strip socket that provides filtered output voltage. A managed unit supports remote administration.

Power over Ethernet (PoE) Specification allowing power to be supplied via switch ports and ordinary data cabling to devices such as VoIP handsets and wireless access points. Devices can draw up to about 13 W (or 25 W for PoE+).

Precision Time Protocol (PTP) Provides clock synchronization to network devices to a higher degree of accuracy than Network Time Protocol (NTP).

Presentation layer OSI model layer that transforms data between the formats used by the network and applications.

pre-shared key (PSK) A wireless network authentication mode where a passphrase-based mechanism is used to allow group authentication to a wireless

network. The passphrase is used to derive an encryption key.

private branch exchange (PBX) Routes incoming calls to direct dial numbers and provides facilities such as voice mail, Automatic Call Distribution (ACD), and interactive voice response (IVR). A PBX can also be implemented as software (virtual PBX). An IP-based PBX or hybrid PBX allows use of VoIP.

private cloud A cloud that is deployed for use by a single entity.

private key In asymmetric encryption, the private key is known only to the holder and is linked to, but not derivable from, a public key distributed to those with whom the holder wants to communicate securely. A private key can be used to encrypt data that can be decrypted by the linked public key or vice versa.

production configuration Configuration settings used when an appliance, instance, or app is booted or started.

protocol analyzer Utility that can parse the header fields and payloads of protocols in captured frames for display and analysis.

protocol data unit (PDU) Network packet encapsulating a data payload from an upper layer protocol with header fields used at the current layer.

proxy server A server that mediates the communications between a client and another server. It can filter and often modify communications as well as provide caching services to improve performance.

public cloud A cloud that is deployed for shared use by multiple independent tenants.

public key During asymmetric encryption, this key is freely distributed and can be used to perform the reverse encryption or decryption operation of the linked private key in the pair.

public key infrastructure (PKI) A framework of certificate authorities, digital certificates, software, services, and other cryptographic components deployed for the purpose of validating subject identities.

public switched telephone network (PSTN) Global network connecting national telecommunications systems.

public versus private addressing Some IP address ranges are designated for use on private networks only. Packets with source IP addresses in public ranges are permitted to be forwarded over the Internet. Packets with source IP addresses from private ranges should be blocked at Internet gateways or forwarded using some type of translation mechanism.

punchdown tool Tool used to terminate solid twisted-pair copper cable to an insulation displacement connector block.

quad small form factor pluggable/enhanced quad small form factor pluggable (QSFP/QSFP+) Fiber optic transceiver module type supporting four individual duplex lanes at 1 Gbps (QSFP) or 10 Gbps (QSFP+) that can be aggregated into a single 4 Gbps or 40 Gbps channel.

quality of service (QoS) Systems that differentiate data passing over the network that can reserve bandwidth for particular applications. A system that cannot guarantee a level of available bandwidth is often described as class of service (CoS).

rack Storage solution for server and network equipment. Racks are designed to a standard width and height (measured in multiples of 1U or 1.75"). Racks offer better density, cooling, and security than ordinary office furniture.

rack diagram Physical plan of appliances installed in a network rack and their power and network connections.

radio frequency attenuation (RF) Loss of signal strength due to distance and environmental factors.

received signal strength indicator (RSSI) Signal strength as measured at the receiver, using either decibel units or an index value.

recovery point objective (RPO) The longest period that an organization can tolerate lost data being unrecoverable.

recovery time objective (RTO) The maximum time allowed to restore a system after a failure event.

recursive lookup DNS query type whereby a server submits additional queries to other servers to obtain the requested information.

registered-jack connector (RJ) Series of jack/plug types used with twisted-pair cabling, such as RJ45 and RJ11.

Remote Authentication Dial-In User Service (RADIUS) AAA protocol used to manage remote and wireless authentication infrastructures.

Remote Desktop Protocol (RDP) Application protocol for operating remote connections to a host using a graphical interface. The protocol sends screen data from the remote host to the client and transfers mouse and keyboard input from the client to the remote host. It uses TCP port 3389.

repeater Layer 1 device that regenerates and retransmits signals to overcome media distance limitations.

reservation (DHCP) DHCP configuration that assigns either a prereserved or persistent IP address to a given host, based on its hardware address or other ID.

resource record (AAAA) Data file storing information about a DNS zone. The main records are as follows: A (maps a host name to an IPv4 address), AAAA (maps to an IPv6 address), CNAME (an alias for a host name), MX (the IP address of a mail server), and PTR (allows a host name to be identified from an IP address).

reverse DNS DNS query type to resolve an IP address to a host name.

RFC1918 Standards document that defines private address ranges.

risk Likelihood and impact (or consequence) of a threat actor exercising a vulnerability.

roaming WLAN configured with multiple access points in an extended service set allowing clients to remain connected to the network within an extended service area.

rogue access point Wireless access point that has been enabled on the network without authorization.

role-based access control (RBAC) An access control model where resources are protected by ACLs that are managed by administrators and that provide user permissions based on job functions.

root bridge selection In Spanning Tree Protocol (STP), the process and metrics that determine which bridge or switch will be identified as root. Selection of an inappropriate root device can cause performance and security issues.

route command Cross-platform command tools used display and manage the routing table on a Windows or Linux host.

router Intermediate system working at the Network layer capable of forwarding packets around logical networks of different layer 1 and layer 2 types.

router advertisement (RA) Packet sent by an IPv6-capable router to notify hosts about prefixes and autoconfiguration methods available on the local link

Router Advertisement Guard (RA) Switchport security feature to block router advertisement packets from unauthorized sources.

Routing Information Protocol (RIP) Distance vector-based routing protocol that uses a hop count to determine the least-cost path to a destination network.

routing loop Troubleshooting issue where a packet is forwarded between routers in a loop until its TTL expires.

routing table Data store on an IP host used to determine the interface over which to forward a packet.

runt Malformed Ethernet frame that is smaller than the permitted 64 byte minimum size.

sanitization The process of thoroughly and completely removing data from a storage medium so that file remnants cannot be recovered.

satellite System of microwave transmissions where orbital satellites relay signals between terrestrial receivers or other orbital satellites.

Satellite internet connectivity is enabled through a reception antenna connected to the PC or network through a DVB-S modem.

scalability Property by which a computing environment is able to gracefully fulfill its ever-increasing resource needs.

scope (DHCP) Range of consecutive IP addresses in the same subnet that a DHCP server can lease to clients.

screened subnet A segment isolated from the rest of a private network by one or more firewalls that accepts connections from the Internet over designated ports.

Secure Access Service Edge (SASE) A networking and security architecture that provides secure access to cloud applications and services while reducing complexity. It combines security services like firewalls, identity and access management, and secure web gateway with networking services such as SD-WAN.

secure erase (SE) Method of sanitizing a drive using the ATA command set.

Secure Shell (SSH) Application protocol supporting secure tunneling and remote terminal emulation and file copy. SSH runs over TCP port 22.

Security Assertion Markup Language (SAML) An XML-based data format used to exchange authentication information between a client and a service.

Security Information and Event Management (SIEM) A solution that provides real-time or near-real-time analysis of security alerts generated by network hardware and applications.

Security Service Edge (SSE) Design paradigm and associated technologies that mediate access to cloud services and web applications.

self-signed certificate A digital certificate that has been signed by the entity that issued it, rather than by a CA.

separation of duties Security policy concept that states that duties and responsibilities should be divided among individuals to prevent ethical conflicts or abuse of powers.

Server Message Block (SMB)

Application protocol used for requesting files from Windows servers and delivering them to clients. SMB allows machines to share files and printers, thus making them available for other machines to use. SMB client software is available for UNIX-based systems. Samba software allows UNIX and Linux servers or NAS appliances to run SMB services for Windows clients.

service level agreement (SLA) An agreement that sets the service requirements and expectations between a consumer and a provider.

Service Set Identifier (SSID) A character string that identifies a particular wireless LAN (WLAN).

Session Initiation Protocol (SIP)

Application protocol used to establish, disestablish, and manage VoIP and conferencing communications sessions. It handles user discovery (locating a user on the network), availability advertising (whether a user is prepared to receive calls), negotiating session parameters (such as use of audio/ video), and session management and termination.

Session layer OSI model layer that provides services for applications that need to exchange multiple messages (dialog control).

shadow IT Computer hardware, software, or services used on a private network without authorization from the system owner.

shellcode A lightweight block of malicious code that exploits a software vulnerability to gain initial access to a victim system.

shielded twisted pair (STP) Copper twisted-pair cabling with screening and shielding elements for individual wire pairs and/or the whole cable to reduce interference.

shoulder surfing Social engineering tactic to obtain someone's password or PIN by observing them as they type it in.

show arp command Command tools used in router operating systems to list the contents of the Address Resolution Protocol (ARP) cache of IP address to MAC address mappings.

show commands Set of commands in a switch OS to report configuration or interface information.

show route command Command tools used in router operating systems to list the contents of routing tables.

Simple Mail Transfer Protocol (SMTP)

Application protocol used to send mail between hosts on the Internet. Messages are sent between servers over TCP port 25 or submitted by a mail client over secure port TCP/587.

Simple Network Management

Protocol (SNMP) Application protocol used for monitoring and managing network devices. SNMP works over UDP ports 161 and 162 by default.

Simultaneous Authentication of

Equals (SAE) Personal authentication mechanism for Wi-Fi networks introduced with WPA3 to address vulnerabilities in the WPA-PSK method.

single sign-on (SSO) Authentication technology that enables a user to authenticate once and receive authorizations for multiple services.

single-mode fiber (SMF) Fiber optic cable type that uses laser diodes and narrow core construction to support high bandwidths over distances of more than five kilometers.

site survey Documentation about a location for the purposes of building an ideal wireless infrastructure; it often contains optimum locations for wireless antenna and access point placement to provide the required coverage for clients and to identify sources of interference.

small form factor pluggable/enhanced small form factor pluggable (SFP/SFP+)

Fiber optic transceiver module type supporting duplex 1 Gbps (SFP) or 10 Gbps (SFP+) links.

small office, home office (SOHO)

Category of network type and products that are used to implement small-scale LANs and off-the-shelf Internet connection types.

social engineering An activity where the goal is to use deception and trickery to convince unsuspecting users to provide sensitive data or to violate security guidelines.

socket Combination of a TCP/UDP port number and IP address. A client socket can form a connection with a server socket to exchange data.

software as a service (SaaS) A cloud service model that provisions fully developed application services to users.

software defined WAN (SD-WAN) Services that use software-defined mechanisms and routing policies to implement virtual tunnels and overlay networks over multiple types of transport network.

software-defined networking (SDN) APIs and compatible hardware/virtual appliances allowing for programmable network appliances and systems.

source control Technologies that manage development of software code by tracking and merging or rejecting changes from multiple authors.

Spanning Tree Protocol (STP) Protocol that prevents layer 2 network loops by dynamically blocking switch ports as needed.

spectrum analyzer Device that can detect the source of interference on a wireless network.

spine and leaf topology Topology commonly used in datacenters comprising a top tier of aggregation switches forming a backbone for a leaf tier of top-of-rack switches.

split tunnel VPN configuration where only traffic for the private network is routed via the VPN gateway.

spoofing Attack technique where the threat actor disguises his or her identity or impersonates another user or resource.

standard operating procedure (SOP) Documentation of best practice and work instructions to use to perform a common administrative task.

star topology In a star network, each node is connected to a central point, typically a switch or a router. The central point mediates communications between the attached nodes. When a device such as a hub is used, the hub receives signals from a node and repeats the signal to all other connected

nodes. Therefore the bandwidth is still shared between all nodes. When a device such as a switch is used, point-to-point links are established between each node as required. The circuit established between the two nodes can use the full bandwidth capacity of the network media.

stateless address autoconfiguration (SLAAC) Mechanism used in IPv6 for hosts to assign addresses to interfaces without requiring manual intervention.

static route Entry in the routing table added manually by an administrator.

storage area network (SAN) Network dedicated to provisioning storage resources, typically consisting of storage devices and servers connected to switches via host bus adapters.

straight-tip connector (ST) Bayonet-style twist-and-lock connector for fiber optic cabling.

Structured Query Language (SQL) Programming and query language common to many relational database management systems.

subinterface Configuring a router's physical interface with multiple virtual interfaces connected to separate virtual LAN (VLAN) IDs over a trunk.

subnet addressing Division of a single IP network into two or more smaller broadcast domains by using longer netmasks within the boundaries of the network.

subscriber connector (SC) Push/pull connector used with fiber optic cabling.

supervisory control and data acquisition (SCADA) A type of industrial control system that manages large-scale, multiple-site devices and equipment spread over geographically large areas from a host computer.

switch Intermediate system used to establish contention-free network segments at OSI layer 2 (Data Link). An unmanaged switch does not support any sort of configuration.

Switch Virtual Interface (SVI) Feature of layer 3 switches that allows a virtual interface assigned with an IP address to act as the default gateway for a VLAN.

syslog Application protocol and event-logging format enabling different appliances and software applications to transmit logs or event records to a central server. Syslog works over UDP port 514 by default.

T568A/T568B (T568A) Twisted-pair termination pinouts defined in the ANSI/TIA/EIA 568 Commercial Building Telecommunications Standards.

tabletop exercise A discussion of simulated emergency situations and security incidents.

tailgating Social engineering technique in which a person gains access to a building by following someone who is unaware of his or her presence.

TCP flag Field in the header of a TCP segment designating the connection state, such as SYN, ACK, or FIN.

tcpdump command A command line packet sniffing utility.

telnet Application protocol supporting unsecure terminal emulation for remote host management. Telnet runs over TCP port 23.

Terminal Access Controller Access Control System Plus (TACACS+)

AAA protocol developed by Cisco that is often used to authenticate to administrator accounts for network appliance management.

test access point (TAP) A hardware device inserted into a cable run to copy frames for analysis.

threat A potential for an entity to exercise a vulnerability (that is, to breach security).

three-tier hierarchal model Paradigm to simplify network design by separating switch and router functionality and placement into three tiers each with a separate role, performance requirements, and physical topology.

throughput Amount of data transfer supported by a link in typical conditions. This can be measured in various ways with different software applications. Goodput is typically used to refer to the actual “useful” data rate at the application layer (less overhead from headers and lost packets).

time to live (DNS) (TTL) Amount of time that the record returned by a DNS query should be cached before discarding it.

Time to Live (IP) (TTL) Counter field in the IP header recording the number of hops a packet can make before being dropped.

tone generator Two-part tool used to identify one cable within a bundle by applying an audible signal.

top-of-rack switch (ToR) High-performance switch model designed to implement the leaf tier in a spine and leaf topology.

topology Network specification that determines the network’s overall layout, signaling, and dataflow patterns.

traceroute/tracert command

Diagnostic utilities that trace the route taken by a packet as it “hops” to the destination host on a remote network. tracert is the Windows implementation, while traceroute runs on Linux.

traffic analysis Processes and tools that facilitate reporting of network communication flows summarized by host or protocol type.

traffic shaper Appliances and/or software that enable administrators to closely monitor network traffic and to manage that network traffic. The primary function of a traffic shaper is to optimize network media throughput to get the most from the available bandwidth.

transceiver Component in a network interface that converts data to and from the media signalling type. Modular transceivers are designed to plug into switches and routers.

Transmission Control Protocol (TCP)

Protocol in the TCP/IP suite operating at the Transport layer to provide connection-oriented, guaranteed delivery of packets.

Transport layer OSI model layer responsible for ensuring reliable data delivery.

Transport Layer Security (TLS)

Security protocol that uses certificates for authentication and encryption to protect web communications and other application protocols.

Trivial File Transfer Protocol (TFTP)

Simplified form of FTP supporting only file copying. TFTP works over UDP port 69.

troubleshooting methodology

Structured approach to problem-solving using identification, theory of cause, testing, planning, implementation, verification, and documentation steps.

trunk Backbone link established between switches and routers to transport frames for multiple virtual LANs (VLANs).

tunneling Encapsulating data from a local protocol within another protocol's PDU to transport it to a remote network over an intermediate network. Tunneling protocols are used in many contexts, including virtual private networks (VPNs) and transport IPv6 packets over IPv4 networks.

twinaxial Media type similar to coax but with two inner conductors to improve performance.

twisted pair cable Network cable construction with insulated copper wires twisted about each other. A pair of color-coded wires transmits a balanced electrical signal. The twisting of the wire pairs at different rates acts to reduce interference and crosstalk.

ultra physical contact (UPC) Fiber optic connector finishing type that uses a slightly curved polish for the ferrule.

unicast A packet addressed to a single host. If the host is not on the local subnet, the packet must be sent via one or more routers.

uninterruptible power supply (UPS) A battery-powered device that supplies AC power that an electronic device can use in the event of power failure.

unshielded twisted pair (UTP) Media type that uses copper conductors arranged in pairs that are twisted to reduce interference. Typically, cables are 4-pair or 2-pair.

URL filtering Type of content filter applied to restrict client queries to particular uniform resource locator (URL) web addresses.

User Datagram Protocol (UDP)

Protocol in the TCP/IP suite operating at the transport layer to provide connectionless, non-guaranteed communication.

variable length subnet masking (VLSM) Using network prefixes of different lengths within an IP network to create subnets of different sizes.

version control Within a source control system, a process that assigns an identification number to each release of an app or script.

virtual appliance A preconfigured, self-contained virtual machine image ready to be deployed and run on a hypervisor.

Virtual Extensible LAN (VXLAN)

Technology used to implement an overlay network so that hosts in separate subnets can establish layer 2 adjacency in a discrete logical segment. The 24-bit VXLAN ID space supports up to 16 million logical segments.

virtual IP Public address of a load balanced cluster that is shared by the devices implementing the cluster.

virtual local area network (VLAN) A logical network segment comprising a broadcast domain established using a feature of managed switches to assign each port a VLAN ID. Even though hosts on two VLANs may be physically connected to the same switch, local traffic is isolated to each VLAN, so they must use a router to communicate.

virtual private cloud (VPC) A private network segment made available to a single cloud consumer on a public cloud.

virtual private network (VPN) A secure tunnel created between two endpoints connected via an unsecure transport network (typically the Internet).

visual fault locator Troubleshooting tool used to identify breaks or imperfections in fiber optic cable.

VLAN hopping Exploiting a misconfiguration to direct traffic to a different VLAN without authorization.

Voice over Internet Protocol (VoIP) A generic name for protocols that carry voice traffic over data networks.

voice virtual local area network (VLAN)

Feature of VoIP handsets and switches to segregate data and voice traffic while using a single network wall port to attach the handset and the computer.

VoIP phone Handset or software client that implements a type of voice over Internet Protocol (VoIP) to allow a user to place and receive calls.

vulnerability A weakness that could be triggered accidentally or exploited intentionally to cause a security breach.

vulnerability assessment Evaluation of a system's security and ability to meet compliance requirements based on the configuration state of the system, as represented by information collected from the system.

warm site An alternate processing location that is dormant or performs noncritical functions under normal conditions, but which can be rapidly converted to a key operations site if needed.

wide area network (WAN) Network scope that spans a large geographical area, incorporating more than one site and often a mix of different media types and protocols plus the use of public telecommunications networks.

Wi-Fi analyzer Device or software that can report characteristics of a WLAN, such as signal strength and channel utilization.

Wi-Fi Protected Access (WPA)

Standards for authenticating and encrypting access to Wi-Fi networks.

wire map tester Tool to verify termination/pinouts of cable.

wireless controller Device that provides wireless LAN management for multiple APs.

wireless local area network (WLAN)

Network scope and type that uses wireless radio communications based on some variant of the 802.11 (Wi-Fi) standard series.

wireless mesh network (WMN)

Wireless network topology where all nodes—including client stations—are capable of providing forwarding and path discovery. This improves coverage and throughput compared to using just fixed access points and extenders.

Wireshark A widely used protocol analyzer.

wiring diagram Documentation of connector pinouts.

work recovery time (WRT) In disaster recovery, time additional to the RTO of individual systems to perform reintegration and testing of a restored or upgraded system following an event.

YAML Ain't Markup Language (YAML)

Language for configuration files and applications such as Netplan and Ansible.

zero trust architecture (ZTA) The security design paradigm where any request (host-to-host or container-to-container) must be authenticated before being allowed.

zero-day A vulnerability in software that is unpatched by the developer or an attack that exploits such a vulnerability.

zone index Parameter assigned by a host to distinguish ambiguous interface addresses within a link local scope.

zone transfer Mechanism by which a secondary name server obtains a read-only copy of zone records from the primary server.

Index

Numbers

0000::/8 block, [146](#)
2.4 GHz channel bandwidth, [381](#), [382–383](#)
3CX, [251](#)
4G cellular standard, [371](#)
4G standard, [386](#)
5 GHz channel bandwidth, [381–382](#)
5G standard, [386](#)
6to4 tunneling, [145](#)
8-bit value, [22](#), [107](#), [140](#)
8P8C connectors, [41](#)
10BASE-T Ethernet, [35](#), [36–37](#)
10-degree antenna, [394](#)
10/forty Gigabit Ethernet (10 GbE), [37](#)
10GBASE-CR, [44](#)
10 Gigabit Ethernet (10 GbE), [37](#), [38](#)
32-bit IPv4 addressing, [114](#)
40GBASE-CR4, [44](#)
/48s, [145](#)
90-degree antenna, [394](#)
100BASE-TX Fast Ethernet standards, [36–37](#)
802.1D standard, [93](#)
802.1p standard, [G-1](#), [303](#)
802.1Q, [G-1](#)
802.1Q standard, [186](#), [187](#), [188](#), [322](#)
802.1Q VLAN, [303](#)
802.1X standard, [G-1](#), [352–353](#), [398](#)
802.3af standard, [93](#)
802.3at (PoE+) standard, [93](#), [101](#)
802.3bt (PoE++) standard, [93](#), [101](#)
802.3 standard, [35](#), [381](#)
802.11 standards, [G-1](#), [381](#)
 802.11a, [381–382](#), [383](#)
 802.11ac, [384](#)
 802.11ax, [384](#), [397](#)
 802.11b, [382–383](#)
 802.11g, [383](#)
 802.11h, [G-1](#), [382](#)

802.11k, [408](#)
802.11n, [383–384](#)
802.11r, [392](#), [408](#)
802.11v, [408](#)
802.11w, [402](#)
1000BASE-T Ethernet, [37](#)
1588 standard, [239](#)

A

A/AAAA records, [G-17](#), [223](#), [224](#), [249](#)
AAA clients or authenticators, [341](#), [399](#)
absorption, [407](#)
AC5300, [384](#)
access control, [332–333](#)
 discretionary access control, [343](#)
 identity and access management, [332–333](#)
 Lightweight Directory Access Protocol, [345–346](#)
 privileged access management, [344–345](#)
 role-based access control, [343–344](#)
access control list (ACL), [G-1](#), [12](#), [173](#), [176–177](#)
 authorization and role-based access control, [343](#)
 configuration, security rules and, [355–357](#)
 content filtering, [358–359](#)
 firewalls, [173](#), [176–177](#), [356–357](#)
 issues, [360](#)
access denied issues, [360](#)
access/edge layer, [G-1](#), [181](#)
access point (AP), [G-1](#), [381](#), [388–389](#), [391–392](#), [393–394](#)
access port, [187](#)
accounting, IAM, [333](#)
acknowledgment (ACK), [380](#)
active-active clustering, [G-1](#), [261](#)
Active Directory, [339](#), [345](#)
active FTP, [242–243](#)
active-passive, [G-1](#)
active-passive clustering, [261](#)
active TAP, [294](#)
address autoconfiguration, IPv6, [212](#)
addressing, defined, [G-1](#), [8](#). *see also* IP addressing
address pool exhaustion, [217](#)
address record (A or AAAA), [223](#), [224](#), [249](#)
Address Resolution Protocol (ARP), [G-1](#), [G-19](#), [109](#)
 tool, [135–136](#)
address scheme design, IPv4, [126–127](#)
ad hoc network, [G-1](#), [395](#)
adjacent channel interference (ACI), [G-1](#), [406](#)
adjacent layer interaction, [8](#)
administrative distance (AD), [G-1](#), [169–170](#)
administratively down, [G-1](#)
Advanced Encryption Standard (AES), [397](#)
advanced persistent threat (APT), [G-1](#), [319](#)
advanced volatile threat (AVT), [319](#)
agents, SNMP, [283–284](#)
agreements, [275](#)
AIC triad. *see* confidentiality, integrity, and availability (CIA)
Aircrack, [402](#)
Aireplay, [402](#)
air-gapped network, [G-14](#)
Airodump, [402](#)
alerting, [289–290](#)
alerts *versus* notifications, [289](#)
alien crosstalk, [69](#), [70](#)
alternating current (AC) voltage, [62](#)
Amazon DynamoDB, [247](#)
Amazon Elastic Compute Cloud, [440](#), [441](#), [445](#)
Amazon Web Services (AWS), [341](#), [444](#), [445](#), [447](#), [448](#)
American National Standards Institute (ANSI), [40](#), [437](#)

- American Standard Code for Information Interchange (ASCII), 14, 295, 328, 388
 - American Wire Gauge (AWG), 40, 101
 - amplification attacks, 317
 - ANDing, 115–116, 128–129
 - Android, 370
 - angled physical contact (APC), G-1, 54, 70
 - Angry IP scanner, 276, 277
 - Ansible, 450
 - antennae
 - cable attenuation, 405
 - placement, 405
 - type, G-1, 394–395
 - anycast addressing, G-1, 112, 143
 - application(s), defined, 3
 - application aware, SDN, 453
 - Application layer, Internet model, 22
 - Application layer (Layer 7), OSI, G-2, 14, 15, 20
 - cloud firewall security, 447
 - functions, 20
 - load balancers, 259
 - logical network diagrams, 274
 - software-defined networking, 453
 - SOHO routers, 20
 - application logs, 286
 - application programming interface (API), G-2, 247, 431
 - cloud access security broker, 459
 - connection methods, 431
 - application security, 236–238
 - Network Time Protocol, 237–239
 - Precision Time Protocol, 239
 - Transport Layer Security, 236–237
 - application services
 - application security, 236–238
 - database services, 246–247
 - disaster recovery services, 254–258
 - email services, 248–250
 - file/print services, 242–246
 - guidelines for supporting, 263–264
 - time synchronization, 239
 - voice and video services, 251–253 (*see also* Voice over Internet Protocol (VoIP))
 - web services, 240–242
 - area border routers, 168
 - ARP. *see* Address Resolution Protocol (ARP)
 - arp command, G-2, 135–136
 - ARP poisoning, 321, 322, 324
 - ARP spoofing, G-2, 321–322
 - assessments, security, 309–310
 - asset inventory documentation, 268–269
 - inventory tools, 268–269
 - licensing, 269
 - warranty support, 269
 - assignment issues, VLAN, 193–194
 - Assisted GPS (A-GPS), 387
 - Asterisk, 251
 - asymmetrical DSL (ADSL), 415
 - asymmetric encryption, 337
 - attack vector, 309
 - attenuation, G-2, 35, 65
 - issues, 69
 - testing, 77
 - Attenuation-to-Crosstalk Ratio, Far End (ACRF), 70
 - Attenuation-to-Crosstalk Ratio, Near End (ACRN), 70
 - audit logs, 287
 - audits, security, 309–310
 - authentication, 332–342
 - access control, 332–333
 - digital certificates, 337–338
 - federated Identity, 339–341
 - identity and access management, 332, 333
 - Kerberos, 336–337, 427
 - key management, 339
 - local, 335–336
 - methods, 334
 - personal, 398
 - policy-based, ZTA, 457
 - public key infrastructure, 338
 - remote, 341–342
 - Security Assertion Markup Language, 340–341
 - single sign-on, 336
 - authentication, authorization, and accounting (AAA), 341
 - authentication factor, 334
 - authentication header (AH), G-2, 107, 419
 - authNoPriv mode, 285
 - authoritative name server, G-2, 221–222
 - authorization, 343–346. *see also* access control
 - identity and access management, 332, 333
 - policy-driven, ZTA, 457
 - authorization creep, 344
 - automated build from a template, 451
 - automatically allocated reservation, 210
 - Automatic Private IP Addressing (APIPA), G-2, 142, 211–212, 216, 323, 324
 - automation, G-2, 450
 - autonegotiation, 87, 98
 - autonomous AP, G-2, 394
 - autonomous system (AS), G-2, 165, 169
 - autonomous system numbers (ASN), 169
 - AUX port, 429
 - availability, CIA triad, 308
 - availability monitoring, G-2, 281–282
- ## B
- baby giant frame, 99
 - backbone (Area 0), 168
 - backbone cabling, 47
 - backdoor Trojan, 318
 - back-reflection, 70
 - backup configuration, G-2, 267
 - backup routers, 262
 - backup strategies, 258
 - badge reader, 373, 374
 - band steering, G-2, 385
 - bandwidth, G-2, 4, 64
 - dynamic routing, 167
 - management, 303–304
 - performance issues, 282, 297–298

speed tester, [G-2, 302](#)
 base [10, 22](#)
 Base64, [328](#)
 baseband radio, [371](#)
 baseband signal (BASE), [35](#)
 baseline configuration, [30](#)
 baseline metrics, [G-2, 281](#)
 basic service area (BSA),
[391–392](#)
 Basic Service Set (BSS), [388](#)
 Basic Service Set Identifier
 (BSSID), [G-2, 388–389, 393, 401](#)
 battery backups, [62](#)
 baud rate, [64](#)
 Bayonet Neill-Concelman (BNC)
 connector, [G-2, 44](#)
 beacon frame, [393](#)
 beamwidth, [394](#)
 behavior aggregates, [303](#)
 behavioral authentication, [334](#)
 behavioral threat research, [316](#)
 bidirectional (BiDi)
 transceivers, [56](#)
 bidirectional wavelength
 division multiplexing (BWDM),
[G-2, 56](#)
 binary conversion, [114–115](#)
 binary value, [22, 107](#)
 IPv4 addressing, [139, 145](#)
 IPv6 addressing, [140](#)
 BIND DNS server, [22, 232](#)
 binding to the server, [346](#)
 biometric authentication, [334](#)
 biometric locks, [373–374](#)
 bit rate, [G-2–G-3, 64](#)
 bits per second (bps), [34](#)
 block tool, [50](#)
 Bluetooth, [381, 390](#)
 Boolean operators, [294](#)
 Border Gateway Protocol
 (BGP), [G-3, 112, 169, 455](#)
 botnet, [G-3, 318](#)
 bottleneck, [G-3, 298](#)
 boundary clock, [239](#)
 branching, [452](#)
 bridge, [G-3, 11](#)
 bridge protocol data unit
 (BPDU), [91, 353](#)
 bridges, [82–83](#)
 root bridge selection, [91–93](#)
 Spanning Tree Protocol,
[90–91](#)

bring your own device (BYOD),
[G-3, 400](#)
 broadband cable modems, [44](#)
 broadband speed checkers, [302](#)
 broadcast, defined, [G-3](#)
 broadcast, SSID, [393](#)
 broadcast addressing, [80, 81,](#)
[110, 121](#)
 broadcast domain, [G-3, 80, 82,](#)
[83, 84](#)
 broadcast storm, [G-3, 100–101,](#)
[282](#)
 brute force, password cracking,
[329](#)
 brute force attack, [G-3](#)
 bugfix, [G-3, 270](#)
 building automation system
 (BAS), [370, 372](#)
 “bump in the wire” appliances,
[355](#)
 burned-in addresses, [79–80](#)
 business continuity (BC), [G-3](#)
 business continuity plan (BCP),
[255](#)
 business impact analysis (BIA),
[G-3, 255, 310](#)
 BWPing, [301](#)
 byte, [22](#)

C

Cable Access TV (CATV), [44, 415](#)
 cable attenuation, antenna,
[405](#)
 cable category issues, [66–67](#)
 cable crimper, [50](#)
 cable-cutting blades, [49](#)
 cabled media, [10](#)
 cable Internet, [415–416](#)
 cable map, [G-3, 272](#)
 Cable Modem Termination
 System (CMTS), [416](#)
 cable stripper, [G-3, 49](#)
 cable tester, [G-3, 67–68](#)
 cable troubleshooting, [64–71](#)
 attenuation and
 interference issues, [69](#)
 cable category issues,
[66–67](#)
 cable issues, [65–66](#)
 cable testers, [67–68](#)
 crosstalk issues, [70](#)

fiber optic cable testing
 tools, [70–71](#)
 guidelines for, [72](#)
 physical inspection, [71](#)
 Power over Ethernet, [101](#)
 reseal the cable, [71](#)
 specification and
 limitations, [64–65](#)
 attenuation, [65](#)
 distance limitations, [65](#)
 interference, [65](#)
 speed *versus*
 throughput, [64–65](#)
 strategies, [71](#)
 tone generator, [69](#)
 verify drivers, [71](#)
 wire map testers, [68–69](#)
 cache-only servers, [226](#)
 caching, DNS, [226](#)
 cameras, [375](#)
 CAM table, [G-12, 99](#). *see also*
 media access control (MAC)
 address table
 canonical name (CNAME), [G-17,](#)
[223, 224](#)
 canonical notation, [G-3, 140](#)
 captive portal, [G-3, 399–400](#)
 cardholder data environment
 (CDE), [312](#)
 card verification value (CVV),
[312](#)
 carrier sense multiple access
 with collision avoidance
 (CSMA/CA), [G-3, 380–381, 385,](#)
[406](#)
 carrier sense multiple access
 with collision detection (CSMA/
 CD), [G-3–G-4, 35–36, 81](#)
 Category (“Cat”) standards, [G-4,](#)
[66–67](#)
 Cat 5e, [41, 66–67, 69](#)
 Cat 6, [37, 41, 50, 67](#)
 Cat 6A, [37, 41, 67](#)
 Cat 7, [37, 41, 67](#)
 Cat 8, [37, 41, 67](#)
 Gigabit Ethernet standards,
[37](#)
 troubleshooting cable
 category issues, [66–67](#)
 twisted pair cable, [35,](#)
[40–41](#)
 untwisting, [50](#)

- CD ROMs, 271
- cellular networks, 371
- cellular radio, G-4, 386
- central policy management, SDN, 453
- central processing unit (CPU), 281
- central repository, 452
- certificate authority (CA), G-4, 237, 242, 338
- Change Advisory Board (CAB), 268
- change management, G-4, 268
- channel, G-4
- channel bandwidths
 - 2.4 GHz, 382–383
 - 5 GHz, 381–382
- channel bonding, G-4, 383–384
- channel link, 65
- channel overlap issues, 406–407
- Chef, 450
- Cipher Block Chaining Message Authentication Code Protocol (CCMP), 397
- cipher suite, G-4, 237
- Cisco
 - data center security, 435
 - Layer 2 Tunneling Protocol, 422
 - NetFlow, 300
 - network icons, 274
 - patches, 270
 - port mirroring, 353
 - switch configuration, 85, 86–93
 - syslog, 288
 - wireless controllers, 394
- Cisco Discovery Protocol (CDP), G-4, 280
- Cisco Unified Communications Manager, 251
- Citrix, 423, 429
- Class F cabling, 42
- classful addressing, G-4, 123–124
- Class II cabling, 42
- classless interdomain routing (CIDR), G-4, 128–129, 444
- class of service (CoS), G-17, 304
- clean agent, 63
- clear to send (CTS), 380
- client authentication, SSH, 427
 - Kerberos, 427
 - public key authentication, 427
 - username/password, 427
- client cache poisoning, DNS, 325
- client disassociation issues, 409
- clientless VPNs, 423
- client ports, 199
- client security, DNS, 228, 229–230
- client-server, G-4, 2–3
- client-to-site VPNs, 422–423
- clock synchronization, 78
- clock types, 239
- closed-circuit television (CCTV), 375
- closed questions, 26
- cloud access security broker (CASB), G-4, 459
- cloud concepts
 - cloud deployment model, 440
 - cloud networking, 443–448
 - cloud service models, 440–441
 - content delivery network, 442
 - datacenter network design, 434–435
 - elasticity, 439–440
 - Fibre Channel, 437–438
 - guidelines for supporting cloud and datacenter architecture, 460
 - modern network environments, 449–459
 - scalability, 439
 - spine and leaf topology, 435–436
 - storage area network, 436–438
- cloud deployment model, G-4, 440
- cloud direct connection, G-4
- cloud gateway, G-4, 445
- cloud hosting, 241
- cloud instances, 443–444
- cloud networking, 443–448
- cloud connectivity options, 445–446
 - colocation, 446
 - direct connect, 446
 - transit gateways, 446
 - virtual private network, 445–446
- cloud firewall security, 446–447
- cloud gateway, 445
- cloud instances, 443–444
- security groups, 447–448
- security lists, 448
- virtual private clouds, 444–445
- cloud service model, G-4, 440–441
- cloud service providers (CSPs), 440
- clustering, G-4, 260–261. *see also* high availability clusters
- cluster services, 258
- coarse wavelength division multiplexing (CWDM), G-4, 57
- coaxial cable, G-4, 35, 43–44
- co-channel interference (CCI), G-4, 406
- cold site, G-4, 257
- collapsed core, G-5, 182
- collision, 35
- collision detection, 35–36
- collision domain, G-5, 35–36, 81, 82, 83, 84
- colocation, G-5, 446
- colocation cages, 374
- command and control (C2), G-5, 318
- command-line interface (CLI), 86, 429, 431
- Commercial Building Telecommunications Cabling Standards, 40, 47–48
- community string, G-5, 285
- Complementary Code Keying (CCK), 382
- compulsory tunneling, 424
- conductor thickness, 101
- confidentiality, CIA triad, 308
- confidentiality, integrity, and availability (CIA) triad, G-5, 308
- (config)# prompt, 86

- (config-if)# prompt, 87
 - configuration baseline, G-5, 266
 - configuration drift, G-5, 267
 - configuration file, 267
 - configuration item (CI), 266
 - configuration management, G-5, 266–267
 - configuration management system (CMS), 266
 - configuration monitoring, G-5, 282
 - conflict identification, 452
 - congestion, 282, 303
 - connectors, testing, 77
 - console connections, 429–430
 - console port, 159, 429
 - consumer-grade smart devices, 370, 372
 - content addressable memory (CAM), G-12, 99
 - content and reputation-based filtering, 359
 - content delivery network (CDN), G-5, 435, 442
 - content filtering, G-5, 358–359
 - contention, 303
 - continuity (open), 68
 - continuity of operations plan (COOP), 255
 - Control And Provisioning of Wireless Access Points (CAPWAP), 394
 - control plane, QoS, 304
 - Converged Enhanced Ethernet, 438
 - converged network adapters (CNAs), 438
 - convergence, G-5, 165
 - Coordinated Universal Time (UTC), 237
 - copper cable and connectors, 35, 39–44
 - Cat cable standards, 35, 40–41
 - coaxial and twinaxial cable and connectors, 35, 43–44
 - plenum and riser-rated cable, 43
 - shielded and screened twisted pair cable, 40
 - twisted pair connector types, 41–42
 - twisted pair copper cabling, 35
 - unshielded twisted pair cable, 39–40
 - copy command, 87
 - core/distribution switch, 181
 - core layer, G-5, 182
 - CORPNET, 389
 - CouchDB, 247
 - Counter Mode, 397
 - course content, mapping. *see* exam objectives for Exam N10-009
 - CPUID's HWMONITOR app, 61
 - crimper, G-5
 - crossed pair (TX/RX transposed), 68
 - crosstalk, G-5, 70
 - cryptographic hash, 335
 - cryptographic hash algorithm, 312
 - cryptographic protocols, updated, 397
 - cryptographic system, 312–313
 - Cumulus VX switch, 354
 - customer edge (CE), 171
 - router, 414
 - customer premises equipment (CPE), 21, 415
 - cybersecurity audits, 310
 - cyclic redundancy check (CRC), G-5, 79, 99
- ## D
- dashboards, configurable, 291
 - data at rest, G-5, 312
 - database services, 246–247
 - datacenter, G-6, 4, 58
 - Data Center Ethernet, 438
 - data center interconnect (DCI), G-5, 455
 - datacenter network design, 434–435
 - datagrams, 12
 - datagram TLS (DTLS), 418
 - Datagram Transport Layer Security (DTLS), 236
 - data historian, 371
 - data in transit, G-5, 312
 - data in use, 312
 - Data Link layer (Layer 2), OSI, G-5, 10–11, 15, 17–18
 - addressing and forwarding, 108
 - cable troubleshooting, 64
 - cisco discovery protocol, 280
 - Ethernet bridges, 82
 - Ethernet frame format, 78
 - functions, 17–18
 - hardware failure issues, 95, 96
 - infrastructure network type, 388
 - interface error counters, 98
 - IP forwarding issues, 150
 - IPv4 address format, 113
 - logical network diagrams, 273
 - packet forwarding, 157
 - Point-to-Point Protocol, 418
 - SOHO routers, 17–18
 - T-carrier system, 415
 - wide area networks, 414
 - data locality, 311
 - Data Over Cable Service Interface Specification (DOCSIS), 416
 - data plane, QoS, 304
 - data remnant, G-5, 271
 - data service units (DSUs), 414
 - data sovereignty, G-5, 311
 - data switches, 181
 - data transport, VoIP, 252
 - deauthentication/disassociation attack, G-6, 402
 - decapsulation, 9
 - decentralized key management model, 339
 - deception technologies, 314
 - decibel loss (dB), G-6, 69, 78
 - decimal conversion, 114–115
 - decommissioning, G-6, 271–272
 - data remnants, 271
 - Instant Secure Erase, 272
 - sanitization, 271
 - Secure Erase, 271
 - dedicated server, 241
 - default gateway, G-6
 - Address Resolution Protocol, 109
 - IPv4 addressing, 119–122

- troubleshooting
 - incorrect IP address, 147
 - incorrect subnet mask, 148
 - ipconfig command, 133
 - IP forwarding issues, 150
 - ping error messaging, 137
- default route, G-6, 155, 156
 - issues, 192
- default VLANs, G-6, 184, 187
- defense in depth, G-6, 347–348
- delay, 65
 - dynamic routing, 167
- demarcation point (demarc), G-6, 21, 47, 415
- demilitarized zone (DMZ), 367
- demodulation, 414
- denial of service (DoS) attack, G-6, 282, 317, 321, 397, 402
- dense wavelength division multiplexing (DWDM), G-6, 57
- density, 58
- designated ports (DP), 91
- desktop Ethernet switch, 85
- device hardening, 348–349
- devices, rogue, 323
- DHCP relay, G-6
- DHCP snooping, G-6, 353
- DHCP Unique Identifier (DUID), 214
- DHCPv6 server configuration, 213–214
- dialog, 14
- dictionary attack, G-6, 329
- Differentiated Services (DiffServ), 303
- Diffie-Hellman algorithm, 421
- DiffServ, G-6
- DiffServ Code Point (DSCP), 303
- dig command, G-6
- digital certificate (DC), G-6, 237, 281, 337–338, 421
- digital subscriber line (DSL), 17, 108, 414, 415
- digital video broadcast satellite (DVB-S), 387
- direct attach copper (DAC), G-6, 44, 76
- direct-attached storage device, 436
- direct connect, 446
- direct current (DC) voltage, 62
- direction, tcpdump, 294
- directly connected route, G-6, 155
- directors, 438
- Direct Sequence Spread Spectrum (DSSS), 382
- dirty optical cables, 71
- disassociation, G-6, 409
- disaster recovery plan (DRP), G-6, 254–255, 257
- disaster recovery services, 254–258
 - concepts, 254–255
 - exam objectives for Exam N10-009, A-10
 - fault tolerance and redundancy, 257–258
 - first hop redundancy protocols, 261–262
 - high availability clusters, 260–261
 - load balancers, 259
 - metrics, 255–256
 - sites, 257
- discards, 98, 299
- discovery protocols, 280
- discretionary access control (DAC), G-6, 343
- distance limitations, 65
- distance vector, G-6, 164
- distinguished name, 345
- distributed control system (DCS), 370
- distributed denial-of-service (DDoS), G-6–G-7, 317–318
 - amplification attacks, 317
 - botnets, 318
 - distributed reflection DoS attacks, 317
 - load balancers, 259
- distributed reflection DoS (DRDoS) attacks, 317
- distribution/aggregation layer, G-7, 181
- distribution system (DS), G-7, 389, 392
- divide and conquer approach, 29
- DNS caching, G-7
- DNS over HTTPS (DoH), G-7, 228
- DNS over TLS (DoT), G-7, 228
- DNS Security Extensions (DNSSEC), 226, 227
- DNS spoofing, G-7
- do copy command, 87
- Domain Information Groper (dig), 232
- DomainKeys Identified Mail (DKIM), 224
- domain names, 218–219
- Domain Name System (DNS), G-7, 218–232
 - attacks, 324–325
 - DNS-based on-path attacks, 324
 - DNS client cache poisoning, 325
 - DNS server cache poisoning, 325
 - canonical name record, 223
 - domain names, 218–219
 - fully qualified domain name, 218–220
 - hierarchy, 219–220
 - host address, 223
 - host name, 125, 218–219
 - internal *versus* external zones, 226–227
 - IP address management, 274
 - ipconfig command, 133
 - mail exchange record, 224
 - name resolution, 220–221
 - network discovery, 276, 277
 - pointer record, 224–225
 - resource records, 221–222
 - reverse DNS query, 224–225
 - security, 227–228
 - client security, 228
 - Security Extensions, 227–228
 - server configuration, 225–226
 - caching, 226
 - server types, 225–226
 - service record, 224
 - text record, 224
 - troubleshooting, 229–232

- Domain Name System Security Extensions (DNSSEC), [G-7, 227–228](#)
 - DORA process, [208](#)
 - dotted decimal notation, [G-7, 114](#)
 - downgrading, [271](#)
 - downlink MU-MIMO (DL MU-MIMO), [385](#)
 - drift, [449](#)
 - drop, [49, 98, 299](#)
 - dry-pipe systems, [63](#)
 - dual band adapter, [383](#)
 - dual stack, [G-7, 144](#)
 - dumpster diving, [G-7, 327–328](#)
 - duplex fiber optic cable, [77](#)
 - duplicate IP and MAC address issues, [149](#)
 - DVD ROMs, [271](#)
 - dynamic ARP inspection (DAI), [353](#)
 - Dynamic Host Configuration Protocol (DHCP), [G-7, 19, 121, 207–210](#)
 - DHCP snooping, [353](#)
 - exclusions, [210](#)
 - IP address management, [274](#)
 - ipconfig command, [133–134](#)
 - IP helper, [216](#)
 - lease time and available leases, [209–210](#)
 - network discovery, [276, 277](#)
 - options, [210](#)
 - process, [207–208](#)
 - relay, [215–216](#)
 - reservations, [210](#)
 - rogue DHCP attacks, [323–324](#)
 - server configuration, [208–209](#)
 - starvation attack, [324](#)
 - troubleshooting, [216–217](#)
 - Unique Identifier, [214](#)
 - dynamic inventories, IaC, [451](#)
 - dynamic link library (DLL), [319](#)
 - Dynamic Rate Switching/Selection (DRS), [390](#)
 - dynamic route, [G-7](#)
 - dynamic routing protocols, [164–170](#)
 - autonomous systems, [165](#)
 - Border Gateway Protocol, [169](#)
 - convergence, [165](#)
 - Enhanced IGRP, [167](#)
 - exam objectives for Exam N10-009, [A-6](#)
 - Open Shortest Path First, [168–169](#)
 - route selection, [169–170](#)
 - Routing Information Protocol, [165–167](#)
 - topology and metrics, [164](#)
- ## E
- EAP over Wireless (EAPoW), [398–399](#)
 - east-west traffic, [G-7, 435](#)
 - eavesdropping, [236, 244, 400](#)
 - edge layer, [181](#)
 - edge routers, [171–172](#)
 - effective isotropic radiated power (EIRP), [G-7, 405–406](#)
 - “egress” traffic, [177](#)
 - Ekahau Site Survey, [392](#)
 - elasticity, [G-7, 439–440](#)
 - electrical environment, monitoring, [61](#)
 - electrician’s scissors (snips), [49](#)
 - electromagnetic interference (EMI), [G-7, 69, 407–408](#)
 - Electronic Industries Alliance (EIA), [40, 58, 67](#)
 - email services, [248–250](#)
 - Internet Message Access Protocol, [250](#)
 - Simple Mail Transfer Protocol, [248–249](#)
 - Encapsulating Security Payload (ESP), [G-7, 107, 419](#)
 - encapsulation, [G-7, 8–9](#)
 - encryption, [G-7, 312–313](#)
 - algorithm, [312](#)
 - end-to-end, [328](#)
 - wireless network attacks, [400–402](#)
 - end of life (EOL), [G-8, 270](#)
 - End of Sale (EOS), [270](#)
 - end of service life (EOSL), [G-8, 270](#)
 - end of support (EOS), [270](#)
 - endpoint security, [347](#)
 - end-to-end encryption, [328](#)
 - end-to-end layer, [13, 15, 20](#)
 - Enhanced IGRP, [167](#)
 - Enhanced Interior Gateway Routing Protocol (EIGRP), [G-8, 107](#)
 - enhanced QSFP (QSFP+), [76](#)
 - enhanced SFP (SFP+), [75–76](#)
 - enterprise authentication, [G-8, 398–399](#)
 - enterprise LAN, [4](#)
 - enterprise mobility management (EMM), [400](#)
 - enterprise network management suites, [276](#)
 - enterprise network topologies, [179–181](#)
 - hybrid topology, [179–180](#)
 - three-tiered hierarchy, [180–182](#)
 - enterprise risk management (ERM), [309](#)
 - enterprise wireless network design, [388–395](#)
 - ad hoc topology, [395](#)
 - antenna types, [394–395](#)
 - heat map, [392](#)
 - infrastructure network type, [388–390](#)
 - mesh topology, [395](#)
 - point to point link, [395](#)
 - range and signal strength, [390–391](#)
 - wireless controllers, [393–394](#)
 - wireless roaming, [392–393](#)
 - wireless survey, [391](#)
 - entrance facilities, [47, 415](#)
 - enumeration, [G-8, 316](#)
 - ephemeral ports, [199](#)
 - Equal Cost Multipathing (ECMP), [436](#)
 - equipment rooms, [58](#)
 - error checking, [79](#)
 - error messaging
 - ICMPv6, [213](#)
 - ping, [137](#)
 - error rate, [299](#)
 - escalation, [G-8, 29](#)
 - Ethernet, [G-8, 34–38](#)
 - 100BASE-TX Fast Ethernet standards, [36–37](#)
 - fiber Ethernet standards, [38](#)

- Gigabit Ethernet standards, [37](#)
- guidelines for deploying and troubleshooting, [72](#)
- media access control and collision domains, [35–36](#)
- network data transmission, [34](#)
- standards, [35](#)
- Ethernet address (EA). *see* Media Access Control (MAC) address
- Ethernet frame format, [78–79](#)
 - error checking, [79](#)
 - EtherType, [79](#)
 - preamble, [78](#)
- Ethernet header, [G-8, 78](#)
- Ethernet networks, [81–87](#)
 - bridges, [82–83](#)
 - hubs, [81–82](#)
 - interfaces, configuring, [74–80](#)
 - switches, [84–87](#)
- Ethernet standards, [35](#)
- Ethernet switches, [17, 84–87](#)
 - interface configuration, [86–87](#)
 - types, [85](#)
 - vendors, [85](#)
- Ethernet trunks, [99](#)
- Ethernet Virtual Private Network (EVPN), [G-8, 455](#)
- EtherType, [79](#)
- ETSI, [444](#)
- EUI-64 interface ID, [142](#)
- event management, [286–292](#)
 - event prioritization and alerting, [289–290](#)
 - log collectors and syslog, [288](#)
 - log reviews, [291–292](#)
 - network device logs, [286–287](#)
 - Security Information and Event Management (SIEM), [290–291](#)
- event prioritization, [289–290](#)
- EventSentry SIEM, [290](#)
- evil twin, [G-8, 401](#)
- exam objectives for Exam N10-009
 - network implementation, [A-6–A-8](#)
 - physical installations, [A-8](#)
 - routing technologies, [A-6](#)
 - switching technologies and features, [A-6–A-7](#)
 - wireless devices and technologies, [A-7](#)
 - networking concepts, [A-1–A-6](#)
 - appliances, [A-1](#)
 - applications, [A-2](#)
 - cloud concepts and connectivity options, [A-2](#)
 - connector types, [A-4](#)
 - functions, [A-2](#)
 - Internet Protocol (IP) types, [A-3](#)
 - modern network environments, [A-5–A-6](#)
 - network architecture, [A-4](#)
 - network topologies, [A-4](#)
 - network types, [A-5](#)
 - protocols/ports, [A-2–A-3](#)
 - related to the Open Systems Interconnection reference model, [A-1](#)
 - traffic types, [A-3](#)
 - transceivers, [A-4](#)
 - transmission media and transceivers, [A-3–A-4](#)
 - wired, [A-3–A-4](#)
 - wireless, [A-3](#)
 - network operations, [A-8–A-11](#)
 - disaster recovery concepts, [A-10](#)
 - IPv4 and IPv6 network services, [A-10–A-11](#)
 - network access and management methods, [A-11](#)
 - network monitoring technologies, [A-9](#)
 - organizational processes and procedures, [A-8–A-9](#)
 - network security, [A-12–A-14](#)
 - attacks, types and impact of, [A-13](#)
 - audits and regulatory compliance, [A-12](#)
 - deception technologies, [A-12](#)
 - device hardening, [A-13](#)
 - key management, [A-13](#)
 - logical security, [A-12](#)
 - network access control, [A-13](#)
 - network segmentation enforcement, [A-13](#)
 - physical security, [A-12](#)
 - security rules, [A-13–A-14](#)
 - security terminology, [A-12](#)
 - zones, [A-14](#)
 - network troubleshooting, [A-14–A-17](#)
 - cable issues, [A-14–A-15](#)
 - hardware issues, [A-15](#)
 - hardware tools, [A-17](#)
 - interface issues, [A-15](#)
 - methodology, [A-14](#)
 - networking device commands, [A-17](#)
 - network services, [A-15–A-16](#)
 - performance issues, [A-16](#)
 - software tools, [A-16–A-17](#)
 - exclusions, DHCP, [210](#)
 - expiration, key management, [339](#)
 - explicit deny, [G-8, 356](#)
 - explicit TLS (FTPS), [244, 249](#)
 - exploit, [G-8, 313](#)
 - extended service area (ESA), [391, 392](#)
 - Extended Service Set (ESS), [389–390](#)
 - Extended Service Set ID (ESSID), [G-8, 389, 391, 392, 393](#)
 - extended unique identifier (EUI), [G-8, 75](#)
 - Extended Unique Identifier-64 (EUI-64), [142](#)
 - Extensible Authentication Protocol (EAP), [G-8, 352, 398](#)
 - eXtensible Markup Language (XML), [267, 340](#)
 - Exterior Gateway Protocol (EGP), [165](#)
 - external DNS zones, [226–227](#)
 - external threats, [315](#)

F

- Facebook, 339
- Far-end crosstalk (FEXT), 70
- Fast Ethernet standards, 36–37
- Fast Link Pulse, 37
- fault tolerance, 257–258
- F-connector coax jack, 17
- FC switch, 438
- FE-13, 63
- federated Identity, 339–341
- fiber distribution panel, G-8, 55
- fiber Ethernet standards, 38
- Fiber Optic Association (FOA), 78
- fiber optic cable, G-8
 - connector types, 53
 - installation, 54–55
 - fiber optic patch cords, 54
 - finishing type, 55
 - patch cords, 54
 - testing tools, 70–71
- fiber optic cables and connectors, 51–57
 - elements for constructing, 51
 - fiber distribution panels, 55
 - fiber optic cable installation, 54–55
 - fiber optic connector types, 53
 - fusion splicing, 55
 - multi-fiber push-on connectors, 56
 - multimode fiber, 52
 - outer jacket designs and materials, 52
 - single mode fiber, 52
 - wavelength division multiplexing, 56–57
- fiber patch cord polarity, 54
- fiber to the curb (FTTC), 416
- fiber to the premises (FTTP), 416
- Fiber to the X (FTTx), 416
- Fibre Channel (FC), G-8, 437–438
 - components, 437–438
 - converged Ethernet, 438
 - defined, 437
 - TCP/IP, 438
- Fibre Channel over Ethernet (FCoE), 438
- Fibre Channel SAN protocol, 76
- Fibre Channel (FC) Switched Fabric (FC-SW), 437
- field devices, 371
- fileless malware, 319
- file/print services, 242–246
 - File Transfer Protocol, 242–243
 - network attached storage, 245–246
 - Secure File Transfer Protocol, 244
 - Server Message Block, 245
- File Transfer Protocol (FTP), G-8, 242–243
 - active *versus* passive FTP, 242–243
 - proxy servers, 357
 - Trivial File Transfer Protocol, 243
- filter syntax, 294
- fingerprinting, 279
- fingerprinting attacks, 316
- finishing type, 55
- fire extinguishers, 63
- fire suppression, 63
- fire triangle, 63
- firewall, G-8, 12, 13, 20
 - access control list, 173, 176–177, 356–357
 - access denied issues, 360
 - cloud firewall security, 446–447
 - Google, 431
 - misconfigured, 360
 - multiple issues, 360
 - ping tests, 150
 - proxy servers, 357–358
 - routing technologies, 176–178
 - packet filtering, 176–177
 - selection and placement, 178
 - stateful inspection, 177
 - screened subnet, 366–367
 - security violation issues, 360
 - SOHO networks, 178
 - stateless, 177
 - virtual, 431
 - web application firewall, 447
- firmware, 270–271
- “first fix” version of IOS, 270
- first hop redundancy protocol (FHRP), G-8, 261–262
 - Hot Standby Router Protocol, 261–262
 - Virtual Router Redundancy Protocol, 262
- fixed address assignment, 210
- fixed Ethernet switch, 85
- flapping interface, 165
- flash devices, 271
- flashing the chip, 270
- flash memory cards, 271
- flooding, 61, 63, 99
- flow data, 300
- fluorinated ethylene propylene (FEP), 43
- FM-200/HFC-227, 63
- foiled twisted pair (FTP), 40
- foiled/unshielded twisted pair (F/UTP), 40
- foil outer shield (F/FTP), 40
- Follow TCP Stream command, Wireshark, 295
- footprinting attacks, 316
- forwarding
 - issues, 149–150
 - Layer 2 *versus* Layer 3, 108
- forward proxies, 357–358, 459
- four-way handshake, 398
- Fox and Hound, 69
- fragmentation, G-8, 158, 213
- frame, G-9, 11
- frame check sequence (FCS), 79
- free space path loss, 404
- frequency band, G-9
- f-type connector, G-9, 44
- full-duplex, G-9, 84
- full tunnel, G-9, 422–423
- fully qualified domain name (FQDN), G-9, 218–220
- functions, 3
- fusion splicing, 55

G

- gain, 394
- gas-based systems, 63
- General Data Protection Regulation (GDPR), G-9, 311
- Generic Routing Encapsulation (GRE), G-9, 107, 145, 418

Generic Security Services Application Program Interface (GSSAPI), 427
 geofencing, G-9, 334, 376
 Geostationary Orbital Satellite Internet Access, 387
 Get command, 283
 Get Next command, 283
 GG45 connectors, 42, 67
 giant, G-9
 giant frame errors, 99
 Gigabit Ethernet standards, 37, 67
 gigabits per second (Gbps), 35
 global addressing, 141–142
 global configuration mode, 86
 Global Information Tracker (GIT), 452
 Global Positioning System (GPS), G-9, 237, 376, 387
 golden configuration, 266
 goodput, 65
 Google
 firewall service, 431
 public DNS resolver, 231
 Google App Engine, 441
 Google Workspace, 339, 441
 grandmaster clock, 239
 graphical user interface (GUI), 122, 246, 277, 425, 428
 gratuitous ARP replies, 321
 group authentication, 398
 guest network, 365, 389, 399–400
 GUI Properties dialog, 121

H

hacking the human. *see* social engineering attacks
 half-duplex, G-9, 81
 half-open scanning, 278, 279
 Halon, 63
 handshake
 four-way, 398
 three-way, 201
 Transmission Control Protocol, 201
 Transport Layer Security, 359
 hard disk drives (HDDs), 271

hardening, G-9, 309, 347–349
 defense in depth, 347–348
 defined, 348
 device and service, 348–349
 hardware
 assets, 268
 failures, 95, 96
 server, 58
 VoIP PBX, 251
 Wi-Fi tester hardware devices, 404
 hashing, G-9, 420
 hdparm utility, 271
 headers, HTTP, 240–241
 heartbeat monitors, 281
 heating, ventilation, air conditioning (HVAC), G-9, 43, 60, 370
 heat map, G-9, 392
 Hertz (Hz), 34
 hexadecimal notation (hex), 23
 hierarchical star, 179
 hierarchical star-mesh, 179–180
 high attenuation, 35
 high availability (HA), G-9, 255
 high availability clusters, 260–261
 active-active clustering, 261
 active-passive clustering, 261
 virtual IP, 260
 high efficiency (HE), 384
 High-Level Data Link Control (HDLC), 415
 holddown timer, 193
 home/residential network, 3
 honeypot, G-9, 314
 hop, G-9, 137, 157–158
 horizontal cabling, 46
 host address, DNS, 223
 host bus adapter (HBA), 437, 438
 host discovery and monitoring, 276–282
 availability monitoring, 281–282
 configuration monitoring, 282
 discovery protocols, 280
 exam objectives for Exam N10-009, A-9

network discovery, 276–277
 Nmap port scanning, 278–279
 Nmap Security Scanner, 277–278
 performance monitoring, 280–281
 Hosted Private cloud
 deployment model, 440
 hosting packages, 241–242
 host key, SSH, 426
 host name, G-9, 218–219
 host number (host ID)
 broadcast addresses, 121
 default gateway, 120
 host address ranges, 118–119
 IPv4 address format, 113
 IPv4 address scheme design, 126, 127
 IPv6 network prefixes, 141
 network masks, 115–116
 subnet masks, 117
 host port, 187
 host routes, 155
 hosts (file), G-9, 2, 230
 host-to-host layer, 13, 15, 20
 hot aisle/cold aisle layout, 59, 60
 hot site, G-9, 257
 hotspot, 399
 Hot Standby Router Protocol (HSRP), 261–262
 HTML5 VPN, G-9, 423
 HTTP Secure (HTTPS), G-10, 242, 357
 hub, G-9, 10
 hub-and-spoke, G-10, 6
 hub/control system, 370
 hubs, 81–82
 human authentication, 334
 human-machine interfaces (HMIs), 371
 humidity, 60–61
 hybrid cloud, G-10, 440
 hybrid fiber coax (HFC), 415
 hybrid topology, G-10, 179–180
 hierarchical star, 179
 hierarchical star-mesh, 179–180
 star of stars, 180

- Hybrid Wireless Mesh Protocol (HWMP), 395
- HyperTerminal, 429
- HyperText Markup Language (HTML), 240
- HyperText Transfer Protocol (HTTP), G-10, 9, 13, 20, 240–242
 - content filtering, 242, 358
 - DNS client security, 228
 - headers and payload, 240–241
 - HTTP Secure, 242, 357
 - Internet Message Access Protocol, 250
 - NoSQL databases, 247
 - open wireless, 400
 - padlock icon, 242
 - proxy servers, 357
 - Security Assertion Markup Language, 340
 - Simple Mail Transfer Protocol, 249
 - state-preserving features, 241
 - web servers, 241–242
- I**
- ICMP Echo Request, 162
- ICMP Port Unreachable response, 162
- ICMP Time Exceeded message, 162, 192
- identification, IAM, 332, 333
- identity and access management (IAM), G-10, 332–333, 347
 - accounting, 333
 - authentication, 332, 333
 - authorization, 332, 333
 - identification, 332, 333
- Identity Association Identifier (IAID), 214
- ifconfig command, G-10, 134–135
- implicit deny, G-10, 356
- implicit TLS, 244, 249
- implicit trust zone, 458
- in-band management, 430
- incorrect IP address, 147–148
- incorrect pin-out/incorrect termination/mismatched standards, 68
- Independent Basic Service Set (IBSS), 395
- Indoor Positioning System (IPS), 376
- industrial control system (ICS), G-10, 370–371, 372
- industrial embedded systems, 370–371
- INERGEN, 63
- informational messaging, 213
- informed consent, 311
- Infrastructure as a Service (IaaS), G-10, 440
- Infrastructure as Code (IaC), G-10, 449–452
 - automation, 450
 - dynamic inventories, 451
 - mastering instructions, 451
 - orchestration, 450
 - playbooks, 450
 - reusable tasks, 452
 - upgrades, 451
- infrastructure layer, SDN, 453
- infrastructure network type, 388–390
- “ingress” traffic, 177
- initiator, Fibre Channel, 437
- insertion loss, 69, 70
- instant secure erase (ISE), G-10, 272
- Institute of Electrical and Electronics Engineers (IEEE), 35. *see also* 802.11 standards
 - ad hoc topology, 395
 - burned-in address, 79
 - extended unique identifier, 75
 - Port-Based Network Access Control (NAC), 352–353
 - Power over Ethernet, 93–94
 - VoIP phones, 253
- insufficient wireless coverage issues, 405–406
- insulation displacement connector (IDC), G-10, 47, 48, 50, 272
- integrity, CIA triad, 308
- Integrity Check Value (ICV), 419
- Intel, 407
- interactive logon, 335
- interface configuration, 86–87
- interface error counters, 98–99
 - cyclic redundancy check errors, 99
 - giant frame errors, 99
 - increasing interface counters, 98
 - run frame errors, 99
- interface identifier, 142
- interface ID/EUI-64, 142
- interface statistics, G-10, 299
- interference, 65
- interference issues, 69, 406, 407–408
- Interior Gateway Protocol (IGP), 165
- Interior Gateway Routing Protocol (IGRP), 167
- intermediate distribution frame (IDF), G-10, 46
- internal DNS zones, 226–227
- internal threats, 315
- International Organization for Standardization (ISO)
 - cable standards, 40, 42, 67
 - ID card standards, 373
 - optical multimode categories, 52
 - reference model, 7
- Internet Assigned Numbers Authority (IANA), 22, 107, 124, 169, 198
- Internet Control Message Protocol (ICMP), G-10, 107, 136, 150
- Internet Corporation for Assigned Names and Numbers (ICANN), 22, 219
- Internet Engineering Task Force (IETF), 22, 236, 300, 453, 454
- Internet eXchange Points (IXPs), 21
- Internet-facing zones, 367
- Internet Group Management Protocol (IGMP), G-10, 107, 111
- Internet Key Exchange (IKE), G-10, 420–421
 - digital certificates, 421
 - network address translation, 421
 - pre-shared key (group authentication), 421

- Internet layer, 22
- Internet Message Access Protocol (IMAP), G-10, 250
- internet of things (IoT), G-10, 369–372
 - botnet, 318
 - devices, 369–370
 - industrial embedded systems, 370–371
 - networks, 371
 - network security, 372
- Internet Protocol (IP). *see* IP headings
- Internet service provider (ISP), G-11, 17, 21
- Internet Small Computer Systems Interface (iSCSI), 438
- Internet standards, 22
- internetwork, 12, 105, 113, 116, 126
- intrusion detection system (IDS), G-11, 13, 354, 367–368
- intrusion prevention system (IPS), G-11, 368
- inventory tools, 268–269
- IOS Software Checker, 270
- IP addressing. *see also* IPv4 addressing; IPv6 addressing
 - address format, 113–115, 140
 - anycast addressing, 112, 143
 - ARP, 108
 - broadcast addressing, 110, 121
 - classful addressing, 123–124
 - forwarding
 - issues, 149–150
 - Layer 2 *versus* Layer 3, 108
 - global addressing, 141–142
 - IPv4 datagram header, 106–107
 - link local addressing, 142–143
 - loopback address, 125, 146
 - multicast addressing, 111, 143
 - subnet masks, 117–118, 130–132, 148
- transition mechanisms, 144–145
- troubleshooting, 147–150
 - duplicate IP and MAC address issues, 149
 - incorrect IP address, 147–148
 - incorrect subnet mask, 148
 - IP configuration issues, 147–148
 - IP forwarding issues, 149–150
- troubleshooting tools, 133–138
 - arp command, 135–136
 - ip command, 135
 - ipconfig command, 133–135
 - ping command, 136–138
 - unicast addressing, 110, 141–142
 - unspecified address, 146
- IP address management (IPAM), G-11, 216, 274, 276
- ip command, G-11, 135
- ipconfig command, G-11, 133–135
- iperf, G-11, 301
- IP filtering, 176
- IP header, G-10, 106–107
- IP helper, G-11, 216
- IP packets
 - ACK, 277, 317
 - cloud concepts
 - Fibre Channel packets, 438
 - security groups, 447
 - virtual extensible LAN, 454–455
 - configuration issues
 - incorrect IP address, 147
 - incorrect subnet mask, 148
 - Data Link layer, 9
 - EtherType value, 79
 - Fast Link Pulse, 37
 - firewalls
 - cloud firewall security, 447
 - misconfigured, 360
 - packet filtering firewalls, 176–177
 - stateful inspection firewalls, 177
- fragmentation, 158, 213
- interface error counters, 98
- Internet Protocol Security, 419–420
- IPv4
 - Address Resolution Protocol, 109
 - anycast addresses, 112
 - arp command, 135
 - broadcast addresses, 110, 121
 - default gateway, 119–120
 - IP protocol type, 107
 - loopback addresses, 125
 - multicast addresses, 111
 - ping command, 136–137, 138
 - unicast addresses, 110
- IPv6, 140
 - anycast addresses, 143
 - loopback addresses, 146
 - multicast addresses, 143
 - NAT64, 145
 - Packet Too Big class of error, 213
 - prefix discovery, 212
 - tunneling, 144–145
- Network layer, 11, 12, 19, 113
- network management
 - bandwidth management, 303
 - interface statistics, 299
 - latency and jitter, 298–299
 - Nmap command, 277
 - Nmap port scanning, 278
 - packet capture, 293–294, 299
 - packet payload, 300
 - protocol analyzers, 295–296
 - quality of service, 304
 - SNMP security, 285
 - tcpdump command, 294
 - traffic flows, 300
 - traffic shapers, 304

- network security
 - access control lists, 356
 - BPDUs, 353
 - DDoS attacks, 317
 - deauthentication attacks, 402
 - intrusion detection and prevention systems, 368
 - narrowband-IoT, 371
 - packet sniffing, 293, 298, 328, 354, 428
 - port guards, 353
 - port mirroring, 353–354
 - proxy servers, 357
 - spoofing attacks, 316, 320–321, 323
 - VLAN hopping attacks, 322
- network services
 - DHCP process, 208
 - DHCP relay, 215
 - DNS client security, 228
 - netstat command, 204
 - Transmission Control Protocol, 200
 - transport layer ports and connections, 198
 - User Datagram Protocol, 203, 223
- packet loss, 65, 69, 98, 289, 298, 299, 304
- packet sniffing, 293, 298, 328, 354, 428
- Physical layer, 9, 11
- routing
 - convergence, 165
 - Enhanced IGRP, 167
 - fragmentation, 158
 - Open Shortest Path First, 168, 169
 - packet forwarding, 157–158
 - Port Address Translation, 175
 - route selection, 169
 - Routing Information Protocol, 166
 - routing loop, 192
 - routing tables, 154, 157
 - switch virtual interfaces, 189
 - traceroute, 162
- SYN, 317
- Telnet, 428
- throughput, 301, 403
- Transport layer, 13
- Transport Layer Security, 237
- tunneling, 144–145, 418
- Virtual Router Redundancy Protocol, 262
- Voice over IP, 187, 252
- wireless networks
 - insufficient wireless coverage issues, 406
 - interference issues, 407
 - Received Signal Strength Indicator, 391
 - Wi-Fi 6, 384
 - Wi-Fi encryption standards, 396
- IP protocol type, G-11, 107
- ip route, 161
- iproute2, 135, 204
- IP scanner, G-11, 276
- IP Security (IPSec), G-10–G-11, 419–420
 - Authentication Header, 419
 - Encapsulating Security Payload, 419
 - header formats, 420
 - transport mode, 419
 - tunneling protocols, 418
 - tunnel mode, 419–420
- IP spoofing, 320–321
- IPv4 addressing, 22, 113–132
 - address format, 113–115
 - 32-bit IPv4 addressing, 114
 - binary/decimal conversion, 114–115
 - host number (host ID), 113
 - network number (network ID), 113
 - address ranges not publicly routable, 126
 - Automatic Private IP Addressing, 142
 - broadcast addressing, 121
 - classful addressing, 123–124
 - Classless Inter-Domain Routing, 128–129
 - default gateway, 119–120
 - exam objectives for Exam N10-009, A-10–A-11
 - host address ranges, 118–119
 - IP interface configuration in Linux, 122
 - IP interface configuration in Windows, 121–122
 - versus IPv6, 139–140
 - loopback address, 125
 - network masks, 115–116
 - public versus private addressing, 124–125
 - scheme design, 126–127, 151–152
 - subnet masks, 117–118, 130–132, 148
 - subnetting, 123–132
 - transition mechanisms, 144–145
 - dual stack, 144
 - NAT64, 145
 - tunneling, 144–145
- IPv4 datagram header, 106–107
- IPv4 link local (IPv4LL), 323, 324
- IPv6 addressing, 139–146
 - address format, 140
 - address prefixes, 145–146
 - loopback address, 146
 - unspecified address, 146
 - anycast addressing, 143
 - exam objectives for Exam N10-009, A-10–A-11
 - global addressing, 141–142
 - interface ID/EUI-64, 142
 - IPv4 versus, 139–140
 - link local addressing, 142–143
 - multicast addressing, 143
 - network prefixes, 141
 - transition mechanisms, 144–145
 - dual stack, 144
 - NAT64, 145
 - tunneling, 144–145
 - unicast addressing, 141–142

IPv6 interface
 autoconfiguration and testing, 212–213
 ICMPv6, 213
 Neighbor Discovery Protocol, 212
 router advertisement, 212
 stateless address autoconfiguration, 212
 IPv6 Rapid Deployment (6RD), 145
 isitdownrightnow.com, 282
 IT contingency planning (ITCP), 255
 iterative lookup, G-11, 221
 IT service continuity planning (ITSCP), 255

J

JavaScript Objection Notation (JSON), 431
 JavaScript Object Notation (JSON), 288
 jitter, G-11, 298–299
 jumbo frame, G-11, 90
 errors, 99
 jump boxes, 430–431
 jump server, G-11, 430–431

K

Kerberos, G-11, 336–337, 427
 Key Distribution Center (KDC), 336–337
 Authentication Service, 336–337
 Ticket Granting Service, 337
 key generation, 339
 key management, 339
 Key Management Interoperability Protocol (KMIP), 339
 key recovery attacks, 397
 Key Signing Key, 227–228
 knowledge-based authentication, 334
 known good duplicate, 30
 Kubernetes, 450

L

LAN switches, 181

laser optimized MMF (LOMMF), 52
 latency, G-11, 65, 298–299
 Layer 1. *see* Physical layer (Layer 1), OSI
 Layer 2. *see* Data Link layer (Layer 2), OSI
 layer 2 broadcast domain, 83
 Layer 2 Tunneling Protocol (L2TP), 422
 Layer 3. *see* Network layer (Layer 3), OSI
 layer 3 switch, G-11, 181
 Layer 4. *see* Transport layer (Layer 4), OSI
 layer 4 switch load balancer, 259
 Layer 5. *see* Session layer (Layer 5), OSI
 Layer 6. *see* Presentation layer (Layer 6), OSI
 Layer 7. *see* Application layer (Layer 7), OSI
 layer 7 switch (content switch) load balancer, 259
 layers, Internet model, 22
 layers, OSI model, 1, 7–8.
see also individual layers
 encapsulation and decapsulation, 8–9
 Internet model layers *versus*, 22
 SOHO routers, 16–20
 troubleshooting, 28–29
 divide and conquer approach, 29
 top-to-bottom/bottom-to-top approach, 28
 upper layers, 14, 15
 LDAP Secure (LDAPS), 346
 lease time and available leases, DHCP, 209–210
 least privilege, G-11, 344, 457
 legacy networks, 245
 legacy system, 313
 licensing, 269
 lifecycle management, 270–271
 lifecycle roadmap, G-11
 light source hazards of fiber optics, 71
 Lightweight Access Point Protocol (LWAPP), 394

lightweight AP, G-11, 394
 Lightweight Directory Access Protocol (LDAP), G-11, 345–346
 link aggregation, G-11, 88–90
 Link Aggregation Control Protocol (LACP), G-11, 89
 link aggregation group (LAG), 89
 Link Layer Discovery Protocol (LLDP), G-11, 280
 link local addressing, G-11, 142–143, 211, 216
 link state, G-11–G-12, 98, 164
 Link State Advertisement (LSA), 168
 link state database (LSDB), 168
 Linux. *see also* UNIX/Linux
 authentication, 336
 client ports, 199
 DHCP issues, 216
 DNS service, 225
 hdparm utility, 271
 IP interface configuration, 122
 iproute2, 204
 login, 335
 netstat command, 203–204
 Nmap Security Scanner, 277
 rogue DHCP, 323
 sudo, 342
 Ttcp, 301
 VoIP PBX, 251
 “live off the land” techniques, 319
 load balancer, G-12, 259
 local address resolution, IPv6, 212
 local area network (LAN), G-12, 3–4
 datacenter, 4
 enterprise LAN, 4
 home/residential network, 3
 SME network, 4
 SOHO router, 4, 16, 17, 19, 20
 local authentication, 335–336
 local connector (LC), G-12, 53–54
 local exchange carrier’s (LEC’s) network, 47
 local/global (L/G) terminology, 80

local loop, 415
 Local Security Authority (LSA), 335
 local sign-in, 335
 location-based authentication, 334
 locks, 373–374
 badge reader, 373, 374
 biometric, 373–374
 rack system, 374
 log aggregation, 290
 log collectors, 288
 logging level, G-12, 289
 logical network diagrams, 273–274
 Application Layer, 274
 Data Link (layer 2), 273
 logical (IP/layer 3), 274
 Physical layer (Layer 1), 273
 logical topology, 10
 logical unit number (LUN), 437
 log-only rules, 356
 log reviews, 291–292
 longest prefix match, 169
 Long Term Evolution (LTE), G-12, 371, 386
 loopback address, G-12, 125, 146
 loopback interface, 159
 loopback tool, 66
 loop issues, 192–193
 loss budget calculator, 78
 lossless Ethernet, 438
 loss of connectivity, 70
 low Earth orbit (LEO), 387
 low-observable characteristics (LOC) attack, 319
 LTE Advanced (LTE-A), 386
 LTE Machine Type Communication (LTE-M), 371
 Lucent Connector, 53–54

M

MAC-derived address, 142
 MAC filtering, G-12, 351
 MAC flooding, G-12, 322
 machine to machine (M2M) communication, 369
 macOS
 Nmap Security Scanner, 277–278
 Secure Shell, 426

MAC spoofing, 320, 353
 mail exchange (MX), G-17, 224, 249
 main distribution frame (MDF), G-12, 47
 malware, G-12
 attacks, 318–319
 code and techniques, 319
 defined, 318
 managed device, 283
 managed Ethernet switch, 85
 management and orchestration (MANO), 444
 Management Frame Protection (MFP), 402
 management frames, 397
 management information base (MIB), G-12, 283, 284
 management plane, QoS, 304
 management plane, SDN, 454
 management port, 429
 Man-in-the-Middle (MitM) attacks. *see* on-path attack
 many-to-one NAT, 174
 MariaDB platform, 246
 massive MIMO, 386
 master image, 451
 mastering instructions, 451
 master key (MK), 399
 master router, 262
 maximum hop count, 193
 maximum tolerable downtime (MTD), G-12, 255–256
 maximum transmission unit (MTU), G-12, 90
 MDI crossover (MDIX), G-12, 81
 mean time between failures (MTBF), G-12, 257, 258
 mean time to failure (MTTF), G-12, 258
 mean time to repair/replace/recover (MTTR), G-12, 258
 media access control (MAC), G-12
 64-bit interface ID determined by, 142
 Address Resolution Protocol, 109, 135, 136
 broadcast addressing, 80, 81, 110
 burned-in addresses, 79–80
 collision domain, 35–36

duplicate issues, 149
 format, 79–80
 interface ID/EUI-64, 142
 Layer 2 *versus* Layer 3 addressing and forwarding, 108
 multicasting addressing, 111
 network interface cards, 75
 patch panels, 48
 spoofing attacks, 149
 media access control (MAC) address table, G-12, 99–100, 136
 media bandwidth, 34
 media converter, G-12, 10
 medium dependent interface (MDI), G-12, 81
 megabits per second (Mbps), 35
 membership, VLANs, 184–185
 memorandum of understanding (MoU), G-13, 275
 memory, 281
 Mesh Basic Service Set (MBSS), 395
 mesh topology, G-13, 6, 395
 Message Digest v5 (MD5), 328
 message submission agent (MSA), 249
 metrics, 164
 Metro Ethernet, 38
 microsegmentation, G-13, 84
 Microsoft
 Active Directory LDAP schema, 345
 logon or sign-in, 335
 Point-to-Point Tunneling Protocol, 422
 Remote Desktop Protocol, 423, 428–429
 Secure Socket Tunneling Protocol, 422
 SQL Server, 246
 Microsoft Azure
 network security groups, 448
 SQL Database, 441
 Virtual Machines, 440
 Microsoft Office 365, 441
 Microsoft Visio, 273, 274
 midspan device, 94

millimeter wave (mmWave), 386
 minimum point of entry (MPOE), 415
 mismatched standards, 68
 missing route, G-13, 191
 mission essential function (MEF), G-13, 310
 MOBIKE multihoming, 421
 Mobile Device Management (MDM), 376
 modern network environments, 449–459
 infrastructure as code, 449–452
 overlay network, 454–455
 Secure Access Service Edge, 459
 software-defined networking, 453–454
 software-defined WAN, 456
 source control, 452
 zero trust architecture, 457–458
 modular Ethernet switch, 85
 modular transceivers, 75–76
 QSFP/QSFP+, 76
 SFP/SFP+, 75–76
 modulation, 414
 MongoDB, 247
 monitor, SNMP, 284–285
 multicast addressing, G-13, 111, 143
 multicast transmissions, 80
 multifactor authentication (MFA), G-13, 334
 multifiber push-on (MPO), G-13, 56, 76
 multihoming, 421
 multimode fiber (MMF), G-13, 38, 52, 56
 multiple input multiple output (MIMO), G-13, 383–384
 Multiprotocol Label Switching (MPLS), 304, 456
 multi-source agreement (MSA), 76
 multitenant cloud deployment model, 440
 Multiuser MIMO (MU-MIMO), G-13, 385
 MySQL platform, 246

N

name resolution, DNS, 220–221
 issues, 230
 name server (NS), 222. *see also* Domain Name System (DNS)
 narrowband-IoT (NB-IoT), 371
 NAT64, G-13, 145
 NAT gateway, 445
 National Electrical Code (NEC), 43
 native VLAN, G-13, 187
 NAT masquerade, 174
 NAT overloading, 174
 Near End (NEXT), 70
 Neighbor Discovery (ND), 277, 322
 inspection, 353
 protocol, G-13, 143, 212
 neighbors, 142, 168
 NetBIOS, 245
 netcat tool, 294
 NetFlow, G-13, 300
 netsh commands, 121–122
 netstat command, G-13, 203–204
 net-tools package, 134–135
 network access control (NAC), G-13, 350–351
 802.1X Port-Based, 352–353
 port security, 350–351
 network access point (NAP), 341
 network access server (NAS), 341
 network adapter, 11. *see also* network interface card (NIC)
 network interface card (NIC)
 Network Address Port Translation (NAPT), 174
 Network Address Translation (NAT), G-13, 125, 145
 cloud gateway, 445
 edge routers, 171–172
 Internet Key Exchange, 421
 Network Address Translation, 173–174
 Port Address Translation, 174–175
 traversal, 421
 network appliances, 3
 backup management, 267
 broadcast storm, 100, 282

configuration management, 282
 edge, 341
 flow data, 300
 hardware assets, 268
 IP address management, 274
 licensing, 269
 passwords, 348
 power management, 62
 Privileged access management, 344
 rack systems, 58
 remote configuration, 425
 Secure Shell, 426
 Simple Network Management Protocol, 283
 Trivial File Transfer Protocol, 243
 unmanaged switches, 429
 network attached storage (NAS), G-13, 245–246
 network data transmission, 34
 network device backup management, 267
 network device logs, 286–287
 audit logs, 287
 performance/traffic logs, 287
 system and application logs, 286
 network discovery, G-13, 276–277
 network edge, 348
 network edge appliances, 341
 Network Functions Virtualization (NFV), G-13, 444
 networking overview, 2–6
 mesh topology, 6
 networking concepts, 2–3
 appliances, 3
 applications, 3
 client-server *versus* peer-to-peer networks, 2–3
 functions, 3
 guidelines for using, 32
 network topology, 4–5
 network types, 3–4
 LAN, 3–4
 WLAN, 4
 star topology, 5–6

- network interface card (NIC), G-14, 11, 18, 55, 74–75, 438
- network interfaces, configuring, 74–80
 - frame format, 78–79
 - Media Access Control address format, 79–80
 - modular transceivers, 75–76
 - network interface cards, 74–75
 - transceiver mismatch issues, 77
 - transceiver signal strength issues, 77–78
- Network layer (Layer 3), OSI, G-14, 12, 15, 18–20
 - addressing and forwarding, 108
 - Address Resolution Protocol, 109
 - Border Gateway Protocol, 169
 - broadcast domain boundaries, 110
 - cloud firewall security, 447
 - content filtering, 358
 - functions, 18–20
 - Generic Routing Encapsulation, 145
 - hardware failure issues, 95
 - Internet Protocol Security, 419
 - IP configuration issues, 147
 - IPv4 address format, 113
 - IPv4 datagram header, 106, 107
 - SOHO routers, 18–20
 - throughput, 65, 403
 - tunneling protocols, 418
 - virtual LANs and subnets, 183
 - VoIP-enabled PBX, 251
 - Wide area networks, 414
- Network Layer Reachability Information (NLRI), 169
- Network Level Authentication (NLA), 428
- network links, 258
- network loop, G-14, 100
- network management
 - event management, 286–292
 - guidelines for, 305
 - host discovery and monitoring, 276–282
 - organizational policies and documentation, 266–275
 - packet capture and analysis, 293–296
 - Simple Network Management Protocol, 283–285
 - traffic monitoring, 297–304
- network mask, G-14, 115–116
- network number (network ID)
 - classful addressing, 123
 - Classless Inter-Domain Routing, 128
 - default gateway, 119–120
 - IPv4 address format, 113
 - IPv4 address scheme design, 126
 - IPv6 network prefixes, 141
 - network masks, 115–116
 - subnet masks, 117
- network protocol, 8–9
 - addressing, 8
 - decapsulation, 9
 - encapsulation, 8–9
- network schematics, 273–274. *see also* logical network diagrams
- network security
 - concepts, 308–314
 - deception technologies, 314
 - design
 - guidelines for supporting, 377
 - internet of things (IoT), 369–372
 - physical security, 373–376
 - zone-based security, 364–368
 - encryption, 312–313
 - exam objectives for Exam N10-009, A-12–A-14
 - features
 - authentication, 332–342
 - authorization and role-based access control, 343–346
 - guidelines for applying network access control solutions, 361
 - hardening, 347–349
 - network security rules, 355–360
 - switch security, 350–354
 - guidelines for supporting security planning and auditing, 330
 - password attacks, 328–329
 - regulatory compliance, 311–312
 - rogue system attacks, 323–325
 - rules, 355–360
 - access control list configuration, 355–357
 - access control list issues, 360
 - content filtering, 358–359
 - misconfigured firewalls, 360
 - proxy servers, 357–358
 - security audits and assessments, 309–310
 - social engineering attacks, 326–328
 - spoofing attacks, 320–322
 - terminology, 308–309
 - threats and attacks, 315–329
 - rogue system attacks, 323–325
 - social engineering, 326–329
 - spoofing attacks, 320–322
 - vulnerability and exploit types, 313
 - zones, 364–365
- network security group, G-14, 343–344, 447–448
- network security list, G-14, 448
- network segmentation enforcement, 364
- network separation, G-14
- network services. *see also* Domain Name System (DNS);

Dynamic Host Configuration Protocol (DHCP)

- Automatic Private IP Addressing, [211–212](#)
- guidelines for supporting, [233–234](#)
- stateless address autoconfiguration, [212–214](#)
- transport and application layer protocols, [198–206](#) (see also Transmission Control Protocol (TCP))
- network sign-in, [335](#)
- network threats and attacks, [315–329](#)
 - deauthentication attack, [402](#)
 - denial of service attacks, [317, 321, 397](#)
 - distributed DoS attacks, [317–318](#)
 - evil twin, [401](#)
 - external *versus* internal threats, [315](#)
 - fingerprinting attacks, [316](#)
 - footprinting attacks, [316](#)
 - malware attacks, [318–319](#)
 - rogue access point, [400](#)
 - rogue system attacks, [323–325](#)
 - social engineering attacks, [326–329](#)
 - spoofing attacks, [316, 320–322](#)
 - threat research, [316](#)
 - wireless network attacks, [400–402](#)
- Network Time Protocol (NTP), [G-14, 216, 237–239, 317](#)
- Network Time Security (NTS), [G-14, 239](#)
- network topology, [4–5](#)
 - networking overview, [2–6](#)
 - OSI model, [7–15](#)
 - SOHO, [16–23](#)
 - troubleshooting, [24–31](#)
- network types, [3–4](#)
 - LAN, [3–4](#)
 - WLAN, [4](#)
- NFPA (National Fire Protection Association), [63](#)
- NFV infrastructure, [444](#)

- ngrep, [294](#)
- nibble, [23](#)
- NIC teaming, [G-14, 88–90](#)
- Nmap, [G-14, 276, 278–279](#)
- Nmap Security Scanner, [277–278](#)
- nodes, [437](#)
- noise, [404](#)
- nondelivery report (NDR), [249](#)
- nondisclosure agreement (NDA), [G-14, 275](#)
- non-persistence, [451](#)
- nontransparent proxy servers, [357](#)
- Nonvolatile Memory Express (NVMe), [437, 438](#)
- north-south traffic, [G-14, 434](#)
- NoSQL database, [247](#)
- no switchport command, [190](#)
- notifications, alerts *versus*, [289](#)
- nslookup command, [G-14, 231](#)
- NT LAN Manager (NTLM) authentication, [335](#)
- NVMe over Fabrics (NVMe-oF), [438](#)
- NVMe over FC (FC-NVMe or NVMe/FC), [438](#)

O

- OAuth, [341](#)
- Object Identifier (OID), [283](#)
- octet, [22](#)
 - 32-bit IPv4 addressing, [114](#)
 - binary/decimal conversion, [114–115](#)
 - boundaries, [116, 117](#)
 - broadcast addressing, [121](#)
 - classful addressing, [123–124](#)
 - IP configuration issues, [148](#)
 - IPv4 address scheme design, [126, 127](#)
 - network masks, [115–116](#)
 - subnet masks, [117–118](#)
- OFDM with multiple access (OFDMA), [384](#)
- onboarding policies, [334](#)
- one-to-many NAT, [174](#)
- on-path attack, [G-14, 320–322](#)
 - ARP poisoning, [321, 322, 324](#)
 - ARP spoofing, [321–322](#)
 - DNS-based, [324](#)
 - IP spoofing, [320](#)
 - MAC spoofing, [320](#)
- open authentication, [G-14, 399](#)
- open questions, [26](#)
- Open Shortest Path First (OSPF), [G-14, 107, 168–169](#)
- OpenSSH, [426, 427](#)
- OpenStack, [440](#)
- Open Systems Interconnection (OSI) model, [7–15](#)
 - data encapsulation/decapsulation, [8–9](#)
 - layers (see layers, OSI model)
 - overview, [7–8](#)
 - summary, [15](#)
 - wide area network, [414](#)
- Open Systems Interconnection (OSI) reference model, [G-14, 7](#)
- operating procedure (SOP), [268](#)
- operational technology (OT), [G-14, 371](#)
- OPNsense security appliance
 - audit logs, [287](#)
 - firewall ruleset, [356](#)
 - IKE, [421](#)
 - IPSec tunnel, [420](#)
 - one-to-one NAT rule, [174](#)
 - as RADIUS client, [342](#)
 - SNMP, [284](#)
 - state table, [177](#)
 - status dashboard, [178](#)
 - transmitting logs to remote syslog server, [288](#)
 - transparent proxy settings, [358](#)
- optical line terminal (OLT), [416](#)
- optical link budget, [G-14, 77](#)
- optical multimode (OM), [G-14, 52](#)
- optical network terminal (ONT), [416](#)
- optical return loss (ORL), [70](#)
- optical transceivers, [52](#)
- option (DHCP), [G-14, 210](#)
- Oracle Cloud Infrastructure (OCI), [448](#)
- Oracle Database, [441](#)
- Oracle remote data access protocol SQL*Net, [246](#)
- orchestration, [G-14, 450](#)

ordinary clock, 239
 organizationally unique identifier (OUI), 79, 80
 organizational policies and documentation, 266–275
 asset inventory
 documentation, 268–269
 change management, 268
 common agreements, 275
 configuration management, 266–267
 decommissioning, 271–272
 exam objectives for Exam N10-009, A-8–A-9
 IP address management, 274
 lifecycle management, 270–271
 logical network diagrams, 273–274
 network device backup management, 267
 physical network diagrams, 272–273
 Orthogonal Frequency Division Multiplexing (OFDM), 382
 out-of-band (OOB) management, G-14, 430
 overcapacity issues, 409
 overlay network, G-15, 454–455
 Data Center Interconnect, 455
 virtual extensible LAN, 454–455
 overwriting, 271
 ownership-based authentication, 334

P

packet capture and analysis, 293–296
 packet sniffer, 293–294
 protocol analyzer, 295–296
 tcpdump command, 294
 packet filtering firewalls, 176–177
 packet forwarding, 157–158
 packet loss, G-15, 298
 packet sniffer, G-15, 293–294
 pairwise master key (PMK), 398, 399
 pan-tilt-zoom (PTZ), 375
 parabolic/dish antenna, 394
 parabolic grid antenna, 394
 parallel fiber optic cable, 77
 parsing, 290
 partial mesh, 6
 passive FTP, 242–243
 passive optical network (PON), 416
 Password Authenticated Key Exchange (PAKE), 398
 passwords
 attacks, 328–329
 cracking, 329
 hardening, 348
 local authentication, 335
 network appliances, 348
 Password Authenticated Key Exchange, 398
 SSH client authentication, 427
 Telnet, 428
 patch, G-15, 270
 patch management, 270
 patch panel, G-15, 48–49, 55, 67
 path selection, 154–155
 payload, HTTP, 240–241
 Payment Card Industry Data Security Standard (PCI DSS), G-15, 312
 peer-to-peer, G-15, 2–3
 Perfect Forward Secrecy (PFS), 237
 performance logs, 287
 performance metric, G-15, 280–281
 performance monitoring, 280–281
 baseline metrics, 281
 performance metrics, 280–281
 Per Hop Behavior (PHB), 303
 perimeter networks, 365
 perimeter security model, 347
 permissions, assigning, 344
 personal authentication, 398
 personal identity number (PIN), 334, 335
 personally identifiable information (PII), G-15, 311
 phased array, 387
 phishing, G-15
 attacks, 326–327
 defined, 326
 physical access control system (PACS), 370
 physical hardware, 3
 physical installation factors, 58–63
 fire suppression, 63
 humidity and temperature, 60–61
 power management, 62
 rack systems, 58–60
 physical interface, 10
 Physical layer (Layer 1), OSI, G-15, 10, 15, 17
 802.11 standards, 380
 cable troubleshooting, 64, 65
 Ethernet frame format, 78
 functions, 17
 hubs, 81
 IP forwarding issues, 150
 logical network diagrams, 273
 SOHO routers, 17
 switch troubleshooting, 97, 98
 wide area network, 414
 wireless troubleshooting, 403
 physical network diagrams, 272–273
 cable map, 272
 rack diagram, 273
 wiring diagram, 272
 physical security, 373–376
 cameras, 375
 geofencing, 376
 locks, 373–374
 physical topology, 10
 piggybacking attacks, 327
 ping command, G-15, 136–138
 basic ping usage, 136–137
 firewalls, 150
 IP forwarding issues, 149–150
 ping error messaging, 137
 destination host unreachable, 137
 no reply (request timed out), 137
 ping switches, 138

pin-out, incorrect, 68
 plain old telephone service (POTS), 251
 Platform as a Service (PaaS), G-15, 441
 playbook, G-15, 450
 plenum, G-15, 43
 pluggable authentication module (PAM), 336
 pointer (PTR) record, G-17, 224–225
 point of presence (PoP), 416
 point-to-point, G-15, 4–5, 395
 Point-to-Point Protocol (PPP), G-15, 415, 418
 Point-to-Point Tunneling Protocol (PPTP), 422
 poisoning, DNS, 324, 325
 polarization, G-15, 395
 policy administrator, 458
 policy-based authentication, ZTA, 457
 policy decision point, 458
 policy-driven authorization, ZTA, 457
 policy engine, 458
 port, G-15
 AUX, 429
 console, 429
 File Transfer Protocol, 243
 management, 429
 Network Time Security, 239
 Nmap port scanning, 278–279
 Nmap Security Scanner, 277–278
 relational database management system, 246–247
 Remote Desktop Protocol, 428
 Secure Shell, 426
 Session Initiation Protocol, 253
 Simple Mail Transfer Protocol, 249
 syslog collector, 288
 TCP, 205–206
 Telnet Daemon, 428
 Transport layer, 198–199
 traps, 284
 UDP, 205–206, 284, 288

Port Address Translation (PAT), G-15, 174–175
 port aggregation, 88
 port filtering/security, 177
 port forwarding, 173
 port guards, 353
 port mirroring, G-15, 353–354
 port naming convention, 272
 port role, G-15, 91
 port scanner, G-16, 278–279
 port scanning, 316
 port security, G-16, 350–351
 disable unneeded switch ports, 350
 MAC filtering, 351
 port-side exhaust/intake, G-16, 59
 port state, G-16, 92
 port status indicators, 96
 port tagging, G-16, 187
 PostgreSQL platform, 246
 posture assessment, G-16, 310
 potentially unwanted applications (PUAs), 318
 potentially unwanted programs (PUPs), 318
 power budget, G-16, 102
 power distribution unit (PDU), G-16, 62
 Powered Devices (PDs), 101–102
 power injector, 94
 power issues, 95
 power load, 62
 power management, 62
 Power over Ethernet (PoE), G-16, 67, 253, 375, 391
 defined in IEEE standards, 93–94
 troubleshooting, 101–102
 cabling issues, 101
 incorrect standard, 101
 power budget exceeded, 102
 PowerShell, 122, 232, 319
 power sourcing equipment (PSE), 93
 power sum crosstalk calculations (PSNEXT, PSACRN, and PSACRF), 70
 power supply units (PSUs), 272, 273

pre-action systems, 63
 preamble, 78
 Precision Time Protocol (PTP), G-16, 239
 prefix discovery, IPv6, 212
 pre-protocol utilization, 299
 Presentation layer (Layer 6), OSI, G-16, 14, 15
 pre-shared key (PSK), G-16, 398, 421
 PRI code, 288
 primary zone, 225
 principals, 336
 privacy extensions, 142
 private branch exchange (PBX), G-16, 251
 private client network, 365
 private cloud, G-16, 440
 private IP address, 124–125
 private key, G-16, 237
 private server administrative network, 365
 privileged access management (PAM), 344–345
 least privilege, 344
 separation of duties, 345
 privileged EXEC mode, 86
 privileges, assigning, 344
 probe, 69
 process, DHCP, 207–208
 process assessment, 310
 process ID (PID), 204
 production configuration, G-16, 267
 programmable logic controllers (PLCs), 370, 371
 PROSet Wi-Fi configuration utility, 407
 protocol, tcpdump, 294
 protocol analyzer, G-16, 293, 295–296
 protocol data unit (PDU), G-16, 9, 12, 13
 Protocol ID/type, 177
 provider edge (PE), 171, 414
 proxy server, G-16, 357–358
 forward proxies, 357–358, 459
 nontransparent, 357
 reverse proxies, 358, 459
 transparent, 357–358
 PRTG scanner, 276

public cloud, [G-16, 440](#)
 public IP address, [124–125](#)
 public key, [G-16, 237](#)
 public key authentication (PKA), [427](#)
 public key cryptography, [337–338](#)
 public key infrastructure (PKI), [G-16, 338](#)
 public network, [365](#)
 public server network, [365](#)
 public switched telephone network (PSTN), [G-17, 21, 251](#)
 public *versus* private addressing, [G-17, 124–125](#)
 pulling cable, [49](#)
 punchdown tool, [G-17, 50](#)
 Puppet, [450](#)
 PuTTY, [429](#)
 PVC (polyvinyl chloride) jackets, [43](#)

Q

quad small form factor pluggable/enhanced quad small form factor pluggable (QSFP/QSFP+), [G-17, 76, 438](#)
 quality of service (QoS), [G-17, 252, 253, 297, 304](#)
 storage area network, [438](#)

R

rack, [G-17, 58](#)
 rack diagram, [G-17, 273](#)
 rack-mounted Ethernet switch, [85](#)
 rack system locks, [374](#)
 rack systems, [58–60](#)
 radio frequency (RF) attenuation, [G-17, 404](#)
 radio frequency interference (RFI), [69, 391](#)
 Radio Grade (RG) designations, [44](#)
 RAID arrays, [437](#)
 Random Early Detection (RED), [304](#)
 range and signal strength, [390–391](#)
 ransomware, [318](#)
 Rapid STP (RSTP), [93](#)

RC4, [396, 397](#)
 RDP Restricted Admin (RDPRA) Mode, [429](#)
 read-only access (query), [346](#)
 read/write access (update), [346](#)
 Real-Time Transport Protocol (RTP), [253](#)
 Received Signal Strength Indicator (RSSI), [G-17, 391](#)
 receive (Rx) wires, [81](#)
 recovery point objective (RPO), [G-17, 256](#)
 recovery time objective (RTO), [G-17, 256, 258](#)
 recursive lookup, [G-17, 221](#)
 recursive resolvers, [221](#)
 redirection, IPv6, [212](#)
 redundancy, [257–258](#)
 Redundant Array of Independent Disks (RAID), [245, 258](#)
 redundant spares, [258](#)
 reflection/bounce (multipath interference), [407](#)
 refraction, [407](#)
 Regional Internet Registries, [228](#)
 registered-jack connector (RJ), [G-17](#)
 RJ11 connectors, [41–42](#)
 RJ45 connectors, [41–42](#)
 regulatory compliance, [311–312](#)
 data locality, [311](#)
 General Data Protection Regulation, [311](#)
 Payment Card Industry Data Security Standard, [312](#)
 personally identifiable information, [311](#)
 relational database management system (RDBMS), [246–247](#)
 relative distinguished name, [345](#)
 relay, DHCP, [215–216](#)
 relying party (RP), [341](#)
 remote access methods
 guidelines for supporting, [432](#)
 remote management, [425–431](#)

virtual private network, [417–424](#)
 wide area network, [414–416](#)
 remote access server (RAS), [417](#)
 remote access Trojan (RAT), [318](#)
 remote authentication, [341–342](#)
 AAA, [341](#)
 RADIUS, [341–342](#)
 TACACS+, [341, 342](#)
 Remote Authentication Dial-In User Service (RADIUS), [G-17, 341–342, 399](#)
 remote configuration of network appliances, [425](#)
 Remote Credential Guard, [429](#)
 remote desktop connections, [425](#)
 remote desktop gateway, [425](#)
 Remote Desktop Protocol (RDP), [G-17, 341, 423, 428–429](#)
 remote host access, [425](#)
 remote management, [425–431](#)
 API connection methods, [431](#)
 console connections, [429–430](#)
 exam objectives for Exam N10-009, [A-11](#)
 jump boxes, [430–431](#)
 out-of-band, [430](#)
 Remote Desktop Protocol, [428–429](#)
 remote host access, [425](#)
 Secure Shell, [426–427](#)
 SSH client authentication, [427](#)
 Kerberos, [427](#)
 public key authentication, [427](#)
 username/password, [427](#)
 SSH commands, [427](#)
 SSH host key, [426](#)
 Telnet, [428](#)
 remote network access, VPN, [417–418](#)
 remote routes, [155](#)
 remote sign-in, [336](#)
 renewal, key management, [339](#)

- repeater, [G-17, 10](#)
- reputational threat intelligence, [316](#)
- Request for Change (RFC), [268](#)
- Requests for Comments (RFCs), [22, 124](#)
- reservation (DHCP), [G-17, 210, 214](#)
- resets, [98](#)
- resetting, [30](#)
- residential cabling standard (TIA 570), [48](#)
- Resolve-DnsName, [231](#)
- resolvers, DNS
 - caching, [226](#)
 - client issues, [229–230](#)
 - client security, [228](#)
 - DHCP options, [210](#)
 - dig command, [232](#)
 - function of, [227](#)
 - name resolution, [220–221](#)
 - name resolution methods, [230](#)
 - nslookup command, [231](#)
- resource record, [G-17, 221–222, 227](#)
- resource units (RUs), [384](#)
- retransmissions, [299](#)
- reusable tasks, IaC, [452](#)
- reverse DNS, [G-17, 224–225](#)
- reversed pair, [68](#)
- reverse proxies, [358, 459](#)
- revocation, key management, [339](#)
- RFC 1542, [215, 217](#)
- RFC 1918, [G-17, 124](#)
- RFC 3927, [211](#)
- riser cabling, [43](#)
- risk, [G-17, 309](#)
- risk assessment, [309](#)
- risk management, [309](#)
- risk posture, [310](#)
- RJ11 connectors, [17, 41–42](#)
- RJ45 connectors, [17, 41–42, 67, 272](#)
- RJ45 patch cord, [49](#)
- roaming, [G-17, 392–393](#)
- roaming misconfiguration issues, [408](#)
- rogue access point, [G-18, 400](#)
- rogue devices and services, [323](#)
- rogue system attacks, [323–325](#)
 - DNS attacks, [324–325](#)
 - rogue devices and services, [323](#)
 - rogue DHCP, [323–324](#)
- role-based access control (RBAC), [G-18, 343–346](#)
- rollback, [271](#)
- root bridge selection, [G-18, 91–93](#)
- root bridge switch, [353](#)
- rootkit, [318](#)
- round robin DNS, [223](#)
- round trip time (RTT), [136, 298, 387](#)
- route command, [G-18, 161](#)
- route flapping, [165](#)
- router, [G-18](#)
- router, SOHO, [16](#)
- router advertisement (RA), [G-18, 212](#)
- Router Advertisement (RA) Guard, [G-18, 353](#)
- router configuration, [159–160](#)
- router implementation, [419](#)
- router solicitation (RS), [212](#)
- route selection, [169–170, 191–192](#)
- routing. *see also* dynamic routing protocols
 - convergence, [165](#)
 - Enhanced IGRP, [167](#)
 - fragmentation, [158](#)
 - Open Shortest Path First, [168, 169](#)
 - packet forwarding, [157–158](#)
 - Port Address Translation, [175](#)
 - route selection, [169](#)
 - Routing Information Protocol, [166](#)
 - routing loop, [192](#)
 - routing tables, [154, 157](#)
 - switch virtual interfaces, [189](#)
 - traceroute, [162](#)
- routing by rumor, [165](#)
- routing entry, parameters defining, [154](#)
- Routing Information Protocol (RIP), [G-18, 165–167](#)
- routing loop, [G-18, 192–193](#)
- routing table, [G-18](#)
 - default routes, [155, 156](#)
 - directly connected routes, [155](#)
 - example, [156–157](#)
 - host routes, [155](#)
 - issues, troubleshooting, [191–192](#)
 - path selection, [154–155](#)
 - remote routes, [155](#)
 - routing entry, parameters defining, [154](#)
 - static routes, [156](#)
- routing table tools, [160–163](#)
 - ip route, [161](#)
 - route command, [161](#)
 - show arp, [160](#)
 - show route, [160](#)
 - traceroute, [162](#)
 - tracert, [162–163](#)
- routing technologies. *see also* routing table
 - default routes, [156](#)
 - dynamic (*see* dynamic routing protocols)
 - enterprise network topologies, [179–181](#)
 - hybrid topology, [179–180](#)
 - three-tiered hierarchy, [180–182](#)
- firewalls, [176–178](#)
 - packet filtering, [176–177](#)
 - selection and placement, [178](#)
 - stateful inspection, [177](#)
- fragmentation, [158](#)
- guidelines for supporting routing and campus network design, [195–196](#)
- network address translation, [171–175](#)
 - edge routers, [171–172](#)
 - Network Address Translation, [173–174](#)
 - Port Address Translation, [174–175](#)
- packet forwarding, [157–158](#)

router configuration, 159–160
troubleshooting, 191–194
 default route issues, 192
 loop issues, 192–193
 routing table issues, 191–192
 VLAN assignment issues, 193–194
trunking and IEEE 802.1Q, 186
VLANs, 183–194
 assignment issues, troubleshooting, 193–194
 default, 187
 IDs and membership, 184–185
 native, 187
 port tagging, 187
 routing, 188–190
 subnets, 183–184
voice, 187–188
RTP Control Protocol (RTCP), 253
runt, G-18, 99

S

Salesforce, 441
Samba software suite, 245
sandbox environments, 30
sanitization, G-18, 271
satellite systems, 386–387
satellite, G-18
scalability, G-18, 439
schematics, 273–274. *see also* logical network diagrams
scope (DHCP), G-18, 209
scope exhaustion, DHCP, 216
scope of a network, 3
screened subnet, G-18, 366–367
screened twisted pair (ScTP), 40
scripting tools, 319
SDN controller, 453
secondary zone, 226
Secure Access Service Edge (SASE), G-18, 459
secure administrative workstation (SAW), 429, 434
secure erase (SE), G-18, 271
secure file transfer protocol (SFTP), 244, 426
Secure Hash Algorithm (SHA), 328
Secure Shell (SSH), G-18, 159, 244, 336, 341, 426–427
 remote host access, 425
 SSH client authentication, 427
 Kerberos, 427
 public key authentication, 427
 username/password, 427
 SSH commands, 427
 SSH host key, 426
 virtual terminal, 429
Secure Sockets Layer (SSL), 237, 346, 422
Secure Socket Tunneling Protocol (SSTP), 422
secure version of SIP (SIPS), 253
secure web gateway (SWG), 459
security, DNS, 227–228
security, SNMP, 285
Security Accounts Manager (SAM), 335
Security Assertion Markup Language (SAML), G-18, 340–341
security association (SA), 420
security audits and assessments, 309–310
security controls, 310
Security Information and Event Management (SIEM), G-18, 290–291
Security Service Edge (SSE), G-18, 459
security violation issues, 360
segments, 10, 13
self-encrypting drives (SEDs), 272
self-signed certificate, G-18, 338
Sender Policy Framework (SPF), 224
sensors, 61
separation of duties, G-18, 345

Serial Advanced Technology Attachment (SATA), 271, 437
Serial Attached SCSI (SAS), 271, 437
server, AAA, 341
server cache poisoning, DNS, 325
server configuration
 DHCPv6, 213–214
 Domain Name Service, 225–226
server configuration, DHCP, 208–209
server hardware, 58
Server Message Block (SMB), G-19, 20, 245
server rooms, 58
server types, DNS, 225–226
service assets, 266
service hardening, 348–349
service level agreement (SLA), G-19, 275
service (SRV) record, 224
services, rogue, 323
Service Set Identifier (SSID), G-19, 388–389, 393, 399, 401
 broadcast and beacon frame, 393
session, 14
session control, VoIP, 252
Session Initiation Protocol (SIP), G-19, 253
Session layer (Layer 5), OSI, G-19, 14, 15, 236, 418
shadow, G-19, 323
shared hosting, 242
shellcode, G-19, 319
shielded/foiled twisted pair (S/FTP), 40
shielded modular plug, 50
shielded twisted pair (STP), G-19, 40
short, 68
shortest path first (SPF), 168
shoulder surfing, G-19
shoulder surfing attacks, 327
show arp command, G-19, 160
show commands, G-19, 97–98
 administratively down/down, 97
 down/down, 97
 down/error disabled, 98

- show config, 97
 - show interface, 97
 - up/down (suspended), 98
 - up/up, 97
 - show power inline command, 102
 - show route command, G-19, 160
 - shutdown command, 87
 - signaling, 10
 - signal mode, 35
 - signal strength, 390–391
 - signal-to-noise ratio (SNR), 391
 - Simple Authentication and Security Layer (SASL), 346
 - simple bind, 346
 - Simple Mail Transfer Protocol (SMTP), G-19, 20, 248–249, 357
 - Simple Network Management Protocol (SNMP), G-19, 283–285
 - agents, 283–284
 - monitor, 284–285
 - network discovery, 277
 - security, 285
 - Simple NTP (SNTP), 238
 - Simple Object Access Protocol (SOAP), 340
 - Simultaneous Authentication of Equals (SAE), G-19, 397, 398
 - single-mode fiber (SMF), G-19, 38, 52, 56
 - single sign-on (SSO), G-19, 336
 - single tenant cloud deployment model, 440
 - SIP Uniform Resource Identifier (URI), 253
 - site resiliency, 257
 - site survey, G-19
 - site-to-site VPNs, 424
 - six position connectors, 42
 - size of a network, 3
 - slam method, 238
 - slew method, 238
 - small and medium-sized enterprise (SME) network, 4
 - Small Computer Systems Interface (SCSI), 437, 438
 - small form factor pluggable/enhanced small form factor pluggable (SFP/SFP+), G-19, 75–76
 - small office/home office (SOHO) network, G-19, 4, 16–23
 - binary notation, 22
 - broadband speed checkers, 302
 - hexadecimal notation (hex), 23
 - Internet, 21–22
 - Internet service providers, 21
 - public switched telephone network, 21
 - standards, 22
 - router firewall, 178
 - routers, 16–21
 - Application layer, 20
 - Data Link layer, 17–18
 - Internet, 21–22
 - Network layer, 18–20
 - Physical layer, 17
 - Transport layer, 20
 - security, 20–21
 - self-signed certificate, 338
 - switches, 85
 - smart buildings, 370, 372
 - smart card login, 336
 - smart cards, 373
 - smart devices, 370
 - snagless cable construction, 42
 - sniffing attacks, 349
 - SNMP v2c, 285
 - SNMP v3, 285
 - snowflake topology, 180
 - social engineering, G-19, 326
 - social engineering attacks, 326–328
 - dumpster diving attacks, 327–328
 - phishing attacks, 326–327
 - piggybacking attacks, 327
 - shoulder surfing attacks, 327
 - tailgating attacks, 327
 - Social Security number (SSN), 311
 - socket, G-20, 199
 - software
 - assets, 268
 - central repository, 452
 - decommissioning, 271–272
 - enterprise mobility management, 400
 - enterprise network management suites, 276
 - flow data, 300
 - inventory tools, 268, 269
 - IP address management, 274
 - lifecycle management, 270–271
 - log collectors and syslog, 288
 - log reviews, 291
 - mapping, 274
 - Mobile Device Management, 376
 - NetFlow, 300
 - Nmap port scanning, 278–279
 - Nmap Security Scanner, 277–278
 - packet sniffer, 293
 - password cracking, 329
 - protocol analyzer, 295–296
 - Security Information and Event Management, 290
 - SNMP agent, 283–284
 - SNMP monitor, 284
 - software-defined networking, 453–454
 - traffic shapers, 304
 - traffic testing tools, 301–302
 - updates, 271
 - VoIP PBX, 251
 - VPN solutions, 422
 - vulnerabilities, 270
 - Wi-Fi analyzer, 404
- Software as a Service (SaaS), G-20, 441
- software-defined networking (SDN), G-20, 453–454
 - architecture, 453–454
 - management plane, 454
 - properties, 453
- software-defined WAN (SD-WAN), G-20, 456
- solid state drives (SSDs), 271
- Something/Anything/Everything as a Service (XaaS), 440
- source control, G-20, 452
- source ports, 199

- Spanning Tree Protocol (STP), G-20, 90–93
 - broadcast storm, 282
 - configuration, 91–93
 - manipulation attack, 322
 - port guards, 353
- SPAN (switched port analyzer)/port mirroring, 293
- spatial diversity, 383
- spatial multiplexing, 383
- spectrum analyzer, G-20, 408
- speed *versus* throughput, 64–65
- Spiceworks IT Support management tool, 31
- spine and leaf topology, G-20, 435–436
- splices, 77
- split horizon, 193
- split pair, 69
- split tunnel, G-20, 422
- spoofing, G-20, 316
- spoofing attacks, 149, 316, 320–322
 - ARP spoofing, 321–322
 - IP spoofing, 320–321
 - MAC flooding attacks, 322
 - MAC spoofing, 320
 - on-path attacks, 320–322
 - VLAN hopping attacks, 322
- spyware, 318
- stackable Ethernet switch, 85, 86
- standard operating procedure (SOP), G-20, 450
- standby group, 261
- standby power supplies, 258
- star of stars, 180
- start frame delimiter (SFD), 78
- Start of Authority (SOA), 221–222
- star topology, G-20, 5–6
- STARTTLS, 249
- starvation attack, DHCP, 324
- state/bare metal, 267
- stateful inspection firewalls, 177
- stateless address
 - autoconfiguration (SLAAC), G-20, 212–214
- static address assignment, 210
- static route, G-20, 156
- stations, 381, 388
- Statistics menu, Wireshark, 295
- stealthy hosts, 278
- sticky MACs, 351
- storage
 - key management, 339
 - performance metrics, 281
- storage area network (SAN), G-20, 38, 76, 436–438, 440
- straight tip (ST) connector, G-20, 53
- stranded-wire patch cords, 49
- “strip” sockets, 62
- structured cable installation, 49
- structured cabling system, 45–47
 - backbone cabling, 47
 - entrance facilities/demarc, 47
 - horizontal cabling, 46
 - telecommunications room, 47
 - work area, 45
- Structured Query Language (SQL), G-20, 246
- stub resolver, 220
- subinterface, G-20, 188–189
- subnet, 183–184
- subnet addressing, G-20, 117, 127, 128, 142, 148
- subnet mask, 117–118, 148
 - incorrect, 148
 - variable length, 130–132
- subnetting, 123–132
 - address ranges reserved for special use and not publicly routable, 126
 - classful addressing, 123–124
 - Classless Inter-Domain Routing, 128–129
 - IPv4 address scheme design, 126–127, 151–152
 - loopback address, 125
 - public *versus* private addressing, 124–125
 - purpose, 119
 - variable length subnet masking, 130–132
- subscriber connector (SC), G-20, 53, 54
- substitute known working hosts, 66
- sudo, 342
- supervisory control and data acquisition (SCADA), G-20, 371, 372
- supplicant, AAA, 341
- supplicant device, 399
- switch, G-20, 11
 - Ethernet, 84–87
 - guidelines for deploying, 103
 - ping, 138
 - port configuration, 88–94
- switched port analyzer (SPAN), 353
- switching, 108
- switchport command, 87
- switch port configuration, 88–94
 - exam objectives for Exam N10-009, A-6–A-7
 - link aggregation, 88–90
 - maximum transmission unit, 90
 - NIC teaming, 88–89
 - Power over Ethernet, 93–94
 - Spanning Tree Protocol, 90–93
- switch security, 350–354
 - Extensible Authentication Protocol, 352
 - IEEE 802.1X Port-Based NAC, 352–353
 - network access control, 350–351
 - port guards, 353
 - port mirroring, 353–354
- switch troubleshooting, 95–102
 - broadcast storms, 100–101
 - hardware failures, 96
 - interface error counters, 98–99
 - MAC address tables, 99–100
 - network loop issues, 100
 - port status indicators, 96
 - power issues, 95
 - Power over Ethernet issues, 101–102
 - restarting, 96
 - switch show commands, 97–98

Switch Virtual Interface (SVI),
[G-20, 189–190](#)
 symbols (series of events), [64](#)
 symmetrical DSL, [415](#)
 symmetric cipher, [337](#)
 symmetric encryption key, [337](#)
 symmetric session key, [337](#)
 SYN packet, [317](#)
 syslog, [G-21, 288](#)
 severity levels, [289](#)
 system lifecycle, [270](#)
 system logs, [286](#)

T

T568A/T568B (T568A), [G-21, 47–48, 49](#)
 tabletop exercise, [G-21, 255](#)
 tactics, techniques, and procedures (TTPs), [314, 316](#)
 tagged ports, [187](#)
 tail drop, [304](#)
 tailgating, [G-21, 327](#)
 TAP, [294](#)
 tape libraries, [437](#)
 tape media, [271](#)
 target, Fibre Channel, [437](#)
 T-carrier, [415](#)
 tcpdump command, [G-21, 294](#)
 TCP flag, [G-21, 201](#)
 TDM-based PBXes, [251](#)
 TeamViewer, [429](#)
 telecommunications
 closets, [58](#)
 companies, [414](#)
 room, [47](#)
 Telecommunications Industry Association (TIA), [40, 67](#)
 teletype (TTY), [426](#)
 Telnet, [G-21, 428, 429](#)
 Telnet Daemon, [428](#)
 temperature, [60–61](#)
 templates, IaC, [451](#)
 Temporal Key Integrity Protocol (TKIP), [396–397](#)
 temporary interface ID or token, [142](#)
 tenants, [440](#)
 TERA connectors, [42, 67](#)
 Terminal Access Controller Access Control System Plus (TACACS+), [G-21, 341, 342, 399](#)

termination
 incorrect, [68](#)
 tools and techniques, [50](#)
 test access point (TAP), [G-21, 293](#)
 text (TXT) record, [224](#)
 theory of probable cause, [27–29](#)
 threat, [G-21, 309](#)
 threat actor, [309](#)
 threat agent, [309](#)
 threat assessment, [309](#)
 threat data, [316](#)
 threat research, [316](#)
 threat scope reduction, ZTA, [457](#)
 three-factor authentication, [334](#)
 three-tiered hierarchy, [G-21, 180–182](#)
 access or edge layer, [181](#)
 collapsed core, [182](#)
 core layer, [182](#)
 distribution or aggregation layer, [181](#)
 three-way handshake, [201](#)
 throughput, [G-21, 281](#)
 speed *versus*, [64–65](#)
 testers, [301](#)
 Ticket Granting Service (TGS), [337, 427](#)
 Ticket Granting Ticket, [337, 427](#)
 time division multiplexing (TDM), [251, 415](#)
 time drift, [238](#)
 time factor authentication, [334](#)
 time-of-day restrictions, [368](#)
 time synchronization, [239](#)
 time to live (TTL), [G-21, 137, 226](#)
 Tip and Ring wires, [42](#)
 tone generator, [G-21, 69](#)
 tools, [160–163](#)
 inventory, [268–269](#)
 ip route, [161](#)
 route command, [161](#)
 routing table, [160–163](#)
 scripting, [319](#)
 show arp, [160](#)
 show route, [160](#)
 termination, [50](#)
 traceroute, [162](#)

tracert, [162–163](#)
 traffic testing, [301–302](#)
 troubleshooting, IP
 addressing, [133–138](#)
 top-level domain (TLD), [218, 219, 220](#)
 top listeners, [302](#)
 top-of-rack (ToR), [G-21, 436](#)
 topology, [G-21, 164](#)
 Topology Change Notifications, [91](#)
 top talkers, [302](#)
 top-to-bottom/bottom-to-top approach, [28](#)
 TP-LINK SOHO access point, [397](#)
 traceroute command, [G-21, 162](#)
 tracert command, [G-21, 162–163](#)
 traffic analysis, [G-21, 295](#)
 traffic flows, [300](#)
 traffic logs, [287](#)
 traffic monitoring, [297–304](#)
 bandwidth management, [303–304](#)
 flow data, [300](#)
 interface statistics, [299](#)
 performance issues, [297–299](#)
 bandwidth, [297–298](#)
 bottlenecks, [298](#)
 latency and jitter, [298–299](#)
 packet loss, [298](#)
 traffic shaping, [304](#)
 traffic testing tools, [301–302](#)
 traffic shaper, [G-21, 304](#)
 traffic testing tools, [301–302](#)
 bandwidth speed testers, [302](#)
 listeners, [302](#)
 throughput testers, [301](#)
 top talkers, [302](#)
 transceiver, [G-21, 10](#)
 mismatch issues, [77](#)
 modular, [75–76](#)
 signal strength issues, [77–78](#)
 transit gateways, [446](#)

- Transmission Control Protocol (TCP), [G-21, 9, 200–202](#)
 - connection teardown, [202](#)
 - data in transit, [313](#)
 - Fibre Channel, [438](#)
 - flags, [201](#)
 - header fields, [200](#)
 - IP protocol types, [107](#)
 - Nmap port scanning, [278](#)
 - ports, [177, 205–206, 428](#)
 - Telnet, [201](#)
 - three-way handshake, [201](#)
- transmit (Tx) wires, [81](#)
- transparent proxy servers, [357–358](#)
- transport agnostic, SDN, [453](#)
- transport and application layer protocols, [198–206](#). *see also* Transmission Control Protocol (TCP)
 - netstat command, [203–204](#)
 - Transport layer ports and connections, [198–199](#)
 - UDP, [203, 205–206](#)
- Transport layer, Internet model, [22](#)
- Transport layer (Layer 4), OSI, [G-21, 13, 15, 20](#)
 - cloud firewall security, [447](#)
 - content filtering, [358](#)
 - functions, [20](#)
 - IPv4 datagram header, [107](#)
 - load balancers, [259](#)
 - NAT64, [145](#)
 - Port Address Translation, [175](#)
 - port filtering/security, [177](#)
 - ports and connections, [198–199](#)
 - SOHO routers, [20](#)
 - throughput, [65](#)
 - Transmission Control Protocol, [200](#)
 - User Datagram Protocol, [203](#)
- Transport Layer Security (TLS), [G-21, 236–237](#)
 - content filters, [359](#)
 - Datagram Transport Layer Security, [236](#)
 - digital certificates, [337–338](#)
 - DNS client security, [228](#)
 - end-to-end encryption, [329](#)
 - File Transfer Protocol, [242–243](#)
 - handshake, [359](#)
 - HTTP Secure, [242](#)
 - Internet Message Access Protocol, [250](#)
 - Network Time Protocol, [239](#)
 - proxies, [359](#)
 - Secure FTP, [244](#)
 - Session Initiation Protocol, [253](#)
 - Simple Mail Transfer Protocol, [249](#)
 - tunneling protocols, [418](#)
 - VoIP phones, [253](#)
 - VPN solution, [422](#)
- transport mode, [419](#)
- Trap command, [283](#)
- traversal, NAT, [421](#)
- tree topology, [179](#)
- triple homed network, [366](#)
- Trivial File Transfer Protocol (TFTP), [G-22, 243, 270](#)
- Trojan, [318](#)
- troubleshooting
 - broadcast storms, [100–101](#)
 - cable, [64–71](#)
 - default gateway, [133, 137, 147–148, 150](#)
 - DHCP, [216–217](#)
 - DNS, [229–232](#)
 - Domain Name System, [229–232](#)
 - Dynamic Host Configuration Protocol, [216–217](#)
 - Ethernet, [72](#)
 - exam objectives for Exam N10-009, [A-14–A-17](#)
 - IP addressing, [147–150](#)
 - network loops, [100](#)
 - OSI model layers, [28–29](#)
 - photos of current configuration, [272](#)
 - Power over Ethernet, [101–102](#)
 - routing table issues, [191–192](#)
 - routing technologies, [191–194](#)
 - switches, [95–102](#)
 - tools, [133–138](#)
 - unresponsive service issues, [281–282](#)
 - VLAN assignment issues, [193–194](#)
- troubleshooting methodology, [G-22, 24–31](#)
 - divide and conquer approach, [29](#)
 - document findings, actions, and outcomes, [31](#)
 - establish a plan of action, [29–30](#)
 - accept, [30](#)
 - repair, [29](#)
 - replace, [29](#)
 - establish a theory of probable cause, [27–29](#)
 - divide and conquer approach, [29](#)
 - top-to-bottom/bottom-to-top approach, [28](#)
 - exam objectives for Exam N10-009, [A-14](#)
 - identify problem symptoms, [26–27](#)
 - approach multiple problems individually, [27](#)
 - determine if anything has changed, [27](#)
 - identify symptoms and duplicate the problem, [26](#)
 - identify the problem, [25–26](#)
 - gather information, [25](#)
 - question users, [26](#)
 - implement the solution, [30](#)
 - network, [24–25](#)
 - test the theory to determine the cause, [29](#)
 - top-to-bottom/bottom-to-top approach, [28](#)
 - verify the solution, [30–31](#)
 - trunks/trunking, [G-22, 186](#)
 - Ttcp, [301](#)
 - tunneling, [G-22, 144–145](#)
 - tunneling protocols, [418](#)
 - Generic Routing Encapsulation, [418](#)
 - Internet Protocol Security, [418](#)

Point-to-Point Protocol, 418
 Transport Layer Security, 418
 tunnel mode, 419–420
 tuple, 300, 356
 twinaxial cable, G-22, 44
 twisted pair cable, G-22, 35
 twisted pair connector types, 41–42
 RJ11 connectors, 42
 RJ45 connectors, 41–42
 two-factor authentication, 334
 TX/RX transposed (crossed pair), 68
 type, tcpdump, 294
 typosquatting techniques, 324

U

Ultra Physical Contact (UPC), G-22, 54, 70
 unauthorized modification, 236
 unicast addressing, G-22, 110, 141–142
 unicast transmissions, 80, 81
 Unicode, 14
 unidirectional antenna, 394
 Uniform Resource Identifier (URI), 253
 uniform resource locator (URL), 240
 filtering, 359
 uninterruptible power supply (UPS), G-22, 62, 95, 258, 272, 273
 universal address, 79
 universal/local (U/L), 79, 80
 UNIX/Linux. *see also* Linux
 DNS poisoning, 325
 Secure Shell, 426
 syslog, 288
 tunneling, 145
 unmanaged Ethernet switch, 85
 unpatched system, 313
 unshielded twisted pair (UTP) cable, G-22, 39–40
 unspecified address, 146
 untagged ports, 187
 updated cryptographic protocols, 397
 upgrades, IaC, 451

uplink MU-MIMO (UL MU-MIMO), 385
 uptime monitors, 281
 URL filtering, G-22, 359
 USB thumb drives, 271
 User Account Control (UAC), 342
 user agents, 253
 User Datagram Protocol (UDP), G-22, 203
 Datagram Transport Layer Security, 236
 device queries, 284
 IP protocol types, 107
 load balancers, 259
 Network Time Protocol, 237
 Nmap port scanning, 278
 overlay networks, 454
 port numbers, 177
 ports, 205–206, 284, 288
 relational database management system, 246
 Session Initiation Protocol, 253
 structure of, 203
 Trivial File Transfer Protocol, 237
 user EXEC mode, 86
 username, 427
 utilization, 281, 299

V

validation tests, 255
 variable length subnet masking (VLSM), G-22, 130–132
 version control, G-22, 452
 vertical cabling, 47
 vertical-cavity surface-emitting lasers (VCSEL), 52
 vertical rod antenna, 394
 very high-speed DSL (VDSL), 416
 very high throughput (VHT), 384
 very small aperture terminal (VSAT), 387
 video, bandwidth requirements, 298
 video conferencing (VTC), 251, 252
 virtual appliances, G-22, 3

Virtual Extensible LAN (VXLAN), G-22, 454–455
 virtual firewall, 431
 virtual IP, G-22, 260
 virtualization, 30
 virtual LANs (VLANs), G-22, 84, 183–194
 assignment issues, 193–194
 automated VLAN pooling, 394
 default, 187
 firewalls, 355
 hopping attacks, 322
 IDs and membership, 184–185
 native, 187
 port guards, 353
 port tagging, 187
 routing, 188–190
 subinterfaces, 188–189
 Switch Virtual Interface, 189–190
 subnets, 183–184
 traffic management, 304
 trunking and IEEE 802.1Q, 186
 voice, 187–188
 VoIP phones, 253
 virtual machine (VM), 443–444
 Virtual Network Computing (VNC), 429
 Virtual network function (VNF), 444
 virtual private cloud (VPC), G-22, 444
 virtual private network (VPN), G-22, 107, 400, 417–424
 clientless VPNs, 423
 client-to-site VPNs, 422–423
 cloud connectivity, 445–446
 exam objectives for Exam N10-009, A-11
 Internet Key Exchange, 420–421
 Internet Protocol Security, 419–420
 remote network access, 417–418
 site-to-site VPNs, 424
 tunneling protocols, 418
 virtual private clouds, 445
 virtual private server (VPS), 241

Virtual Router Redundancy Protocol (VRRP), 262
 virtual terminal, 429
 viruses, 318
 visio cafe.com, 273
 visual fault locator, G-22, 71
 VLAN hopping, G-22, 322
 VLAN ID (VID), G-13, 184–185, 186, 253
 VLANs 2–1001 (normal range), 185
 VLANs 006–4094 (extended), 185
 VLANs 1002–1005 (reserved), 185
 voice and video services, 251–253. *see also* Voice over Internet Protocol (VoIP)
 voice or auxiliary VLAN, 187–188
 Voice over Internet Protocol (VoIP), G-22, 251–253
 bandwidth requirements, 298
 handsets, 101
 phones, 253
 private branch exchange, 251
 protocols, 252–253
 data transport, 252
 quality of service, 252, 253
 Real-Time Transport Protocol, 253
 RTP Control Protocol, 253
 session control, 252
 Session Initiation Protocol, 253
 voice VLANs, 187–188
 VoIP-enabled PBX, 251–252
 VoIP phones, 253
 voice virtual local area network (VLAN), G-23, 187–188
 VoIP-enabled PBX, 251–252
 VoIP phone, G-23, 253
 voltage, 62
 vulnerability, G-23, 313
 vulnerability assessment, G-23, 309, 313
 VXLAN network identifier (VNI), 454
 VXLAN Tunnel Endpoint (VTEP), 454

W

warm site, G-23, 257
 warranty support, 269
 Wavelength Division Multiplexing (WDM), 56–57
 bidirectional, 56
 coarse, 57
 dense, 57
 modular transceivers, 76, 77
 Wazuh SIEM dashboard, 291
 web application firewall (WAF), 447
 web servers, 241–242
 web services, 240–242
 HTTP Secure, 242
 HyperText Transfer Protocol, 240–242
 website performance checkers, 302
 WebSockets, 423
 wet-pipe systems, 63
 wet-to-dry method, 71
 wide area network (WAN), G-23, 4, 414–416
 fiber to the curb, 416
 fiber to the premises, 416
 guidelines for supporting, 432
 Internet access types, 415–416
 optical network terminal, 416
 OSI model, 414
 point to point link, 5
 SOHO router, 16, 17, 20, 21, 22
 very high-speed DSL, 416
 Wi-Fi. *see also* 802.11 standards
 4/5/6 standards, 407
 analyzer, 404
 encryption standards, 396–398
 tester hardware devices, 404
 Wi-Fi 5 (802.11ac), 384
 Wi-Fi 6 (802.11ax), 384, 397
 Wi-Fi Alliance, 381
 Wi-Fi analyzer, G-23, 404
 Wi-Fi Direct, 395
 Wi-Fi Enhanced Open, 398

Wi-Fi Protected Access (WPA), G-23, 396–397
 Windows
 Active Directory, 224, 225, 226, 336, 339, 345
 authentication, 335–336
 Automatic Private IP Addressing, 211–212, 323, 324
 canonical name records, 223
 client ports, 199
 DHCP issues, 216
 dig command, 232
 DNS poisoning, 325
 DNS Security Extensions, 228
 domain controller, 427
 host records, 223
 local sign-in, 335
 netstat command, 203–204
 network sign-in, 335
 Nmap Security Scanner, 277
 patches, 270
 PowerShell, 122, 232, 319
 remote sign-in, 336
 Secure Shell, 426
 single sign-on, 336
 Start of Authority record, 222
 system and application events, 289
 Teredo protocol, 145
 tunneling, 145
 User Account Control, 342
 VoIP PBX, 251
 WDI driver model, 395
 Windows Management Instrumentation (WMI), 319
 WinSCP SFTP client, 244
 wired equivalent privacy (WEP), 396–397
 wireless access point (AP), 11, 17
 wireless concepts and standards, 380–387
 2.4 GHz channel bandwidth, 382–383
 5 GHz channel bandwidth, 381–382
 802.11 standards, 380–384
 band steering, 385

- cellular technologies, 386
 - channel bonding, 383–384
 - exam objectives for Exam N10-009, A-7
 - multiple input multiple output, 383–384
 - Multiuser MIMO, 385
 - satellite technologies, 386–387
 - Wi-Fi 5 (802.11ac), 384
 - Wi-Fi 6 (802.11ax), 384
 - wireless controller, G-23, 393–394
 - wireless distribution system (WDS), 393
 - wireless intrusion detection system (WIDS), 401
 - wireless intrusion prevention system (WIPS), 401
 - wireless local area network (WLAN), G-23, 381
 - wireless media, 10
 - wireless mesh network (WMN), G-23, 395
 - wireless networks, 410–411
 - attacks, 400–402
 - deauthentication attack, 402
 - evil twin, 401
 - rogue access point, 400
 - enterprise wireless network design, 388–395
 - wireless concepts and standards, 380–387
 - wireless security, 396–402
 - wireless troubleshooting, 403–411
 - wireless performance assessment, 403–404
 - wireless roaming, 392–393
 - wireless security, 396–402
 - Bring Your Own Device, 400
 - captive portals, 399–400
 - enterprise authentication, 398–399
 - guest networks, 399–400
 - personal authentication, 398
 - Wi-Fi encryption standards, 396–398
 - wireless network attacks, 400–402
 - wireless survey, 391
 - wireless troubleshooting, 403–411
 - channel overlap issues, 406–407
 - client disassociation issues, 409
 - guidelines, 410–411
 - insufficient wireless coverage issues, 405–406
 - interference issues, 407–408
 - overcapacity issues, 409
 - roaming misconfiguration issues, 408
 - wireless performance assessment, 403–404
 - wire map tester, G-23, 68–69
 - Wireshark, G-23, 80, 295–296, 299, 302, 321
 - wiring diagram, G-23, 272
 - wiring implementation, 45–50
 - patch panels, 48–49
 - structured cable installation, 49
 - structured cabling system, 45–47
 - T568A and T568B termination standards, 47–48
 - termination tools and techniques, 50
 - work area, 45
 - workflow and process automation systems, 370–371
 - workgroup, 2
 - workgroup switch, 85
 - work recovery time (WRT), G-23, 256
 - WorldWide Names (WWN), 437
 - WorldWide Node Name (WWNN), 437
 - WorldWide Port Name (WWPN), 437
 - worms, 318
 - WPA2, 397, 398
 - four-way handshake, 399
 - open authentication, 399
 - PSK, 398
 - WPA3, 398
 - Enterprise, 398–399
 - features, 397–398
 - SAE, 397, 398, 399
 - W-Fi Enhanced Open mechanism, 399
- ## X
- X.500 directory, 345
 - xBASE-y, 35
 - Xirrus Wi-Fi Inspector, 401
- ## Y
- Yagi antenna, 394
 - YAML Ain't Markup Language (YAML), G-23, 122
- ## Z
- Zenmap, 277, 278, 279
 - Zeroconf, 323, 324
 - zero-day vulnerabilities and exploits, G-23, 313
 - zero filling, 271
 - zero-touch provisioning, SDN, 453
 - zero trust architecture (ZTA), G-23, 435, 457–458
 - zone-based security, 364–368
 - intrusion detection system, 367–368
 - intrusion prevention system, 368
 - network security zones, 364–365
 - perimeter networks, 365
 - screened subnets, 366–367
 - zone index, G-23, 143
 - zones
 - DNS, 225–226
 - network security, 364–365
 - Zone Signing Key, 227
 - zone transfer, G-23, 226