



APPENDIX B

CCNA 200-301 Official Cert Guide, Volume 2 Exam Updates

Over time, reader feedback enables Pearson to gauge which topics give our readers the most problems when taking the exams. To assist readers with those topics, the authors create new materials clarifying and expanding on those troublesome exam topics. As mentioned in the Introduction, the additional content about the exam is contained in a PDF on this book's companion website, at www.ciscopress.com/title/9781587147135.

This appendix provides you with updated information if Cisco makes minor modifications to the exam topics during the life of the 200-301 exam. In particular, this appendix does the following:

- Mentions technical items that might not have been mentioned elsewhere in the book
- Covers new topics if Cisco adds new content to the exam over time
- Provides a way to get up-to-the-minute current information about content for the exam

Note that this appendix shows updated information related to the subset of CCNA 200-301 exam topics covered in this book. Refer also to the *CCNA 200-301 Official Cert Guide, Volume 1*, for more details about the rest of the exam topics and for an Appendix B similar to that of this book.

Always Get the Latest at the Book's Product Page

Many of you are reading the version of this appendix that was available when your book was printed or when you downloaded the e-book. However, given that the main purpose of this appendix is to be a living, changing document, it is important that you look for the latest version online at the book's companion website. To do so, follow these steps:

- Step 1.** Browse to www.ciscopress.com/title/9781587147135.
- Step 2.** Click the **Updates** tab.
- Step 3.** If there is a new Appendix B document on the page, download the latest Appendix B document.

NOTE The downloaded document has a version number. Comparing the version of the print Appendix B (**Version 1.0**) with the latest downloadable version of this appendix, you should do the following:

- **Same version:** Ignore the PDF that you downloaded from the companion website.
- **Website has a later version:** Ignore this Appendix B in your book and read only the latest version that you downloaded from the companion website.

Technical Content

This appendix may be updated over time. To that end, we assign version numbers to it. This document is at Version 1.1. Reviewing the version history:

Version 1.0: Version 1.0 of this appendix was the original version of this appendix as of the publication of the book. It held no technical content.

Version 1.1: This document. Published in the second quarter of calendar year 2023, this version adds technology content plus information about exam topic changes, per the following list:

- Updates to the CCNA 200-301 Version 1.0 Exam Blueprint
- AAA Addendum
- LLDP MED
- FHRP Functions and Concepts
- Ethernet Cabling Addendum
- IPsec Remote Access VPNs
- Containers and VRFs

The rest of this Version 1.1 adds one section for each topic in the preceding list.

NOTE The following section should be read first, once you learn of this updated appendix, no matter where you are in your progression through reading the books.

Updates to the CCNA 200-301 Version 1.0 Exam Topics

Before you panic, Cisco did not add a lot of new technologies to the CCNA blueprint when it changed the exam blueprint in 2022. It did, however, reword some exam topics in 2022, so it is useful to look at those changes, reflect, and potentially study a few more topics. This section discusses the details.

NOTE This section is identical in the Version 1.1 Exam Updates appendix of both the Volume 1 and Volume 2 books.

Exam Topic Rewording Gives Us Insights

Cisco tells us upfront that it might change the exam topics. If you go to www.cisco.com/go/ccna, click to find the exam topic page, and download the PDF version of the exam topics, you'll find a couple of paragraphs above the list of exam topics. Many people ignore those paragraphs. Interestingly, you find this statement in them:

“To better reflect the contents of the exam and for clarity purposes, the guidelines below may change at any time without notice.” (Emphasis is mine.)

So, Cisco could leave the exam topics unchanged. Instead, it made slight changes to the wording. Why? There are two reasons, both of which give us a better understanding of the meaning of the exam topics. Personally, I always prefer to know more about what might be on an exam, and Cisco has done that with these exam topic changes.

From a practice point of view, with the changes the question becomes: What else do you need to study? That becomes a question of what, if anything, the exam topic changes add to CCNA. Then, if added, do the books already cover them? As it turns out, there are a few small topics that I think the book needs in reaction to the exam topic changes, and you will find them in the Exam Updates appendix online for the Volume 1 and Volume 2 books, respectively. This section walks through the details of what I've added.

The Exam Topic Changes

First, Cisco numbers exam topic documents, called exam blueprints, with a two-number convention of *version.release*. This book was originally written for the CCNA 200-301 exam, specifically blueprint version 1.0. That is, we wrote this book, published in late 2019, just before the CCNA 200-301 exam released in February, 2020.

Keep in mind that Cisco uses a style for exam topics with brief general wording. As a result, course authors, video creators, book authors, and even the folks creating the exam need to further interpret the exam topics. In my estimation—which might differ from yours or anyone else's—less than half the exam topic wording changes change the literal meaning of the exam topic. For those that have a change in meaning, some shrink the scope of the exam topic. Some expand the topic, but the books already cover the topic well enough. But a few impact what you need to study when using the two-volume *CCNA 200-301 Official Cert Guides*.

Table B-1 lists the exam topic wording changes Cisco made to the CCNA 200-301 Exam Blueprint in 2022. Note that the rows with gray highlights receive more attention in the following section.

Table B-1 Exam Topic Wording Changes, 2022, CCNA 200-301 V1.0

Number	New Wording	Old Wording
1.1.h	PoE	No wording change. The change renumbers from 1.3.c to 1.1.h.
1.2.a	Two-tier	2 tier
1.2.b	Three-tier	3 tier
1.9.a	Unicast (global, unique local, and link local)	Consolidates three exam topics: 1.9.a Global unicast 1.9.b Unique local 1.9.c Link local (Also renumbers 1.9.d, e, f to 1.9.b, c, d)
1.12	Explain virtualization fundamentals (server virtualization, containers, and VRFs)	Explain virtualization fundamentals (virtual machines)
2.1.c	InterVLAN connectivity	Connectivity

Number	New Wording	Old Wording
2.5	Interpret basic operations of Rapid PVST+ Spanning Tree Protocol	Describe the need for and basic operations of Rapid PVST+ Spanning Tree Protocol and identify basic operations
2.6	Describe Cisco Wireless Architectures and AP modes	Compare Cisco Wireless Architectures and AP modes
2.9	Interpret the wireless LAN GUI configuration for client connectivity, such as WLAN creation, security settings, QoS profiles, and advanced settings	Configure the components of a wireless LAN access for client connectivity using GUI only such as WLAN creation, security settings, QoS profiles, and advanced WLAN settings
3.2.a	Longest prefix match	Longest match
3.5	Describe the purpose, functions, and concepts of first hop redundancy protocols	Describe the purpose of first hop redundancy protocol
5.3	Configure and verify device access control using local passwords	Configure device access control using local passwords
5.5	Describe IPsec remote access and site-to-site VPNs	Describe remote access and site-to-site VPNs
5.8	Compare authentication, authorization, and accounting concepts	Differentiate authentication, authorization, and accounting concepts
5.10	Configure and verify WLAN within the GUI using WPA2 PSK	Configure WLAN using WPA2 PSK using the GUI
6.3	Describe controller-based, software defined architecture (overlay, underlay, and fabric)	Describe controller-based and software defined architectures (overlay, underlay, and fabric)
6.3.b	Northbound and Southbound APIs	North-bound and south-bound APIs
6.7	Recognize components of JSON-encoded data	Interpret JSON encoded data

The *CCNA 200-301 Official Cert Guides* (both Volume 1 and Volume 2) have extensive references to the original exam topics. Be aware: We will not be changing the original exam topic wording throughout the books; instead, now that you know about the small set of wording changes, you can refer to Table B-1 as needed to review the wording.

Note that the Cisco website does not list the old and new wording. Instead, when they changed the wording, it changed, with no announcement—and no one really noticed. Most of the changes simply clarify the meaning of the old wording; however, a few of the changes make me think you need to study one or two additional small topics. So, let me walk you through the changes.

Analysis of Wording Changes

I will review with you the plain meaning of the words in the exam topic, but let me tell you what I mean by that. Almost anyone who works with an exam—to write questions or create learning content—has to take the brief wording of an exam topic and create a much more detailed mental model of what is and is not included. For example, in my books, there might be a chapter that exists for a single exam topic that has a dozen words in it. Obviously, I had to expand on the meaning of that one tiny exam topic to choose what to include (or not include) in that chapter. So the comments here are about changes to the plain meaning of the words rather than my interpretation of them, or Cisco's, or anyone else's.

Now look back at Table B-1, noting the rows that do not have a gray highlight. For those, I think the plain meaning remains unchanged, or maybe it changes slightly to reduce the scope of the exam topic—meaning you have nothing more to do.

For instance, 1.2.a, 1.2.b, and 1.9.a are all rewording or changing conventions. Or, for 2.6 and 2.9, the verbs change the plain meaning ever so slightly—but to reduce what you might need to know.

However, in the rows with gray highlights, I think the plain meaning changes, and for some of those, when using these books, I think you have more to study.

ET 1.12: Explain virtualization fundamentals (server virtualization, containers, and VRFs)

This exam topic replaces “virtual machines” with “server virtualization, containers, and VRFs.” That changes the plain meaning of the exam topic. But more importantly, these books did not originally discuss containers and VRFs.

Your action: Make sure you read the new content about Containers and VRFs in the Version 1.1 Appendix B, “CCNA 200-301 Volume 2 Exam Updates,” in Volume 2.

ET 2.1.c: InterVLAN Connectivity

This exam topic replaces “Connectivity” with “InterVLAN connectivity.” I am thrilled about this change because it takes a formerly ambiguous exam topic and makes it clear to someone who's creating learning content.

If you're just studying CCNA, “InterVLAN Connectivity” might not be clear until you learn the topic. To add some meaning:

1. Layer 2 switches create VLANs.
2. Layer 2 switching logic does not forward Layer 2 Ethernet frames between VLANs.
3. IP routing creates connectivity between VLANs by using logic at the next higher layer (the network layer) to route packets.
4. The devices in one VLAN reside in one IP subnet; the devices in another VLAN reside in another subnet, so you use IP routing to route the IP packets between the subnets.

In short, “InterVLAN connectivity” refers to supporting communications between devices in different VLANs by routing between the subnets that exist on those VLANs.

Your action: Nothing more. Volume 1 Chapter 16, “Configuring IPv4 Addresses and Static Routes” (the whole chapter), explains interVLAN connectivity in depth. If you are tracking your progress with exam topics, you should associate exam topic 2.1.c with Chapter 16 as well.

ET 3.5 Describe the purpose, functions, and concepts of first hop redundancy protocols

This exam topic replaces “purpose” with “purpose, functions, and concepts” in a clear expansion of the meaning. To prepare you for that expansion, read the new section about First Hop Redundancy Protocols (FHRPs) in the Version 1.1 Volume 2 “Exam Updates” appendix.

As for the change itself, all FHRPs have the same purpose, so the existing *CCNA 200-301 Official Cert Guide, Volume 2*, Chapter 12 devotes one major section, about seven pages, to FHRPs. That section focuses on one FHRP—HSRP—in part because you can understand the purpose of all by understanding one of the three FHRPs.

Your action: Read the FHRP coverage in the new Volume 2 Exam Updates appendix. The wording change to add “functions and concepts” begs for some additional detail about all three FHRPs beyond Volume 2’s Chapter 12.

ET 5.5 Describe IPsec remote access and site-to-site VPNs

This exam topic keeps the same wording as before, except that it injects the word *IPsec*. VPNs include several protocols, with IPsec being one. Focusing on the plain meaning, the new wording limits the scope versus the old wording, keeping the discussion to only IPsec-based VPNs.

However, from a practical perspective, the *CCNA 200-301 Official Cert Guide, Volume 2*, Chapter 14, “WAN Architecture,” happens to show an IPsec site-to-site VPN but not an IPsec remote access VPN. (It shows a TLS-based remote access VPN instead.)

Your action: To round out the coverage, read the new Version 1.1 of Volume 2’s “Exam Updates” appendix about IPsec remote access VPNs.

ET 5.10 Configure and verify WLAN within the GUI using WPA2 PSK

This exam topic keeps the same wording other than changing “Configure” to “Configure and verify.” I think that expands the plain meaning; however, from a practical perspective, I think it matters very little. It is almost impossible to separate the configuration and verification tasks when writing books or creating other learning content. It is also difficult to separate them when learning. So in the rare cases in which a Cisco exam topic lists only *configure*, I tend to create materials that also help you learn to verify the topic.

However, I already wanted to clarify a few points about wireless LANs in the updated Exam Updates elements. So, we added some new content, emphasizing what to look for when verifying wireless LAN configuration using the WLC.

Your action: In the new Version 1.1 of Volume 1’s Exam Updates appendix, read the section, “Wireless LAN Configuration Verification and Analysis.”

(As an aside, note that exam topic 5.3 also has the same change—from “configure” to “configure and verify”—so it probably has a literal change to the plain meaning as well. But the existing content already covered the verification angle well, so I didn’t list it here as an exam topic that deserved more analysis.)

NOTE The following section should be read after reading or reviewing Chapter 4’s section titled “Controlling and Monitoring User Access,” just after Figure 4-9, before the content in the section, “Developing a Security Program to Educate Users.”

AAA Addendum

(For topics in this appendix whose titles end with “Addendum,” a stronger need exists for you to look at the related context from the chapter. In other words, the extra content here makes more sense if you take the time to read or review the topics listed in the note just above the heading “AAA Addendum.”)

Comparing RADIUS and TACACS+, RADIUS tends to be used more often for end-user AAA services, while TACACS+ (created by Cisco) is used more often to protect networking devices. TACACS+ can authorize the specific CLI commands allowed by a user, whereas RADIUS does not, making it useful for AAA with devices that have a CLI. Table B-2 summarizes the differences between the two protocols.

Table B-2 Comparisons Between TACACS+ and RADIUS

Features	TACACS+	RADIUS
Most often used for	Network devices	Users
Transport protocol	TCP	UDP
Authentication port number(s)	49	1645, 1812
Protocol encrypts the password	Yes	Yes
Protocol encrypts entire packet	Yes	No
Supports function to authorize each user to a subset of CLI commands	Yes	No
Defined by	Cisco	RFC 2865

Also, to review and expand a bit upon the meaning of the AAA acronym, you can think of AAA in the following manner:

- **Authentication:** Who is the user?
- **Authorization:** What is the user allowed to do?
- **Accounting:** What did the user do?

AAA begins with centralized management of login credentials used for authentication. *Authentication* refers to confirming whether the user is who they claim to be. The user can supply information (username and password) or use something they have (like responding to a text from their phone).

Using a centralized authentication process has many advantages over the distributed configuration options found in network device configuration. For instance, users should update passwords from time to time. Companies that rely on per-device password configuration on thousands of networking devices seldom systematically update those passwords. A centralized AAA system would solve that problem.

The second A in AAA, *authorization*, defines the capabilities allowed for the user. For example, on routers and switches, you should be familiar with user mode, reached via simple login, and privileged mode, reached with the **enable** command. In effect, these modes act as two authorization levels. IOS supports additional authorization levels you can configure, defining different commands allowed at each level.

As an example, user mode does not allow the **show running-config** or **configure terminal** commands. With IOS authorization features, you could define a new level that includes all user mode commands plus the **show running-config** command, which lets the user see the configuration. However, you might not want to allow the **configure terminal** command, which would let the user change the configuration. That authorization level can be associated with the user as identified during the authentication process.

Accounting, the third A in AAA, defines the need to record information about user activity. For network device users, that often comes in the form of log messages (as discussed further in Chapter 9, “Device Management Protocols.”) Devices generate information messages—log messages—when events of note occur. Some log messages track user actions, like configuring the device or reloading it.

NOTE The following section should be read at the end of Chapter 9, just before “Chapter Review.”

LLDP MED

The LLDP 802.1AB standard defines a method to expand LLDP for new functions using a type-length-value (TLV) concept. The term refers to three fields in a data structure. The first defines the type of data held there, the length lists the length of the TLV, and the value lists the data. Any device processing received messages with TLVs can quickly identify one TLV after another based on the specified length.

Figure B-1 shows a sample LLDP message. Devices encapsulate LLDP messages directly in the data link protocol. The LLDP message consists of a series of TLVs.




Figure B-1 Type-Length-Value Concept in LLDP Messages

Over time, new endpoint-focused LLDP TLVs emerged. The LLDP Media Endpoint Discovery protocol (LLDP-MED) uses a variety of endpoint-focused LLDP TLVs useful for operations between a switch and an endpoint device. (While IEEE standard 802.1AB defines LLDP, the Telecommunications Industry Association [TIA] defines LLDP-MED as an industry standard in its TIA-1057 document.) For instance, Chapter 17, “Cisco Software-Defined Access (SDA),” discusses Power over Ethernet (PoE), with devices using LLDP-MED to exchange data about the power needed by the endpoint device.

IP Phones connected to LAN switches can also use LLDP-MED or CDP to learn information such as the voice and data VLAN IDs used by the switch. Figure B-2 shows an example for context, with three phones, each with a connected PC. The design uses VLAN 20 as the voice VLAN and VLAN 21 as the data VLAN. (Refer to the *CCNA 200-301 Official Cert Guide, Volume 1*, Chapter 8, Figure 8-13, and surrounding text to review voice and data VLAN configuration.)




Figure B-2 Sample Access Switch with IP Phones

The phones need to learn which voice VLAN to use because the phone adds an 802.1Q trunking header to any Ethernet frames generated by the phone. You do not preconfigure the phone with the voice VLAN ID; instead, the phone learns it using LLDP-MED or CDP. Originally, Cisco IP Phones supported only CDP to learn the voice VLAN ID. Cisco IP Phones have supported both LLDP-MED and CDP for more than a decade. Non-Cisco IP phones, which do not support Cisco-proprietary CDP, use only LLDP-MED.

On a final note about LLDP-MED, you configure LLDP-MED with the same commands as LLDP. The devices on the link dynamically discover that they need to include the TLVs defined by LLDP-MED—no additional configuration is required.

Aside: CDP and LLDP Timers

As a brief aside, note that both CDP and LLDP use a hello and hold timer to manage the information they advertise and keep. The hello time tells the device how frequently to announce CDP details on the interface. The hold time tells a device how long to wait, if no longer hearing from a neighbor, before clearing its CDP data learned from a neighbor.

For instance, with CDP’s default hello and hold timers of 60 and 180 seconds, respectively, each neighbor on a link would send CDP messages every 60 seconds. On receipt of a message, that device would set its per-neighbor hold timer to 180 and count down. If there was no packet loss, over time, that hold timer will count down from 180 toward 120 and get reset back to 180 on receipt of the next CDP message from that neighbor. (Chapter 9’s Example 9-15 displays the CDP holdtime for two neighbors, both between 120 and 180.)

CDP and LLDP use the same timer concepts, with per-interface settings and per-neighbor holdtimes tracked by IOS. LLDP defaults to 30 and 120 seconds, respectively. You can override the defaults with the commands listed in Table B-3.

Table B-3 CDP and LLDP Timer Configuration

CDP Command	LLDP Command	Description
<code>cdp timer</code> <i>seconds</i>	<code>lldp timer</code> <i>seconds</i>	Defines how often CDP or LLDP sends messages on each interface
<code>cdp holdtime</code> <i>seconds</i>	<code>lldp holdtime</code> <i>seconds</i>	Defines how long to wait after the most recent incoming message from a neighbor before deleting that neighbor’s information

NOTE The following section should be read after reading or reviewing Chapter 12's Figure 12-7, just before starting the next major section at the heading "Simple Network Management Protocol."

B

FHRP Functions and Concepts

Although all three FHRPs have the same purpose, they have some different functions and concepts. This section adds more content that focuses on functions and concepts for the other two FHRP options that are discussed little in Chapter 12 of the book: VRRP and GLBP. While all three share the same primary purpose, VRRP has more similarities with HSRP, while GLBP goes beyond HSRP with better load-balancing features.

Virtual Router Redundancy Protocol (VRRP)

HSRP and VRRP emerged in the 1990s when TCP/IP and routers first became common in corporate networks. As is often the case, Cisco saw a need, but with no standards-based solution, it defined HSRP as a proprietary solution for first hop router redundancy. Later, the IETF created VRRP, providing similar features. However, unlike many stories of Cisco-proprietary pre-standard features, HSRP has not faded into history—you will still find both HSRP and VRRP support in many Cisco product families.

NOTE While VRRP includes versions 1, 2, and 3, all references in this chapter refer to VRRPv3 (RFC 5798).

For similarities, note that VRRP supports all the same functions as HSRP, as described in Chapter 12. The purpose remains to provide a standby backup for the default router function and load balancing the default router role by using multiple VRRP groups. The differences come with default settings, protocol details, and addresses used. Table B-4 lists some comparison points between HSRP, VRRP, and GLBP (ignore GLBP for now.)

Table B-4 Comparing Features of the Three FHRP Options

FHRP Option	HSRPv2	VRRPv3	GLBP
Cisco Proprietary	Yes	No	Yes
VIP must differ from the routers' interface IP addresses	Yes	No	Yes
Preemption off by default	Yes	No	Yes
Allows preemption (or not)	Yes	Yes	Yes
Active/active load balancing with multiple active routers in one group	No	No	Yes
IPv4 multicast address used	224.0.0.102	224.0.0.18	224.0.0.102
Group numbers supported in IOS	0–4095	1–255	0–1023
Virtual MAC address pattern	0000.0c9f.fzzz	0000.5e00.01xx	0007.b40x.xxrr

You can configure VRRP so that it appears to work like HSRP. Two or more VRRP routers form a group within one subnet. VRRP routers define one VIP, use multicast messages to

communicate with each other, use an active/standby approach, select the active router with the same logic as HSRP, allow tracking, and failover when the master (active) router fails. (Note that VRRP uses the terms *master* and *backup* rather than active and standby.)

One difference comes in the choice of VIP. You can use the same IP address as one of the VRRP routers' interface addresses or, like HSRP, use another IP address in the subnet. For example, the Chapter 12 HSRP discussion around Figure 12-5 and Figure 12-6 uses VIP 10.1.1.1, with router addresses 10.1.1.9 and 10.1.1.129. You could do the same with VRRP or use 10.1.1.9 (the same IP address as router R1's interface IP address).

All three FHRP tools use priority and preemption to affect the election process. When an election occurs within the group, the router with the best (highest) priority setting wins. For example, HSRP uses priority to select the router to act as the active router.

The preemption setting determines whether a new router arriving in the subnet can take over (or not) after the election. For example, Router R1 wins an election versus Router R2. Later, Router R3 connects to the same subnet. R3's preemption setting determines whether it can immediately take on the primary role. With preemption enabled, a new router can take over immediately if it has the highest priority. Without preemption, the new router must wait for a new election, which occurs when the current primary router fails.

As for VRRP, while it works much like HSRP, there are differences. It defaults to use preemption (HSRP does not). It uses a different multicast IPv4 address (224.0.0.18) for its messages. Like HSRP, it uses a single virtual MAC per group, but with a different pattern to form the MAC address. The VRRP virtual MAC address uses the hex equivalent (two-digit) of the configured decimal VRRP group number at the end of the virtual MAC, using this pattern:

VRRPv3: 0000.5e00.01xx, where xx is the hex group number

In comparison, HSRP has had two versions over time, with HSRPv2 now decades old and commonly used. HSRPv1, like VRRP, allows 256 groups, so the hex-equivalent values require two digits. HSRPv2 allows more groups, so it needs a three-digit hex value to represent the group number in the virtual MAC address. HSRP uses these patterns:

HSRPv1: 0000.0C07.ACxx, where xx is the hex group number

HSRPv2: 0000.0C9F.Fxxx, where xxx is the hex group number

GLBP Concepts

Cisco-proprietary Gateway Load Balancing Protocol (GLBP), defined later than HSRP and VRRP, provides the same benefits as HSRP and VRRP but with different implementation details. But it also includes different internals that allow much more effective load balancing. So, while used for redundancy (the "R" in FHRP), GLBP also adds robust load balancing, per its name.

This GLBP section begins with comparisons to the other FHRPs and then discusses its improved approach to load balancing.

Similarities of GLBP, HSRP, and VRRP

Like HSRP and VRRP, GLBP provides redundancy for the default router function while hiding that redundancy from the hosts using that default router address. The core features follow a familiar theme:

- It uses a virtual IP address (VIP), which is the address used by endpoints as their default router.
- It identifies the best router in the group based on the highest priority.
- It allows for the preemption of the best router when a new router with a better (higher) priority joins the group.
- It sends messages using multicasts but uses a different address: 224.0.0.102.

However, GLBP uses virtual MAC addresses differently than the other FHRPs as part of the underlying support for load balancing. Like HSRP and VRRP, a GLBP group has one VIP. Unlike HSRP and VRRP, the routers in a group do not use one virtual MAC address whose function resides with the one active router. Instead, GLBP uses a unique virtual MAC address per GLBP router.

The GLBP MAC address value needs three hex digits to represent the hex equivalent of the decimal GLBP group number. It also assigns a unique last two digits (01, 02, 03, or 04) for the up to four allowed GLBP routers in a group. The MAC address pattern is 0007.b40x.xxrr. For instance, for two routers in the same GLBP group:

Router R1: 0007:b400:1401 (decimal group 20, which is hex group 014, assigned router number 01)

Router R2: 0007:b400:1402 (decimal group 20, which is hex group 014, assigned router number 02)

GLBP Active/Active Load Balancing

With a name like Gateway Load Balancing Protocol, you would expect load balancing to be an important feature. The term *gateway* refers to the alternative term for default router (default gateway), so, by name, GLBP claims to load balance across the default routers in a subnet—and it does.

GLBP manipulates the hosts' IP ARP tables in a subnet so that some hosts forward packets to one router and some to another. As usual, all the hosts use the same VIP as their default router address. Under normal conditions, with multiple GLBP routers working in the subnet, GLBP spreads the default router workload across all GLBP group members. When one of those routers fails, GLBP defines the methods by which the remaining router(s) takes over the role of the failed router.

To achieve this active/active load balancing, one GLBP router performs the role of active virtual gateway (AVG). The AVG handles all ARP functions for the VIP. Knowing the virtual MAC addresses of all the routers in the group, the AVG replies to some ARP requests with one virtual MAC and some with the other. As a result, some hosts in the subnet send frames to the Ethernet MAC address of one of the routers, with different hosts sending their frames to the MAC address of the second router.

All routers serve as a GLBP active virtual forwarder (AVF) to support load balancing. All the AVFs sit ready to receive Ethernet frames addressed to their unique virtual MAC addresses and to route the encapsulated packets as usual. Note that the AVG router also serves as an AVF.

Figures B-3 and B-4 show the results of two ARP Reply messages from AVG R1. First, Figure B-3 shows how a GLBP balances traffic for host A based on the ARP Reply sent by the AVG (R1). The two AVF routers support virtual IP address 10.1.1.1, with the hosts using that address as their default router setting.

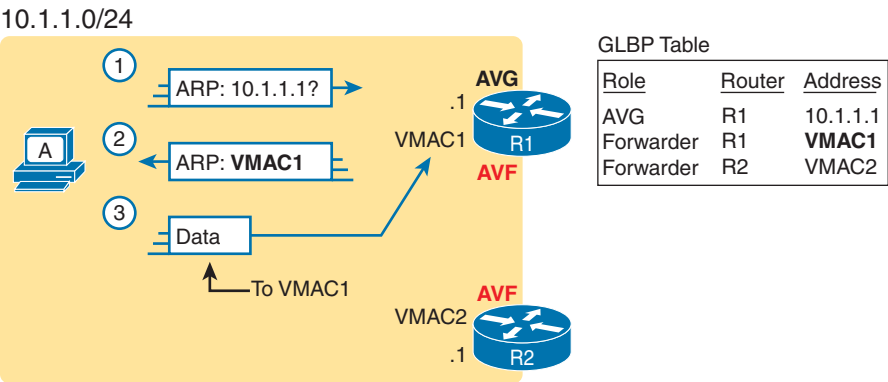


Figure B-3 GLBP Directs Host A by Sending Back ARP Reply with R1’s VMAC1

The figure shows three messages, top to bottom, with the following action:

1. Host A has no ARP table entry for its default router, 10.1.1.1, so host A sends an ARP Request to learn 10.1.1.1’s MAC address.
2. The GLBP AVG, R1 in this case, sends back an ARP Reply. The AVG includes its virtual MAC address in the ARP Reply, VMAC1.
3. Host A encapsulates future IP packets in Ethernet frames destined for VMAC1, so they arrive at R1 (also an AVF.)

To balance the load, the AVG answers each new ARP Request with the MAC addresses of alternating routers. Figure B-4 continues the load-balancing effect with host B’s ARP Request for 10.1.1.1. The router acting as AVG (R1) still sends the ARP Reply, but this time with R2’s virtual MAC (VMAC2).

Here are the steps in the figure:

1. Host B sends an ARP Request to learn 10.1.1.1’s MAC address.
2. The GLBP AVG (R1) sends back an ARP Reply, listing VMAC2, R2’s virtual MAC address.
3. Host B encapsulates future IP packets in Ethernet frames destined for VMAC2, so they arrive at R2.




Figure B-4 GLBP Directs Host B by Sending Back ARP Reply with R2's VMAC2

Finally, to deal with router failures, if the AVG fails, the remaining routers in a GLBP group elect a new AVG. When a router serving as only an AVF fails, the AVG recognizes the failure and causes a still-functional AVF to begin receiving frames sent to the failed AVF's virtual MAC address.

NOTE The following section should be read after reading or reviewing Chapter 13's section titled "Topology Design Terminology," just before beginning the section titled "Small Office/Home Office."

Ethernet Cabling Addendum

(For topics in this appendix whose titles end with "Addendum," the extra content here makes more sense if you take the time to read or review the topics listed in the note just above heading "Ethernet Cabling Addendum.")

In the two-tier and three-tier designs discussed in this chapter, the individual links use the Ethernet data link protocol, supporting Ethernet frames—but often use different physical layer standards. Those standards define the maximum link segment length and transmission speeds. This next section takes a closer look at some of the considerations and decisions when choosing the cabling to use in a campus LAN, first focusing on access links, which often use UTP cabling, followed by distribution and core links, which use UTP and fiber optic links.

Ethernet UTP Links at the Access Layer

Most, if not all, access links use unshielded twisted pair (UTP) cabling—and that is no accident. From the early days of Ethernet, the IEEE set about to ensure that Ethernet would be commercially viable, making it highly useful and affordable, with UTP cabling as the least expensive option. For example:

- Knowing that UTP cabling costs less than the other options, the IEEE emphasized standards that supported UTP for the most common links in a typical network: access links.
- Early studies showed that a 100-meter cable could reach from the wiring closet to any point on the floor for most office buildings.

- Therefore, the IEEE used a 100-meter maximum length convention in its successive UTP-based standards over the years.

Most building construction plans include a structured cabling system. On each building floor, the cables run from a wiring panel in a central wiring closet to most locations around the floor. For example, on a floor used for office cubicles, the structured cabling system includes UTP cables from the wiring closet to every cubicle location, often terminating with an RJ-45 connector in a wall plate.

Not all UTP cables have the same physical transmission characteristics, so the cabling industry and standards bodies have long-defined standards for the UTP cables used by Ethernet. Ethernet standards refer to UTP cable rating categories defined by the Telecommunications Industry Association (TIA) and the US American National Standards Institute (TIA/ANSI). These categories define the physical transmission qualities when using the cable.

While the TIA and ANSI do not define Ethernet, and the IEEE does not define cabling standards, the Ethernet UTP-based standards refer to the minimum quality of UTP cable category that supports each Ethernet standard. For the TIA/ANSI categories, the higher the number and letter, the higher the quality, the more recent the standard release, and the more recent the IEEE standard that uses the cabling. The categories include CAT 3, CAT 5, CAT 5E, CAT 6, CAT 6A, and CAT 8.

Table B-5 lists the data about the cable categories and matching Ethernet standards. Note that the higher the cable category number/letter, the better for supporting faster Ethernet standards. Also, the table does not list every combination, but it does list standards that support a 100-meter cable length.

Table B-5 Ethernet and Cable Standards to Support a 100-Meter Segment

Standard (Common)	Standard (Original Document)	Year of Standard	(Minimum) ANSI/TIA Category	Max Speed
10BASE-T	802.3	1990	CAT 3	10 Mbps
100BASE-T	802.3u	1995	CAT 5	100 Mbps
1000BASE-T	802.3ab	1999	CAT 5E	1 Gbps
10GBASE-T	802.3an	2006	CAT 6A ¹	10 Gbps
40GBASE-T	802.3ba	2010	CAT 8	40 Gbps
2.5GBASE-T	802.3bz	2016	CAT 5E	2.5 Gbps
5GBASE-T	802.3bz	2016	CAT 5E	5 Gbps

¹ 10GBASE-T can also use a CAT 6 cable, but with a distance limit of 55 meters.

The table lists the minimum cable category, with better cable categories also working. For example, the original 10BASE-T worked on then-current CAT 3 cabling (or better). 100BASE-T, the next IEEE UTP Ethernet standard, required CAT 5 (or better). The following Ethernet UTP standard, 1000BASE-T, required even better cabling (CAT 5E) to reach 100 meters.

Today, the IEEE and others work hard to improve Ethernet standards for UTP cabling, with that 100-meter access link as a critical design point. As shown by the dates in Table B-5, about every five years from 1990–2010, the IEEE supplied a new UTP-based standard faster

than UTP. Those standards deliver plenty of speed for the access layer, assuming you have the required cable types installed.

Fiber Uplinks

For all uplinks, you must choose the required speed and the physical layer Ethernet standard. The installed cabling influences the choices because using installed cabling rather than installing new cables can significantly reduce the cost.

If possible, use UTP cables for all uplinks. For example, Table B-5 lists 10GBASE-T as needing CAT 6A UTP cabling to handle distances to 100 meters. If your structured cabling system had plenty of CAT 6A already installed, using 10GBASE-T uplinks makes great sense.

However, several factors will drive a decision to use fiber Ethernet options. Why fiber? The original cabling design might have needed cable segments longer than 100 meters, so they anticipated your choice to use Ethernet standards that use fiber optic cabling. With pre-installed fiber cabling, you can then concentrate on what Ethernet standards will work on that cabling, choosing the best switch hardware (for example, SFPs to add to the switch), and so on.

Multimode fiber cabling has quality standards akin to UTP cable categories. Those standards, called Optical Multimode (OM), include OM1, OM2, OM3, and OM4, with the higher numbers representing newer more-capable cabling standards. The cable standards define some attributes of the cabling—for example, the diameter of the core and cladding. (See the *CCNA 200-301 Official Cert Guide, Volume 1*, Chapter 2's section titled "Fiber Cabling Transmission Concepts" for some figures that show a fiber cable core and cladding.) Table B-6 lists the OM standards and some supported Ethernet standards.

Table B-6 Optical Multimode (OM) and Related Ethernet Standards

(Minimum) ISO Cable Category	Core/Cladding Diameter	1000BASE-SX Max Distance per Standard	10GBASE-SR Max Distance per Standard
OM1	62.5/125	220m	33m
OM2	50/125	550m	82m
OM3	50/125	N/A	300m
OM4	50/125	N/A	400m

When attempting to use an existing installed base of multimode fiber, determine the OM category required by the Ethernet standard to meet the desired speeds. You can then analyze the data in the table, or similar data about other Ethernet standards, to determine the maximum cable lengths supported by each Ethernet standard. (Note that vendors often suggest that their optical transceivers—SFP, SFP+, and so on—can support distances longer than the standards.) As you can see in the table, the 1000BASE-SX supports distances longer than 1000BASE-T's 100 meters, which can be helpful. The 10GBASE-SX standard supports similar distances but at 10 Gbps.

Also, use the numbers in the table for initial planning but be aware of a couple of essential points:

- The standards typically give a conservative distance estimate, but you might be able to make links work at longer distances.

- In practice, 1000BASE-SX works on OM3 and OM4. The table lists “N/A” (not applicable) because the 1000BASE-SX standard predates the OM3 and OM4 standards. But formally, the 1000BASE-SX part of the Ethernet standard does not mention OM3 and OM4.

After identifying the specific cables, you can plan for the Ethernet standard and switch ports. For example, you might make a standard of buying access and distribution switches with modular ports that support 1 or 10 Gbps. You could then buy a 1000BASE-SX SFP or 10GBASE-SR SFP+ to match the standard that meets each case’s needs.

NOTE The following section should be read after reading or reviewing all of Chapter 14.

IPsec Remote Access VPNs

A site-to-site VPN exists to support multiple devices at each site, with the IT staff creating a permanent VPN connection to support all users. In contrast, user devices can dynamically initiate VPN connections in cases where a permanent site-to-site VPN does not exist. Such a VPN connection, terminated by and initiated from the end user device, is called a *remote access VPN*.

For example, a user can walk into a coffee shop and connect to the free Wi-Fi with a tablet or laptop, but that coffee shop does not have a site-to-site VPN to the user’s enterprise network. Instead, software on the endpoint device creates a secure VPN connection back to the enterprise network, as shown in Figure B-5.




Figure B-5 IPsec Remote Access IPsec VPN

Many companies use the remote access VPN model in the figure for employees doing remote work, connecting from their home, a hotel, or any off-site location. The company can choose whether to use IPsec or TLS to secure the traffic.

To support all traffic sent by the computer, the end-user device requires VPN client software; for example, the Cisco AnyConnect Secure Mobility Client. The networking staff also installs and configures a device to act as a VPN concentrator, defining whether to use TLS or IPsec. Employee devices then connect to the VPN concentrator. Although routers can play that role, companies typically use firewall products.

IPsec works a little differently when used for remote access versus site-to-site VPNs. IPsec uses *tunnel mode* for site-to-site encryption, which encrypts the entire original packet. IPsec uses *transport mode* for remote access VPNs, which encrypts the data of the original IP packet—that is, everything after the IP header—but not the IP header itself. Figure B-6 shows a visual comparison of the two modes.




Figure B-6 IPsec Tunnel and Transport Mode—What Is Encrypted

Site-to-site and remote access IPsec VPNs also have other differences, given their different roles. Table B-7 summarizes the differences for easier study.

Table B-7 Comparisons of Site-to-Site and Remote Access IPsec VPNs

Attribute	Site-to-Site IPsec VPN	Remote Access IPsec VPN
Does the end-user device need VPN client software?	No	Yes
Devices supported by one VPN: one or many?	Many	One
Typical use: on-demand or permanent?	Permanent	On-demand
Does the VPN use IPsec tunnel mode?	Yes	No
Does the VPN use IPsec transport mode?	No	Yes

NOTE The following section should be read after reading or reviewing through Chapter 15's Figure 15-4, just before the heading titled "The Physical Data Center Network."

Containers

Software containers have the same goal as VMs but use different methods. Compared to VMs, containers take less CPU and memory while also taking less time to initialize and shut down—making them more appealing in some cases.

To appreciate the differences, consider VMs again for a moment. A single VM exists on disk, waiting to run, as one large file—typically many gigabytes, because it holds an entire OS. Starting a VM takes minutes—think of VM initialization time as similar to the time it takes to boot your desktop or laptop computer. Also, VMs require some work: Both the OS and the application must be installed, and you must apply software fixes over time.

Some of those perceived drawbacks of VMs led to a second wave of server virtualization, called software containers, or simply containers.

First, consider the word *container* as a generic term to understand the fundamentals. What do you imagine? Maybe you think of shipping containers that fill huge ships or ride behind tractor-trailer trucks. Perhaps you think of the plastic container you use to bring your lunch to school. Generically, containers hold other things, typically multiple things.

Software containers hold an application plus every other file it needs to work—other than the OS. For example, the primary executable file for the application is in the container. The container has all the related files that come with the app, plus any libraries (standard code often used by applications) at the prescribed versions. It includes files used by the application; for instance, maybe a text file with application settings. The container—a file with a defined format that collects and includes all the component files—holds all the files. However, because a container does not include the OS, it is usually smaller than a VM, often measured in megabytes rather than gigabytes.

NOTE By convention, the term *container image* refers to the file on disk that holds all the files that make up the application. The term *container* refers to a container after it has been started. For example, you might use a container image for a web server application. When you start that image five times on a server, you have five containers; that is, five instances of the web server running on the server.

So, if the container does not include the OS, how does it execute? It seems like something is missing. The answer: Architecturally, server hardware runs one instance of the OS—an OS that supports containers, like Windows or Linux. Additionally, you need a *container engine*: software that can install, start, stop, and monitor containers. Figure B-7 shows the general idea.

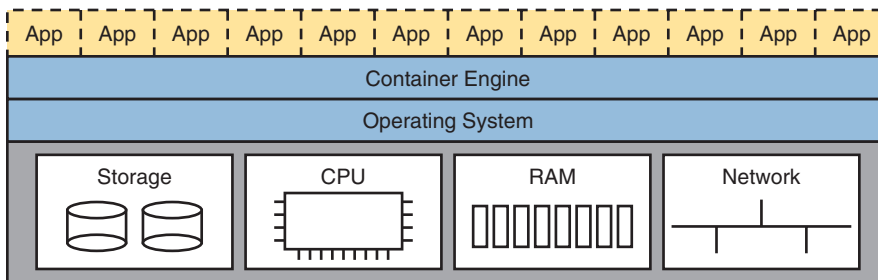


Figure B-7 Twelve Containers Running on One Host Managed by Container Engine

Starting a container requires a container engine, software that understands container file formats and operations. The container engine supplies a GUI to start, stop, and monitor the containers. However, most data center operational activities happen remotely, so the container image includes shell commands and APIs to aid remote control and automation by virtualization software. And when you start a container, it takes the time typical of starting an application (seconds or tens of seconds) rather than the minutes required to boot an OS.

Although Linux and Windows include support for containers, containers became popular in the 2010s when some companies began offering related software services. Some of the most popular and common companies and sites include

- Docker (Docker.com)
- Kubernetes (k8s.io)
- Terraform (hashicorp.com)

Docker (www.docker.com) probably had the most significant early impact on popularizing containers, with the name Docker becoming almost synonymous with containers. For instance, Docker created rules for packaging container images. It also offers the docker engine as a container engine.

Docker also helps speed application development through Docker Hub (hub.docker.com), a website that offers more than 100,000 container images. To begin developing an application, a developer can find an existing Docker container image that includes most of what the application needs. For example, the Docker container image for the world's most popular web server software (Apache) has been downloaded over one billion times per Docker Hub.

NOTE Docker has long had free and paid accounts and software. You can download Docker on your computer at no cost, download Docker containers from Docker Hub, and start running and using containers on your desktop if you want to learn more.

NOTE The following section should be read after reading or reviewing through Chapter 15's Figure 15-12, just before the heading titled "WAN Traffic Paths to Reach Cloud Services."

VRFs

Public cloud services must support many customers concurrently; however, those different customers may, and often do, use overlapping IP subnets. Many companies use private IPv4 networks internally, and the subnets they use for their private and public cloud VMs and containers use addresses from those subnets. Unsurprisingly, those customers use overlapping subnets and addresses.

Overlapping subnets causes problems for a router (or Layer 3 switch) when using traditional conventions. Working through some of the key points:

- Typically, one router has one IP routing table.
- Typically, a router allows only one interface connected to the same subnet.
- If an engineer attempts to connect a second interface to the same subnet, the router will not bring up the interface.
- Data Center virtualization software can locate and move VMs and reprogram networking, so VMs from multiple customers can exist on one physical server—creating a case of overlapping subnets within that physical server.

As an example, consider customers A and B, whose VMs reside in the server shown in Figure B-8. The customers use private class A network 10.0.0.0, and both use subnets 10.1.1.0/24 and 10.1.2.0/24. In this example, the virtual switch performs Layer 2 switching only, with no IP routing. Instead, it uses four separate VLANs, as shown. Within that limited scope, no problems exist.




Figure B-8 *Overlapping Subnets with Two Customers' VMs on the Same Server*

Next, focus on router R1, outside the host. R1 needs to have interfaces configured in the four subnets shown in the figure so it can receive data from those VMs; however, the overlapping subnets confuses R1. For example, the router and virtual switch define their link as a VLAN trunk. The router uses a router-on-a-stick (ROAS) configuration as shown in the figure (see Volume 1's Chapter 17, "IP Routing in the LAN," for more details). That configuration gives router R1 an interface with an IP address in all four subnets shown inside the server.

However, one set of R1 interfaces would fail. If R1 were to configure the statements on the left first and then the ones on the right, IOS would not bring up the two interfaces on the right side of the figure due to the overlapping subnets implied by the **ip address** subcommands.

Virtual Routing and Forwarding (VRF) instances solve this problem by expanding the logic used within a router. VRFs create multiple virtual routers inside a single router or Layer 3 switch. The router configuration associates interfaces and routing protocol neighbors to VRFs, with a separate routing table per VRF. In this case, router R1 would create two VRFs, one for each customer. On the ROAS link, R1 will assign the interfaces for VLANs 10 and 20 into one VRF and the ROAS interfaces for VLANs 30 and 40 into another.

By using a separate IP routing table per VRF, the router can support overlapping subnets by placing them in different VRFs. Figure B-9 expands Figure B-8 to represent that concept, adding VRFs to the router. It places two VLAN interfaces in each VRF and shows the per-VRF routing tables to the left and right of the router. Note the exact same subnets in each routing table, now allowed by using VRFs.




Figure B-9 *Overlapping Subnets with Two Customers' VMs on the Same Server*

You will find a use for VRFs throughout the world of networking. For example, MPLS, as discussed in this book's Chapter 14, "WAN Architecture," makes use of VRFs. One MPLS provider can support thousands of customers with the same MPLS network, even those using overlapping IP addresses, using VRFs. Summing up some of the critical points about VRFs:

- The VRF must be created via configuration in each device that performs routing (router or Layer 3 switch).
- Each router will have a separate IP routing table for each VRF, each holding routes for only that VRF.
- The configuration assigns each interface to a VRF so that the router places the associated connected route into that VRF's routing table.
- The routing protocol configuration defines VRFs to associate neighbor relationships and VRFs. Routes learned by a given routing protocol neighbor result in new routes in that VRF's routing table.
- The router keeps its original routing table, called the global routing table. The global routing table holds routes related to interfaces and routing protocol neighbors not associated with any VRF.