Winter 12-25-2018

A Companion Study Guide for the Cisco DCICN Data Center Certification Exam (200-150)

Miguel Ramlatchan
Old Dominion University, mramlatc@odu.edu

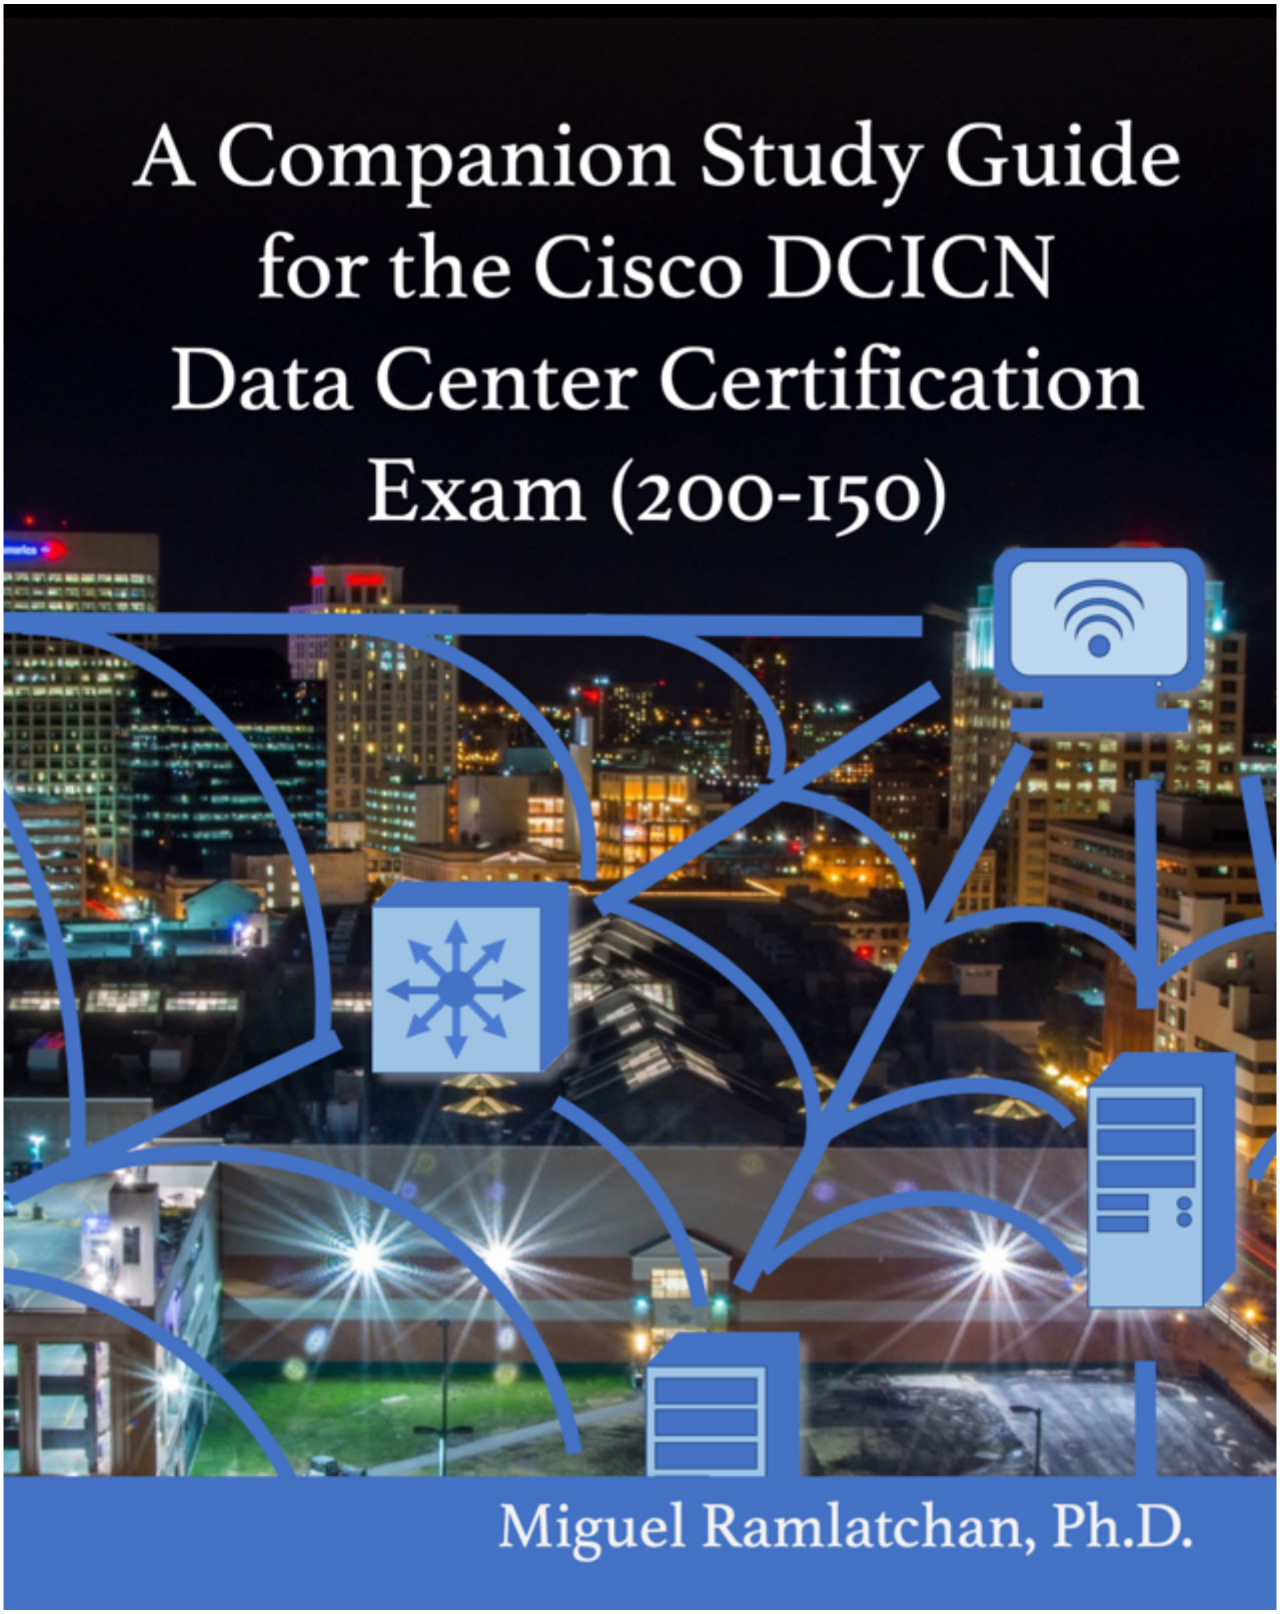

A Companion Study Guide for the Cisco DCICN Data Center Certification Exam (200-150)

Miguel Ramlatchan, Ph.D.

A Companion Study Guide for the Cisco DCICN Data Center Certification Exam (200-150)

By Miguel Ramlatchan, Ph.D.

First Edition, Rev 1.0: December 2018

The CCNA, CCENT, DCICT, CCNP, CCIE, and DCICN are all trademarks of Cisco Systems, Inc. and are only used here for educational and informational purposes. This book and this author are not endorsed by or affiliated with Cisco Systems, Inc.

Cover Design, by William I. Ramlatchan.

ISBN: 9781790260430

This book is dedicated to the proposition that
"...the harder you work ...the luckier you get..."

Acknowledgments

I'd like to thank my wife and kids for their patience and support, Andy at ODU for your encouragement, Debbie at Cisco, Media Inventive for your help, to the experienced and talented instructors I've had along the way, and thanks to my early proofreaders and technical reviewers!

About the Author

Dr. Miguel Ramlatchan is the Assistant Vice President for Technology in Old Dominion University's Office of Distance Learning. Old Dominion University is a public doctoral research institution with its main campus in Norfolk, VA., and has been a pioneer in technology-delivered distance learning since the mid-1980s. ODU has graduated over 16,000+ students through online, broadcast, video conferencing, video streaming, and telepresence technologies from across the U.S. and the world. At ODU, Miguel is responsible for the instructional systems, infrastructure, and operational and technical support services for distance and online students and instructors. The heart of this support infrastructure is a Cisco-centric data center utilizing Cisco Nexus, UCS, TelePresence, WebEx, and Catalyst systems (as well as VMware, Dell, Promise, Kaltura, AWS, and NetApp). Miguel received his Bachelor’s degree in Electrical Engineering in 1998, his Master’s degree in Engineering Management in 2000, and his Ph.D. in Instructional Design and Technology in 2016, all from Old Dominion. His research includes multimedia design theory, technology and social presence, and practical applications of audio and video interaction in online classes and programs.

Introduction

Hello,

The purpose of this book is not to replace the high-quality exam review guides, texts, or classes that are available. Instead, the goal of this book is to complement those resources and serve as additional, review, and supportive material. This book summarizes and consolidates much of the subject matter and provides additional context and content. Much of this book is based on the notes I made while studying, and is an effort to help others study (this guide is available at essentially the same costs it takes to print and ship it, this is not a money-making effort).

This book is not a "braindump"; these are not the specific questions that you will see on the DCICN exam. Instead, these examples are meant to illustrate the content and provide a source for practice and review. I found it extremely difficult to find additional, reliable, and trustworthy material to study when I was working towards my CCNA Data Center certification. The Pearson practice exams and the Cisco CCNA Data Center DCICN 200-150 Official Cert Guide by Chand Hintz, Cesar Obediente, and Ozde Karakok are excellent resources, and I highly recommend them. Also, please stay away from the many braindump websites and 'practice exams' that are widely available online. They are essentially illegal copies and screen grabs made from someone who has taken the test. These are riddled with misspellings, grammar issues, are missing diagrams and exhibits, and in many cases are just outright wrong. While I've tried to include much of the data center topics, you'll also want to supplement your studies with the textbook (http://www.ciscopress.com), and I encourage you to take the official practice exams (www.PearsonTestPrep.com).

Please note, I am not affiliated with Cisco Systems (though I am a fan), and this guide is not endorsed by Cisco. The CCNA, CCENT, DCICT, and DCICN are all trademarks of Cisco and are only used here for educational purposes.

In general, as with all standardized tests, you'll want to read each question twice (especially the questions that ask what options are not correct or not true), and double check your answer before you go on to the next question. In some cases, I've grouped similar items together to make learning or remembering the concepts easier, in other cases I've randomized the order of questions. Please remember, there could be several feasible answer options, but look for and select the best overall answer. Throughout the book it will look like I've combined several questions into one, this is also done on purpose (that's the instructional designer in me) to help you compare and contrast topics and concepts. I found it helpful when using guides like this to take a sheet of paper, cover the answer (and the adjacent page) while reading the question, think about or write down your response in a notebook, and then remove the sheet to check your answers.

Hopefully, a lot of these questions will serve as a reminder of terms and technology that you already know but may not use every day. For the items that you aren't familiar with or didn't know the answers, please be sure to study up on those topic areas! I've tried to stay within the scope and context of the DCICN (my ultra-technical readers out there will notice when I've skipped some details that are not in the scope of the DCICN), except when some solutions required a little more detail to make more sense. I've also researched and reviewed the answer to each question; please let me know if you see any errors or typos (and your suggested corrections) and I'll correct them and issue a book revision.

I hope this study guide helps and good luck on your exam!

Miguel Ramlatchan

Miguel Ramlatchan
mramlatc@mediainventive.com

Let's get started:

1. In Cisco's Collapsed Core system design, which two levels of the traditional Cisco three-tier design have been combined?

    a. Aggregation and Access layers into the Collapsed Core

    b. Core and Distribution layers into the Collapsed Core

    c. Distribution and Access layers into the Collapsed Core

    d. Distribution and Aggregation layers into the Collapsed Core

    e. Expanded Core and Access layers into the Collapsed Core

    f. None of the Above

Answer: b

The Core and Distribution layers of Cisco's traditional three-tier hierarchy can be combined by the use of dense Layer 3 switches into the Collapsed Core layer in collapsed designs. You'll see in some data center applications the Distribution layer also referred to as the Aggregation layer.

Cisco’s traditional (and very successful) three-tiered model emphasizes a Core layer, an Aggregation/Distribution layer, and an Access layer that connects to users/clients:

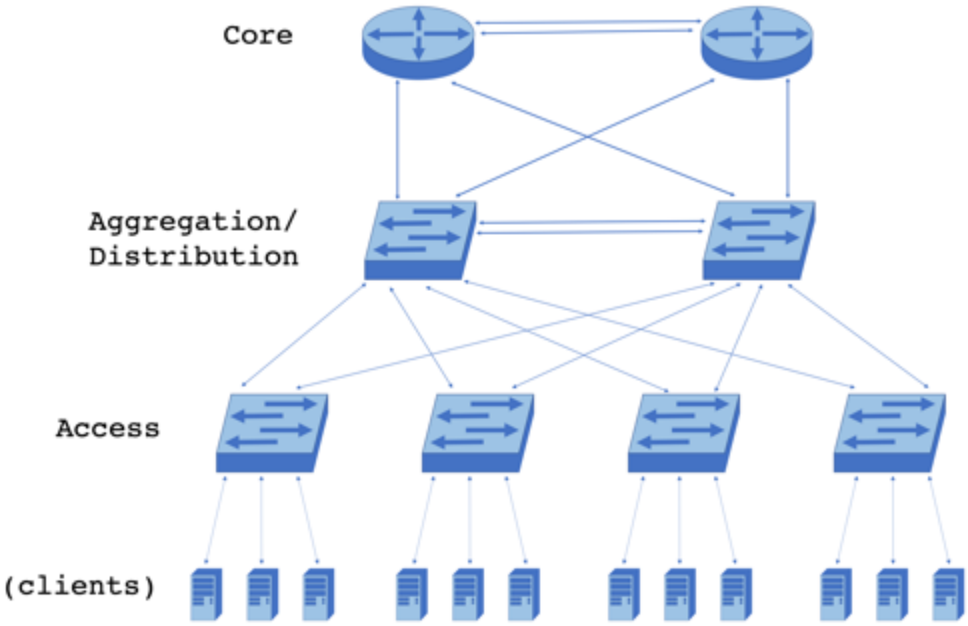

Cisco’s Nexus Layer 3 switches are dense enough to allow for design flexibility, especially in data centers, and thus the Core and Aggregation/Distribution layers can be combined. This is also known as a Spine and Leaf design:

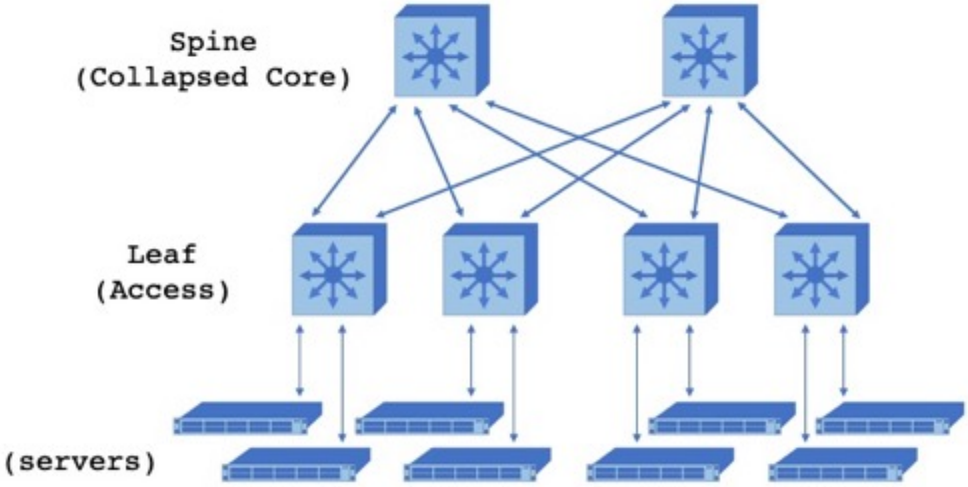

2. What does not happen at the Leaf layer?

a. Server and host connections to the network

b. Network connections to clients

c. The LAN’s connections to the Internet

d. VLANs are implemented on Leaf switches

e. Connections to Spine or Collapsed Core switches

f. Connections between Leaf switches

Answer: c, f

Please be sure to read the questions and answer options on the exam at least twice before answering, in this example the question is asking which options are wrong. In Cisco’s newer Spine (or Collapsed Core layer) and Leaf (or Access layer) topology, each Leaf switch connects to each Spine switch, and so there is no need to connect Leaf switches to each other directly. Also, the Layer 3 switches and routers in the Spine connect the LAN to the outside Internet.

3. What is the purpose of a Cisco Fabric Extender? (choose two)

a. To provide routing for IP packets

b. To provide switching for Ethernet frames

c. To scale out the interfaces of its parent switch

d. To scale out the interfaces of its parent router

e. To provide both Layer 2 switching and Layer 3 routing on the backplane of a UCS 5108 Chassis

f. To store the primary and backup NX-OS

g. To store the primary and backup UCS Manager

Answer: c, g

A Cisco Fabric Extender (FEX), such as a Nexus 2000 series FEX, does not do any internal switching. Instead, it provides additional interfaces for connections to host servers and port channel connections back to its parent switch, such as a Nexus 5000 or 7000 series switch, that does the switching. Production instances of the Cisco Unified Computing System (UCS) Manager also run on the Fabric Extenders, one instance serves as the primary, other instances serve as the secondary or backups for the primary in the case of a failure to the Fabric Extender that is running the primary UCS Manager.

A Cisco Fabric Extender is a cost-effective expansion to the capabilities of a parent Nexus switch by providing additional interfaces to hardware hosts and virtual servers without the need for additional, and rather expensive, Nexus data center switches:

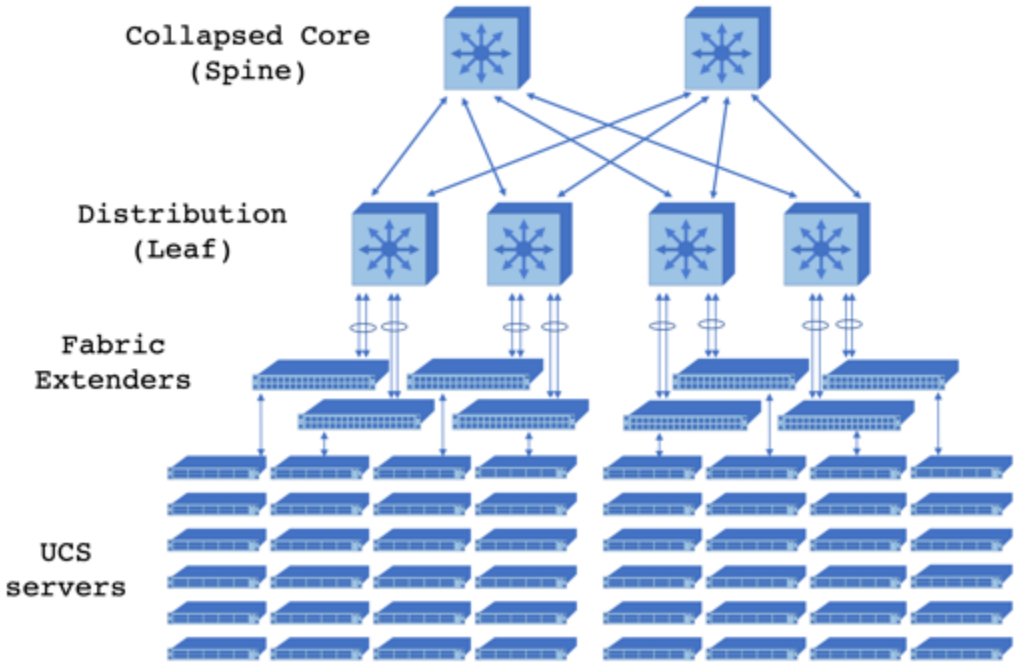

4. What is the primary purpose of the data center’s Collapsed Core or Spine?

   a. Firewall management

   b. SAN integration

   c. High-speed switching

   d. Optimized VLANs

   e. Redundant connections from hosts to the network

Answer: c

The primary functional goal of the data center’s Collapsed Core or Spine is high-speed switching. Connection to the SAN (Storage Area Network), VLANs (Virtual Local Area Networks), and the connections to host servers or workstations are primarily functions of the Access or Leaf layer.

5. When would you need to use a crossover cable? (choose three)

a. When connecting a Switch to a Router

b. When connecting a Host PC to a Router

c. When connecting a Hub to a Switch

d. When connecting a Router to another Router

e. When connecting a PC to another PC

Answer: b, d, e

Ok yes, most modern devices have an auto MDIX feature (Medium Dependent Interface Crossover) where a specific, physical crossover cable is not required, but Cisco wants us to know the fundamentals. And fundamentally a crossover cable reverses the pinouts in a standard Ethernet cable so that the NICs (Network Interface Cards) in devices have correct transmit to receive pairs. These devices will then get transmit signals on their receive pins and transmit signals to the other device's receive pins. In a standard 10/100 Mbps Ethernet crossover cable, Pin 1 connects to Pin 3, Pin 2 to Pin 6, Pin 3 to Pin 1, and Pin 6 to Pin 2 (see diagram).

Note, the term "host" is used throughout this book to describe a hardware workstation, laptop, server, or any physical client device connected to the network. In many cases, especially in the context of data centers, a physical "host" will be a hardware platform for virtual machines and servers.

Standard Straight-through Ethernet Cable:

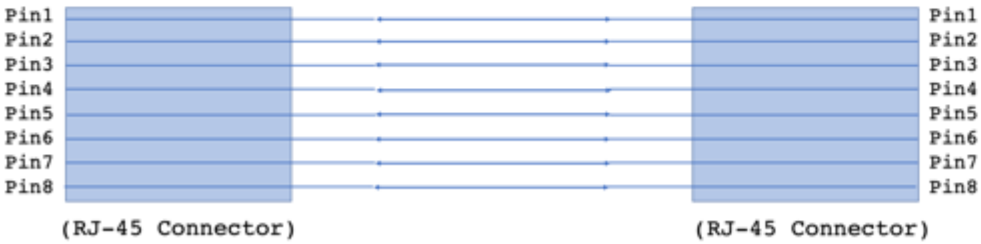

10/100 Mbps Ethernet Crossover Cable:

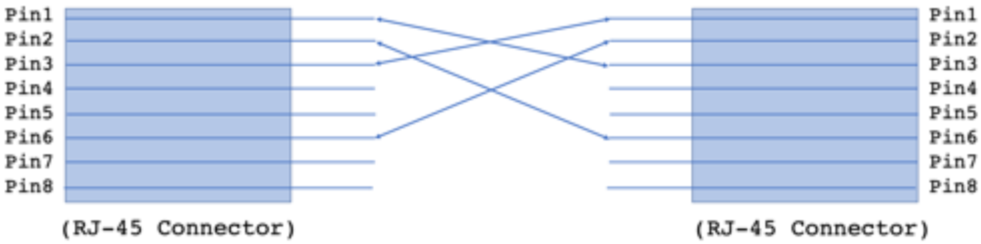

1000 Mbps Ethernet Crossover Cable:

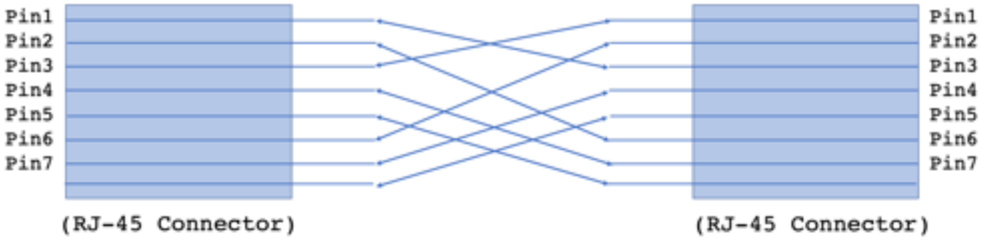

Rollover Cable (Router Console):

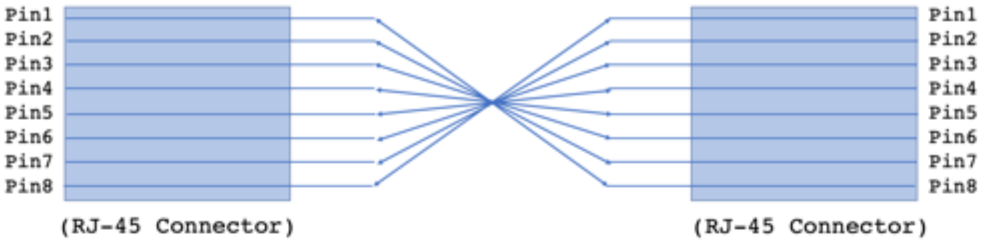

6. What is the IP address 192.168.100.1 when converted into binary?

   a. 11100000.10101000.10011011.00000001

   b. 11000001.01010111.01100100.00000001

   c. 11000011.01010111.01100100.00000001

   d. 11000000.01010111.01100100.10000001

   e. None of the above

Answer: e

192.168.100.1 into binary is 11000000.10101000.01100100.0000001. To answer this question we have to recall our decimal to binary conversion, which is still very important for the DCICN because the binary octets are the foundation for many other concepts in our data centers and networks. And so, a quick refresher:

Imagine a chart with two rows of eight columns, with each column representing powers of 2 from $2^0=1$ to $2^7=128$:

128	**64**	**32**	**16**	**8**	**4**	**2**	**1**

First, start with the 192 in “192.168.100.1” and subtract the largest power of 2 that you can, in this case 128 can be subtracted from 192, so put a 1 in the 128 column. Now subtract 192 – 128 = 64, then look for the largest power of 2 that can be subtracted from 64 which will be... 64, so 64 – 64 = 0 and so we’re done subtracting, and we put a 1 in the 64 column, and zeros in the other columns:

128	**64**	**32**	**16**	**8**	**4**	**2**	**1**
1	1	0	0	0	0	0	0

And now we have the binary equivalent of 192 or 11000000.

Next, we move on to the 168, we know we can subtract 128 so we do that and put a 1 in the 128 column. We have 168 – 128 = 40, we can't subtract 64 so we put a 0 in that column and move onto the next. The next column is 32, which we can subtract from 40, so we put a 1 in that column and we have 40 – 32 = 8. We can't subtract 16 from 8, so we skip that and put a zero in the 16 column, we can subtract 8 from 8, which means we put a 1 in the 8 column, and 8 – 8 = 0 so we're done.

128	**64**	**32**	**16**	**8**	**4**	**2**	**1**
1	0	1	0	1	0	0	0

The binary equivalent of 168 is 10101000.

Next, we have 100, once we subtract 64 we have 36, once we subtract 32 we have 4, and finally we subtract 4 from 4 to get 0 and our conversion is complete:

128	**64**	**32**	**16**	**8**	**4**	**2**	**1**
0	1	1	0	0	1	0	0

The binary equivalent of 100 is 01100100.

Finally, the last octet of the IP address to convert is the 1, we subtract that from the last column to get 1 - 1 = 0, and so our short conversion is complete:

128	**64**	**32**	**16**	**8**	**4**	**2**	**1**
0	0	0	0	0	0	0	1

The binary equivalent of 1 as an 8-bit octet is 00000001.

When we combine those four octets together we see that the binary equivalent of the IP address 192.168.100.1 is 11000000.10101000.01100100.00000001

The process to convert decimal subnet masks to binary is identical. For instance, 255.255.255.128 is 11111111.11111111.11111111.10000000.

7. Can you match the data encapsulation PDU (Protocol Data Unit) to the OSI layer? (some layers may be used more than once)

Segments

Packets

Frames

Datagrams

Layer 1

Layer 2

Layer 3

Layer 4

Layer 5

Layer 6

Layer 7

Answer: Segments = Layer 4
Packets = Layer 3
Frames = Layer 2
Datagrams = Layer 3

Note, the terms “packets” and “datagrams” are often used interchangeably to describe Layer 3 PDUs. However, historically datagrams referred specifically to connectionless best-effort communications. Also, at Layer 1 or the Physical layer, PDUs are converted into digital bits for transmission over the physical medium.

On the exam you will be able to drag the items on the left and drop them onto the items on the right, the closest analogy in a printed book is to draw the lines:

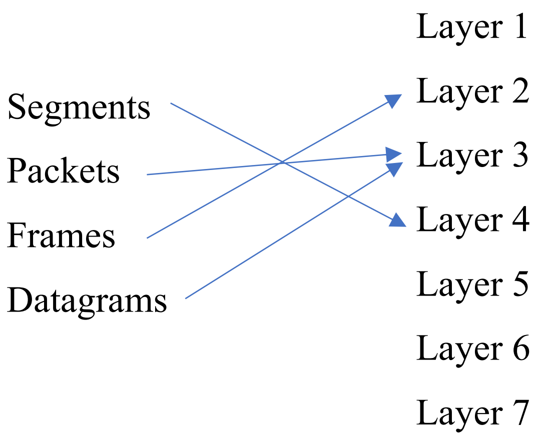

8. In the OSI model, at what layer is UDP and TCP utilized?

a. Application

b. Presentation

c. Session

d. Transport

e. Network

f. Data Link

g. Physical

h. None of the Above

Answer: d

UDP (User Datagram Protocol) and TCP (Transmission Control Protocol) are both part of the Internet Protocol (IP) suite and are used at the Transport layer of the OSI (Open Systems Interconnect) model to handle data transfer and communications between devices, such as between PCs and servers. UDP is commonly used for real-time communications and services like video and voice applications, while TCP is commonly used for applications that are not specifically time sensitive like email or file transferring.

The OSI model:

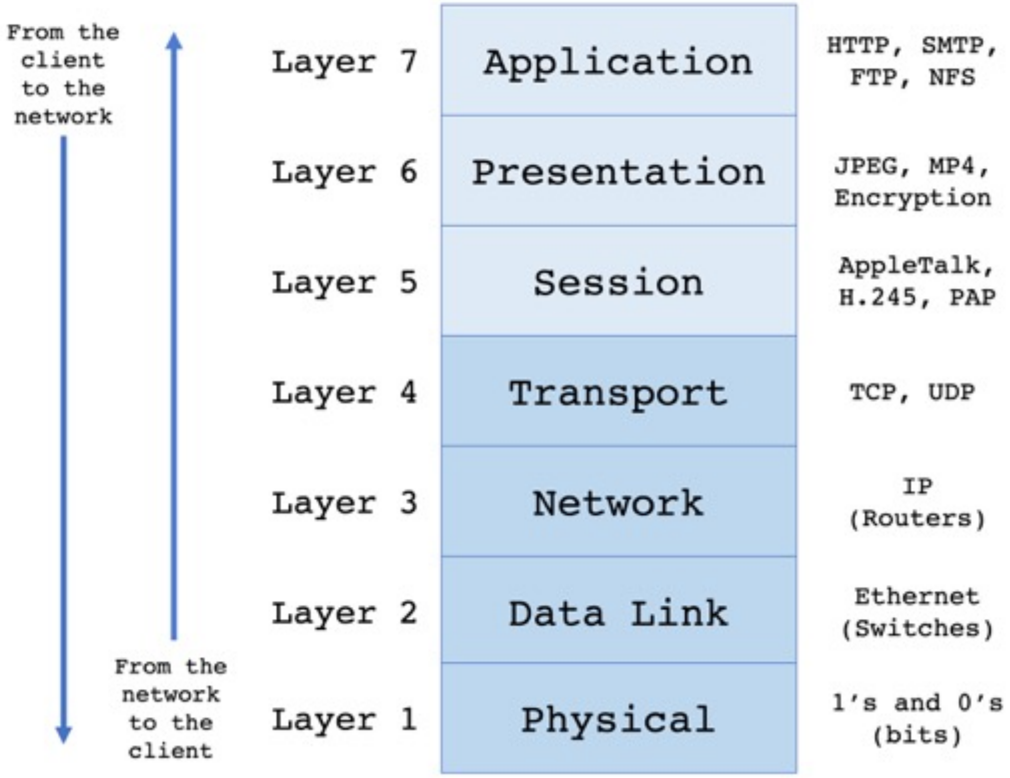
From the client to the network
From the network to the client

Layer 7	Application	HTTP, SMTP, FTP, NFS
Layer 6	Presentation	JPEG, MP4, Encryption
Layer 5	Session	AppleTalk, H.245, PAP
Layer 4	Transport	TCP, UDP
Layer 3	Network	IP (Routers)
Layer 2	Data Link	Ethernet (Switches)
Layer 1	Physical	1's and 0's (bits)

Information ‘flows’ from the network to the client from Layer 1 to Layer 7, then from the client to the network from Layer 7 to Layer 1. Examples of Application layer protocols include HTTP, SMTP, FTP, and NFS. Examples of Presentation layer protocols include JPEG, MP4, and data is encrypted and decrypted at this layer (when encryption is used). Examples of Session protocols include AppleTalk, H.245 for multimedia, and PAP for authentication. Transport layer protocols include TCP and UDP. Routing based on IP addresses occurs at Layer 3, and switching based on MAC hardware addresses happens at Layer 2. Frames are converted into 1’s and 0’s for transmission over copper wire, fiber optic cable, or wireless transmitters and receivers at the Physical layer.

9. Which of these options are layers in the TCP/IP model (select four):

a. Application
b. Presentation
c. Session
d. Transport
e. Network
f. Internet
g. Data Link
h. Link
i. Physical
j. Wireless

Answer: a, d, e, h

Okay, there are a lot of answer options here, but I wanted to highlight the differences between each layer of both models. The layers of the TCP/IP model are the Application, Transport, Internet, and the Link layer. The Application layer combines the OSI's Application, Presentation, and Session layers, while the Transport layers of each model are similar. The TCP/IP Internet layer is similar to the OSI's models Network layer, and the OSI model breaks the TCP/IP model's Link layer into the Data Link and Physical layer.

The TCP/IP model was developed by DARPA (the United States' Department of Defense's Defense Advanced Research Projects Agency) during the development of what would become the ARPANET (and what would evolve into the Internet). It is sometimes referred to as the DoD (Department of Defense) model. However, the OSI model is typically what we use to describe the flow of information from the network to the client. Here are how these layers match up in terms of their comparable functionality:

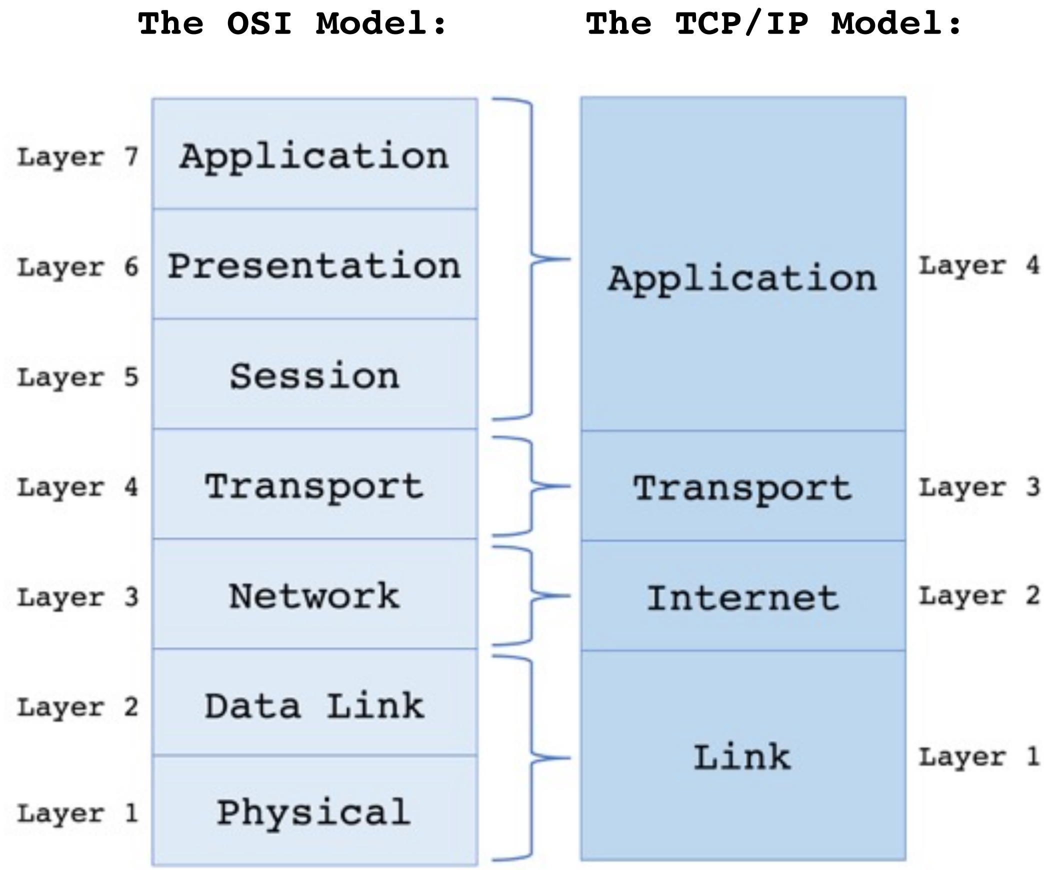

10. A company that specializes in video production has a number of staff that require 24x7 high-speed access to their large, uncompressed video files for their projects. They all work in the same building with other staff in the same company, and often share files and projects with each other. Other company staff in other departments have more generalized office workflow needs and do not require high-speed access, or any access, to these large video files. What technology can we implement to meet the needs of our video production staff? (choose the best answer)

a. a LAN

b. a WAN

c. a SAN

d. a WLAN

e. None of the above

Answer: c

We use Storage Area Networks (SANs) in our data centers when our clients require shared storage, they need more storage than their office workstations are capable of supporting, they need high performance and high throughput, they need to share large files, and we want to protect that data and ensure reliable access by implementing a fabric topology (interconnected switches with redundant paths). Other examples would be research institutions with big data needs, financial institutions with large databases, labs that are running complex simulations, or any example where some staff need high speed, real-time access to large, business-critical files and other staff do not.

11. Which of these options best describe the basic workings of a Layer 2 Switch? (choose the best answer)

a. learning, flooding, forwarding, filtering

b. listening, learning, forwarding, blocking

c. learning, optimizing, routing, dropping

d. none of the above

Answer: a

When a switch receives an Ethernet frame, it reads the MAC addresses and compares that or adds that to its lookup table. If it doesn’t already have the destination MAC address in its table, it will flood the frame out all of its active interfaces (except the interface that the frame was received on). Once the switch gets a response back from the intended destination, it will save that information to its lookup table (learning), and forward frames to only that destination on that interface. The switch will filter, or not send frames out the other interfaces. Listening, learning, forwarding, and blocking better describe the transitions and interface functions in Spanning Tree Protocol. Learning, optimizing (routes to other networks), routing, and dropping better describes Layer 3 routing.

12. In terms of Cisco UCS servers, what are some basic issues to look out for when troubleshooting servers?

a. Adapters not supported by software

b. Mismatched CPUs

c. BIOS and CPU mismatches

d. Using different/mismatching memory modules

e. Failure to boot due to RAID configurations

f. All of the above

Answer: f

As with many technologies, when troubleshooting it is best practice to start with the basics, or the issues that are most likely to occur. While you would think that engineers will double check the capabilities or specifications of BIOS updates, processors, NICs/adapters, or memory before they install them in a box, mistakes do happen (which is also why it is helpful to test upgrades in a pre-production or a development environment first if you can). Any of these issues can occur in Unified Computing System (UCS) servers.

13. What algorithm does this process describe?

Step 1: a network device has a frame to send, but first checks the network to confirm that no other devices are already sending frames,

Step 2: the device sends the frames and listens for any collisions,

Step 3: if a collision occurs (another device is also sending frames) the device sends a jamming signal to all devices in the network,

Step 4: all devices stop sending and each wait a random time before they try to send frames again

Step 5: the process restarts at Step 1…

a. CMSA/CD

b. CSMA/CD

c. CSMA/CA

d. Token passing

e. TDMA

f. None of the above

Answer: b

Carrier Sense Multiple Access with Collision Detection, or CSMA/CD, is the process that Ethernet networks used to use to communicate between devices. However, the process becomes more and more inefficient as the network grows larger, which is why bridges were developed, to separate Layer 2 collisions domains. Then along came multipoint bridges, which allowed for a single multipoint bridge to provide independent collision domains for connected devices. Then along came Ethernet switches, which are multiport bridges that implement Layer 2 processing using specialized hardware to process and switch frames in near real time. Switches separate collision domains such that each device connected to a switchport on the switch is on its own collision domain, eliminating collisions from the device to the switch.

14. What type of header is shown below?

Bit#	0	8	16	24
0	Preamble…			
32	...Preamble (continued)			Start of Frame Delimiter
64	48-bit Destination Address…			
96	...48-bit Destination Address (continued)		48-bit Source Address…	
128	...48-bit Source Address (continued)			
160	802.1Q tag (optional)			
192	Length		DSAP	SSAP
224	CTRL	(payload…)		
…	(...payload...)			
12192	(...payload)		Frame Check Sequence…	
12224	…Frame Check Sequence (cont)		(Start of Interframe gap...)	

a. An IP packet header

b. A UDP segment header

c. A TCP segment header

d. An Ethernet II frame header

e. An 802.3 frame header

Answer: e

This is a frame; the giveaway is the FCS (Frame Sequence Check) field at the end. As for the specific type of frame, IEEE 802.3 frames use a Length field as well as Destination Service Access Points (DSAP) and Source Service Access Points (SSAP) fields that describe the protocol being used by the frame at the Network layer. However, Ethernet II frames are now used more often in IP networks.

15. What type of header is shown below?

Bit#	0	8	16	24
0	Preamble…			
32	...Preamble (continued)			Start of Frame Delimiter
64	48-bit Destination Address…			
96	...48-bit Destination Address (cont.)		48-bit Source Address…	
128	...48-bit Source Address (continued)			
160	802.1Q tag (optional)			
192	Ethertype		(payload...)	
…	(...payload...)			
12192	(...payload)		Frame Check Sequence…	
12224	…Frame Check Sequence (cont.)		(Start of Interframe gap...)	

a. An IP packet header

b. A UDP segment header

c. A TCP segment header

d. An Ethernet II frame header

e. An 802.3 frame header

Answer: d

Ethernet II frames use the Ethertype field rather than the Length, DSAP, and SSAP fields found in IEEE 802.3 frames. Ethernet II frames are what we typically use today in IP networks. IP packets (and TCP/UDP segments) would be encapsulated in the payload section. Also, note the 48-bit Address fields refer to MAC Addresses and there is a Frame Check Sequence footer.

16. What type of header is shown below?

<table>
<tr><th>Bit#</th><th colspan="2">0</th><th>8</th><th colspan="2">16</th><th>24</th></tr>
<tr><td>0</td><td>Ver</td><td>Length</td><td>Type of Service</td><td colspan="3">Total Length</td></tr>
<tr><td>32</td><td colspan="3">Identification</td><td>Flags</td><td colspan="2">Fragment Offset</td></tr>
<tr><td>64</td><td colspan="2">Time to Live</td><td>Protocol</td><td colspan="3">Header Checksum</td></tr>
<tr><td>96</td><td colspan="6">32-bit Source Address</td></tr>
<tr><td>128</td><td colspan="6">32-bit Destination Address</td></tr>
<tr><td>160+</td><td colspan="5">Options</td><td>Padding</td></tr>
</table>

a. An IP packet header

b. A UDP segment header

c. A TCP segment header

d. An Ethernet II frame header

e. An 802.3 frame header

Answer: a

The giveaway that this is an IP packet among other things is the use of a 32-bit Source Address and a 32-bit Destination IP Address (IP addresses are 32 bits long, MAC addresses are 48 bits long). UDP and TCP segments would be encapsulated in the IP Data section that follows this header.

17. What type of header is shown below?

Bit #	0	8	16	24
0	Source Port		Destination Port	
32	Length		Header and Data Checksum	

a. An IP packet header

b. A UDP segment header

c. A TCP segment header

d. An Ethernet II frame header

e. An 802.3 frame header

Answer: b

The UDP header field is pretty simple as compared to a TCP header, this makes sense as UDP is a connectionless communication protocol while TCP looks to establish a reliable two-way communication connection between hosts. Note, the Checksum field is used for error detection, not error correction (if not used this field will be all zero's). Session, Presentation, and Application data would follow this Layer 4 header in the segment data encapsulation process.

18. What type of header is shown below?

<table>
<tr><th>Bit#</th><th colspan="2">0</th><th>8</th><th>16</th><th>24</th></tr>
<tr><td>0</td><td colspan="3">Source Port</td><td colspan="2">Destination Port</td></tr>
<tr><td>32</td><td colspan="5">Sequence Number</td></tr>
<tr><td>64</td><td colspan="5">Acknowledgment Number</td></tr>
<tr><td>96</td><td>Data Offset</td><td>Res</td><td>Flags</td><td colspan="2">Window Size</td></tr>
<tr><td>128</td><td colspan="3">Header and Data Checksum</td><td colspan="2">Urgent Pointer</td></tr>
<tr><td>160+</td><td colspan="5">Options</td></tr>
</table>

a. An IP packet header

b. A UDP segment header

c. A TCP segment header

d. An Ethernet II frame header

e. An 802.3 frame header

Answer: c

There are ten required fields in the connection-oriented TCP segment header used to establish communications and data transfers. Note, there are nine Control Bits in the Flags field that are used to help establish and confirm a reliable connection between hosts (TCP Flags bits include SYN, ACK, FIN, URG, PSH, RST, ECE, CWR, and NS).

19. Match the cable type to the implementation specification:

Passive Twinax	100 m (10GBASE-T)
Active Twinax	55 m (10GBASE-T)
Category 5 UTP	10 m (10GBASE-T)
Category 5e UTP	550 m (1000BASE-LX)
Category 6 UTP	100 m (100BASE-T)
Category 6a UTP	5 m (10GBASE-T)
Category 7a STP	100 m (1000BASE-T)
Multimode Fiber	10 km (1000BASE-LX10)
Single Mode Fiber	50 m (40GBASE-T)

Answer:

Twinax, or twinaxial cabling, is a cost-effective variation of coaxial cable (what we used to use for 10BASE-5 and 10BASE-2), except it has two conductor cores and can be used for short 10 Gbps connections in data centers. Category 5, 5e, 6, 6a (UTP = Unshielded Twisted Pair), and 7a (STP = Shielded Twisted Pair) are all variations of eight conductor twisted pair cable, and of course, fiber optic cables can send signals the farthest. Note, these are Cisco's maximum recommended implementation specifications:

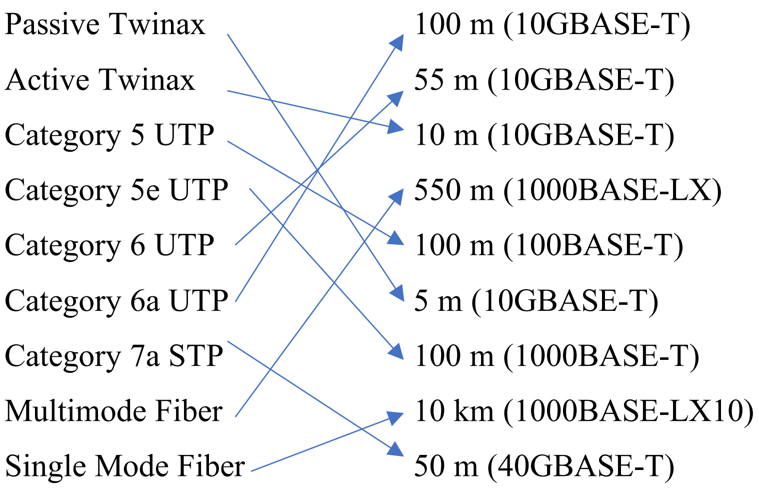

And if you don't want to follow all those arrows:

Passive Twinax	=	5 m (10GBASE-T)
Active Twinax	=	10 m (10GBASE-T)
Category 5 UTP	=	100 m (100BASE-T)
Category 5e UTP	=	100 m (1000BASE-T)
Category 6 UTP	=	55 m (10GBASE-T)
Category 6a UTP	=	100 m (10GBASE-T)
Category 7a STP	=	50 m (40GBASE-T)
Multimode Fiber	=	550 m (1000BASE-LX)
Single Mode Fiber	=	10 km (1000BASE-LX10)

20. In a network with redundant paths between switches, what are some of the problems that Spanning Tree Protocol helps prevent? (choose the best three answers)

a. Broadcast storms

b. Multicast storms

c. Multiple frame transmissions

d. Multiple packet transmissions

e. MAC database instability

f. IP address table instability

Answer: a, c, e

Spanning Tree Protocol helps prevent broadcast storms, multiple frame transmissions, and MAC database instability. Broadcast storms occur when broadcasted frames are received and retransmitted to other switches continuously. Multiple copies of the same frame, or multiple frame transmissions, will unnecessarily use switch RAM and CPU resources. Receiving the same frame from different ports on the same switch will confuse the switch and make its MAC database instable/unstable as it tries to update itself continuously. Option "c" is a better answer than "d", and "e" is a better answer than "f", because technically we're talking about frames and MAC addresses when we talk about switching.

21. These switches have Spanning Tree Protocol enabled, which one will be elected the root bridge?

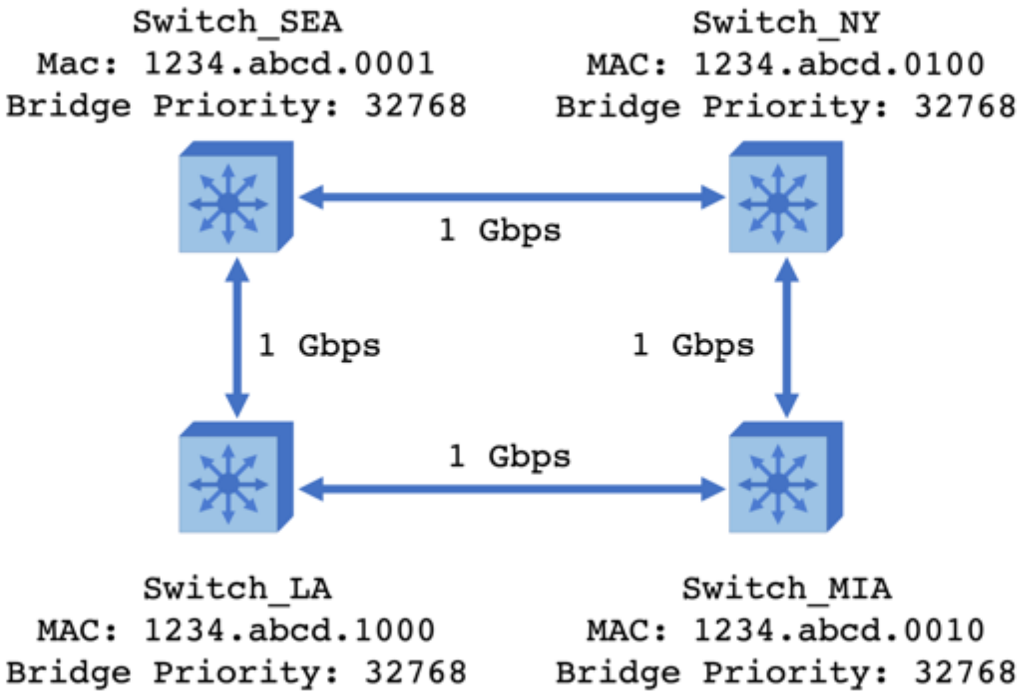

a. Switch_SEA

b. Switch_NY

c. Switch_LA

d. Switch_MIA

e. None of the above

Answer: a (Switch_SEA)

All the links between the switches are the same speed and so will all have the same Spanning Tree Protocol (STP) cost (1 Gbps = a cost of 4 according to 802.1D-1998). All bridge priorities are also set to the Cisco default (32768). So, the switches will look at the MAC address section of their STP bridge IDs, and the switch with the lowest address will become the root bridge. In this case, the BID (Bridge ID) of Switch_SEA is Bridge Priority & MAC or 32768.1234.abcd.0001.

In this scenario, Switch_SEA will send frames to Switch_NY and to Switch_LA, Switch_NY will send frames to Switch_SEA and Switch_MIA, and Switch_LA will send frames to Switch_SEA and Switch_MIA. Switch_MIA will send frames to Switch_NY, but it will block traffic from Switch_LA. This will prevent broadcast storms, multiple frame transmissions, and MAC database instability in our network:

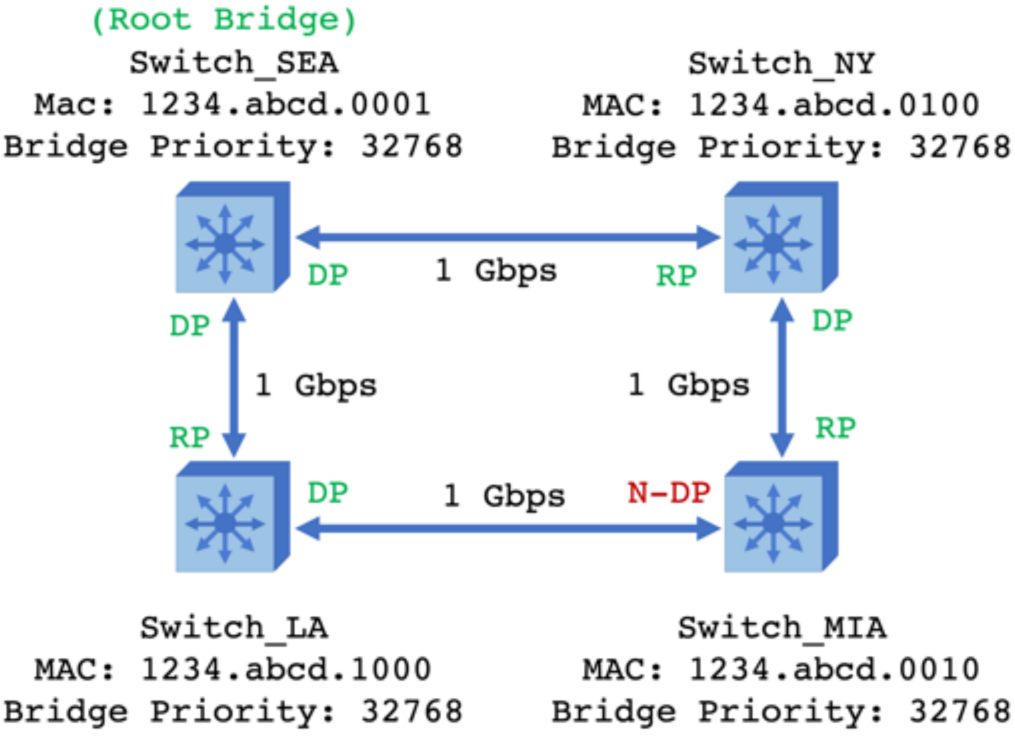

DP = Designated Port (forwards data frames and BPDUs)
RP = Root Port (forwards data frames and BPDUs)
N-DP = Non-Designated Port (blocks and does not forward frames but will forward BPDUs)

Note, when there is a STP cost tie, like when the links from non-root bridges are the same speed and there are an equal number of links between the non-root bridge and the root bridge, then the tie breaker is the lowest BID of the connected switches.

For instance, a frame from Switch_MIA to Switch_SEA could take either:

1. Switch_LA with a BID = 32768.1234.abcd.1000, and with a total cost of 2 x 1 Gbps (4) = 8, or it could take

2. Switch_NY with a BID = 32768.1234.abcd.0100, also with a total cost of 2 x 1 Gbps (4) = 8.

The BID of Switch_NY is lower that the BID of Switch_LA, and so STP will choose the path from Switch_MIA to Switch_NY to Switch_SEA. The Switch_MIA connection from Switch_LA will become a Non-Designated Port and block frames (except for BPDUs) from Switch_LA, thus preventing loops.

22. These switches have Spanning Tree Protocol enabled, which one will be elected the root bridge?

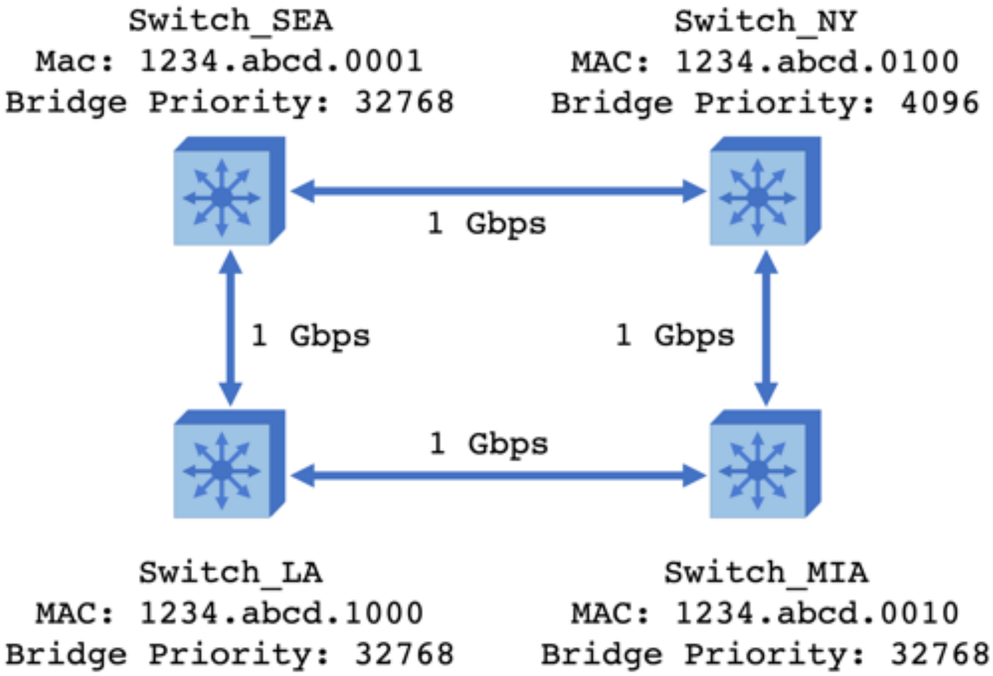

a. Switch_SEA

b. Switch_NY

c. Switch_MIA

d. Switch_LA

e. None of the above

Answer: b (Switch_NY)

The MAC addresses of the switches are the same as in the previous example, though the Bridge Priority of Switch_NY has been changed to 4096, the other switches still have the Cisco default bridge priorities. The BID (Bridge ID) of Switch_NY is Bridge Priority & MAC or 4096.1234.abcd.0100; it now has the lowest BID and will become the root bridge.

In this example, a frame from Switch_LA could take either Switch_SEA (BID = 32768.1234.abcd.0001, with a total cost of 2 x 1 Gbps (4) = 8), or it could take Switch_MIA (BID = 32768.1234.abcd.0010, also with a total cost of 2 x 1 Gbps (4) = 8), on its way to Switch_NY. The BID of Switch_SEA is lower that the BID of Switch_MIA, and so STP will choose the path from Switch_LA to Switch_SEA to Switch_NY.

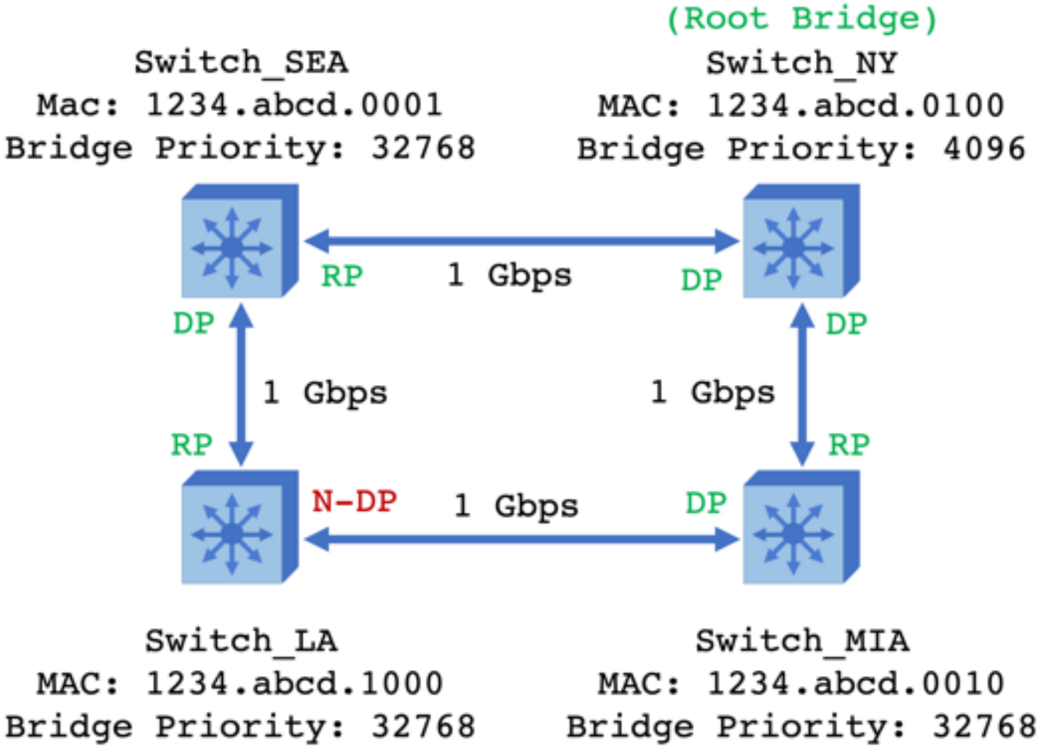

DP = Designated Port (forwards data frames and BPDUs)
RP = Root Port (forwards data frames and BPDUs)
N-DP = Non-Designated Port (blocks and does not forward frames but will forward BPDUs)

23. These switches have Spanning Tree Protocol enabled, which one will be elected the root bridge?

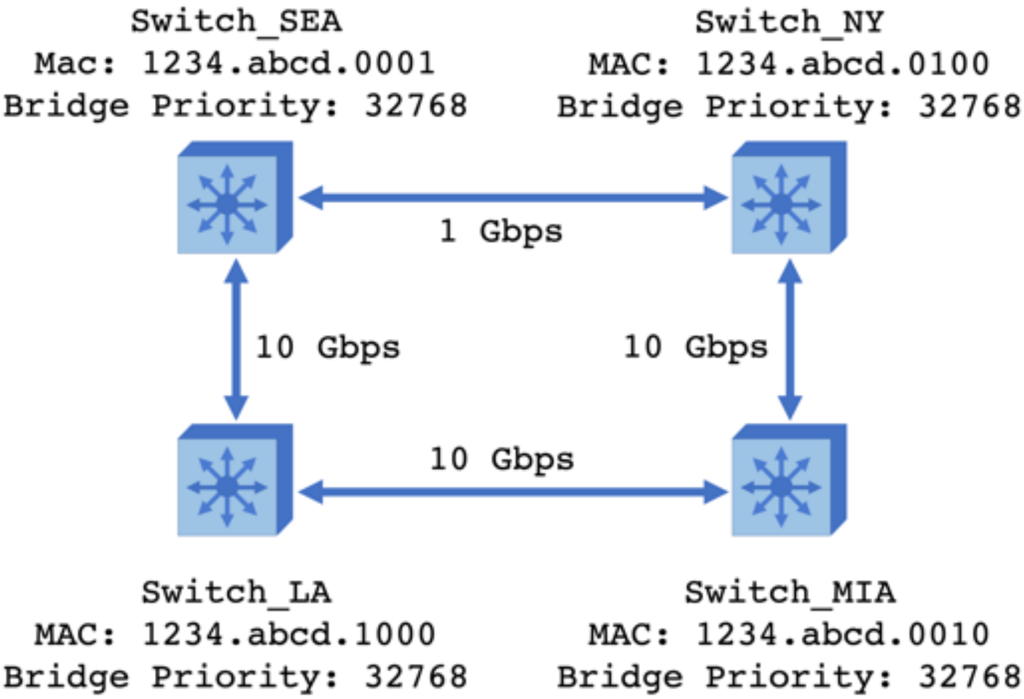

a. Switch_SEA

b. Switch_NY

c. Switch_MIA

d. Switch_LA

e. None of the above

Answer: c (Switch_MIA)

The MAC addresses of the switches are the same as in the previous examples, and the Bridge Priority of the switches are now back to the Cisco defaults. In this new scenario, the link between Switch_SEA and Switch_LA, between Switch_LA and Switch_MIA, and between Switch_MIA and Switch_NY are now all 10 Gbps, or an STP cost of 2 each. The link between Switch_SEA and Switch_NY is 1 Gbps or an STP total cost of 4. Switch_MIA will become the root bridge based in part on these new costs (the cost from Switch_SEA to Switch_NY to Switch_MIA = 4 + 2 = 6, while the cost from Switch_SEA to Switch_LA to Switch_MIA = 2 + 2 = 4). Either Switch_LA or Switch_MIA switch could serve as the root bridge, but there can only be one, so the tie breaker is BID and Switch_MIA has a lower BID than Switch_LA so Switch_MIA becomes the root bridge:

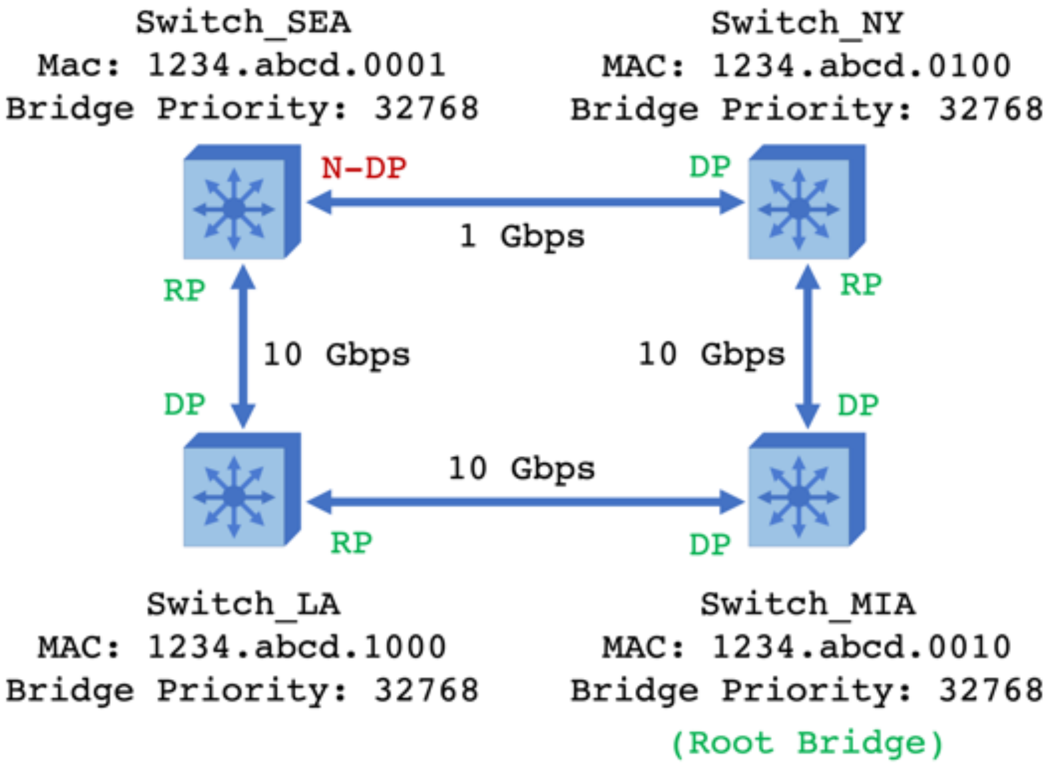

24. In this example, refer to the diagram below and use the Command Line Interface (CLI) on each switch to determine which switch is the Rapid Spanning Tree Protocol root bridge on VLAN 1:

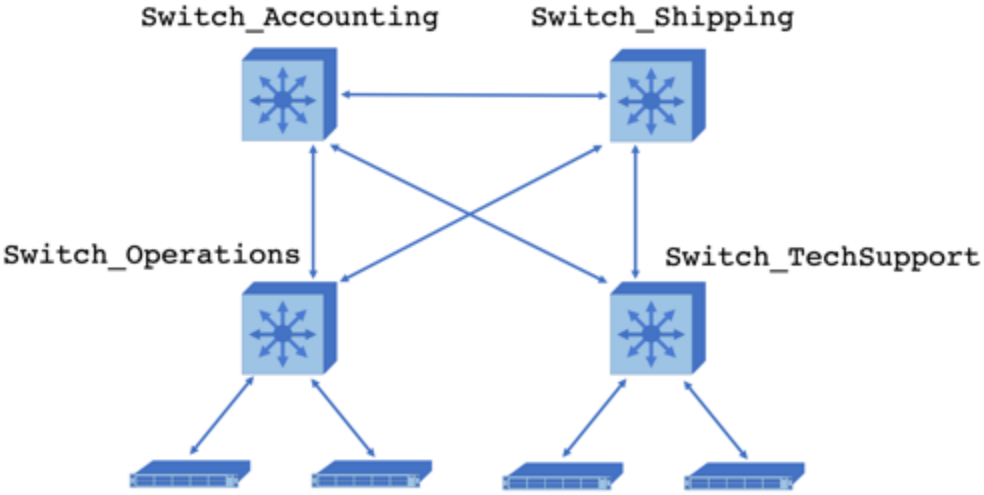

a. Switch_Accounting

b. Switch_Shipping

c. Switch_TechSupport

d. Switch_Operations

One way to find the answer would be to log into each switch and run a “Show Spanning-tree <vlan#>” command where “vlan#” is VLAN 1 in this example:

Starting with Switch_Accounting:

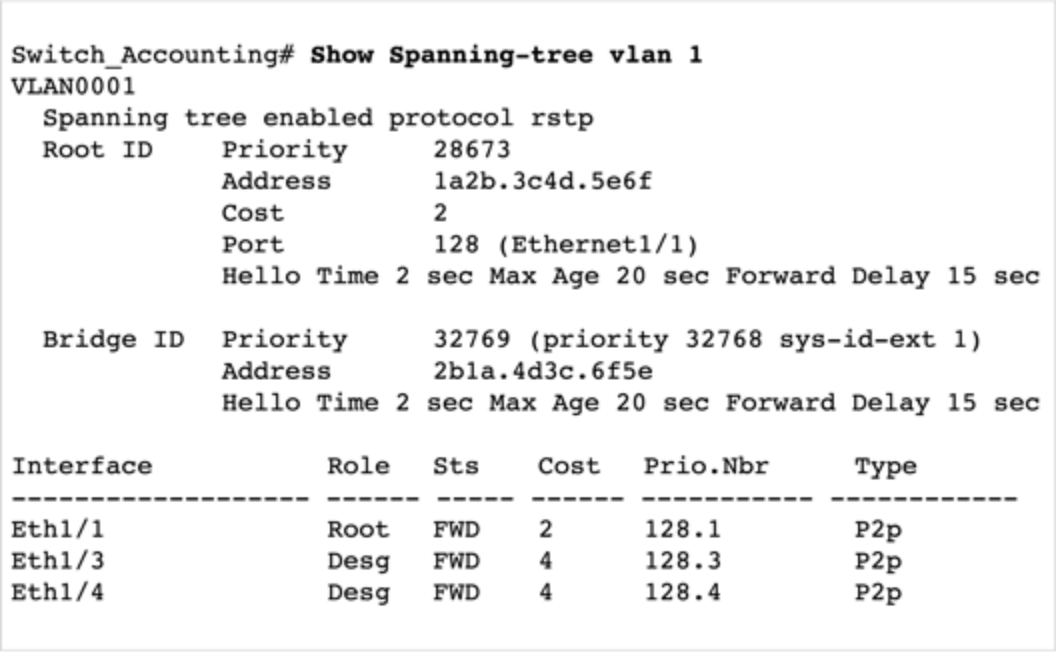

```
Switch_Accounting# Show Spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol rstp
  Root ID     Priority      28673
              Address       1a2b.3c4d.5e6f
              Cost          2
              Port          128 (Ethernet1/1)
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID   Priority      32769 (priority 32768 sys-id-ext 1)
              Address       2b1a.4d3c.6f5e
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface           Role   Sts    Cost   Prio.Nbr      Type
------------------- ------ ----- ------ ----------- ------------
Eth1/1              Root   FWD    2      128.1         P2p
Eth1/3              Desg   FWD    4      128.3         P2p
Eth1/4              Desg   FWD    4      128.4         P2p
```

Then moving to Switch_Shipping:

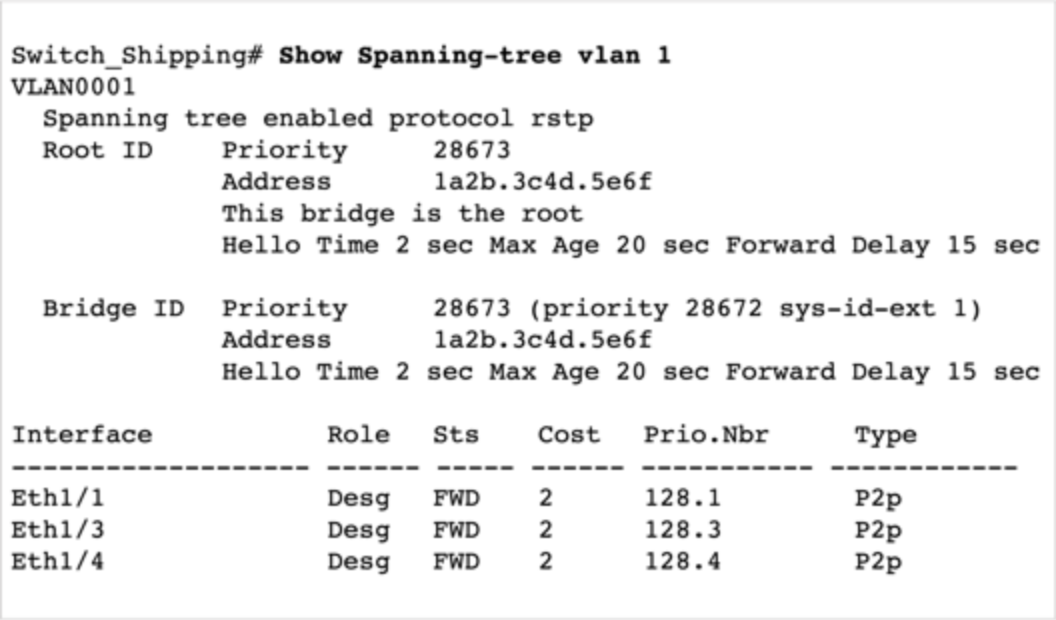

```
Switch_Shipping# Show Spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol rstp
  Root ID     Priority      28673
              Address       1a2b.3c4d.5e6f
              This bridge is the root
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID   Priority      28673 (priority 28672 sys-id-ext 1)
              Address       1a2b.3c4d.5e6f
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface           Role   Sts    Cost   Prio.Nbr      Type
------------------- ------ ----- ------ ----------- ------------
Eth1/1              Desg   FWD    2      128.1         P2p
Eth1/3              Desg   FWD    2      128.3         P2p
Eth1/4              Desg   FWD    2      128.4         P2p
```

We can see that Switch_Shipping is the root bridge, but to complete the example here is Switch_TechSupport:

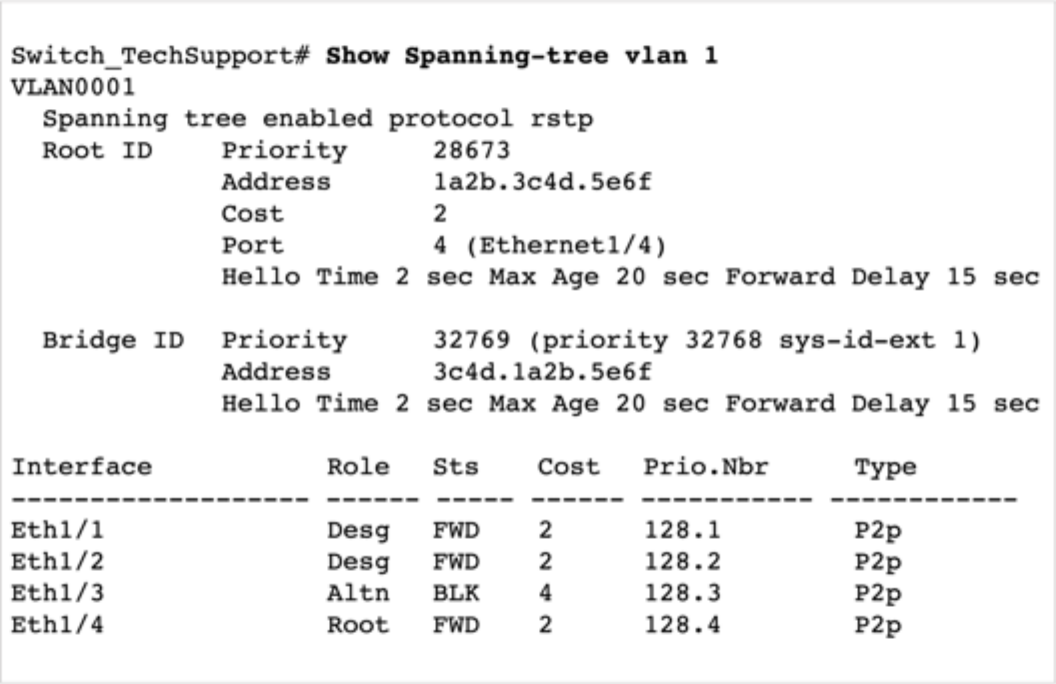

```
Switch_TechSupport# Show Spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol rstp
  Root ID     Priority      28673
              Address       1a2b.3c4d.5e6f
              Cost          2
              Port          4 (Ethernet1/4)
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID   Priority      32769 (priority 32768 sys-id-ext 1)
              Address       3c4d.1a2b.5e6f
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface           Role   Sts    Cost   Prio.Nbr      Type
------------------- ------ ----- ------ ----------- ------------
Eth1/1              Desg   FWD    2      128.1         P2p
Eth1/2              Desg   FWD    2      128.2         P2p
Eth1/3              Altn   BLK    4      128.3         P2p
Eth1/4              Root   FWD    2      128.4         P2p
```

Finally, moving to Switch_Operations:

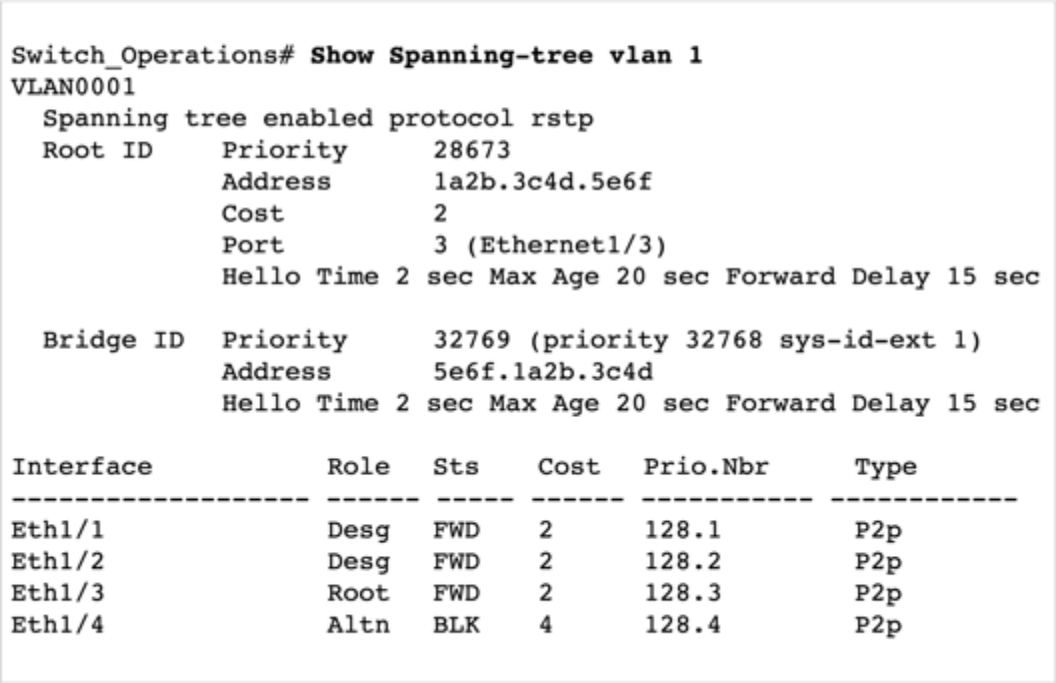

```
Switch_Operations# Show Spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol rstp
  Root ID     Priority      28673
              Address       1a2b.3c4d.5e6f
              Cost          2
              Port          3 (Ethernet1/3)
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID   Priority      32769 (priority 32768 sys-id-ext 1)
              Address       5e6f.1a2b.3c4d
              Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface           Role   Sts    Cost   Prio.Nbr      Type
------------------- ------ ----- ------ ----------- ------------
Eth1/1              Desg   FWD    2      128.1         P2p
Eth1/2              Desg   FWD    2      128.2         P2p
Eth1/3              Root   FWD    2      128.3         P2p
Eth1/4              Altn   BLK    4      128.4         P2p
```

So, which switch is the Rapid Spanning Tree Protocol root bridge on VLAN 1?

Answer: b (Switch_Shipping)

And so, looking at all four switches we can map out the Spanning Tree for our example network:

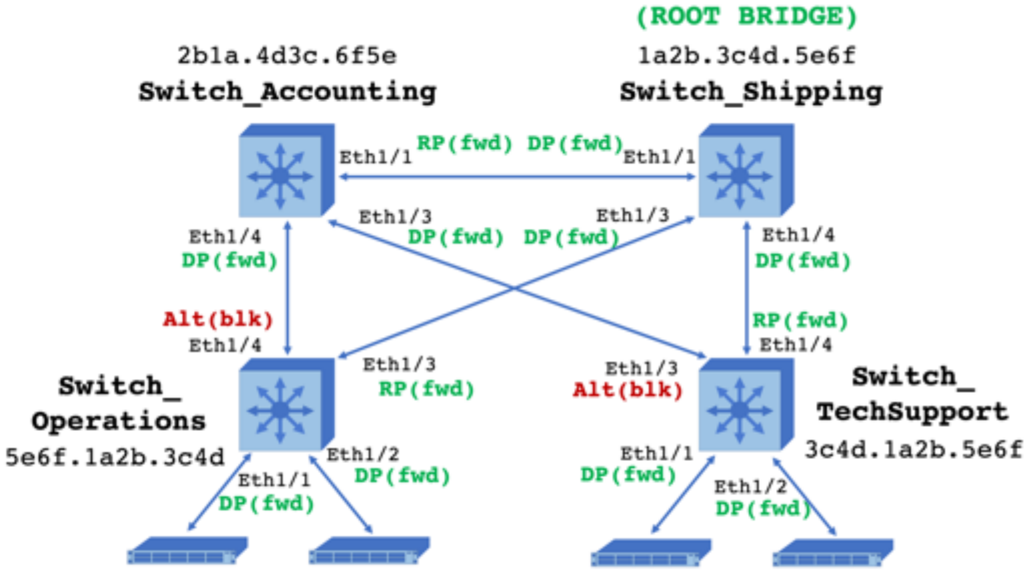

Switch_Shipping is the root bridge and its Ethernet 1/1, Ethernet 1/3, and Ethernet 1/4 interfaces are designated ports in RSTP and will be forwarding frames.

Switch_Accounting has a lower bridge ID than Switch_TechSupport or Switch_Operations, and so its Ethernet 1/4 and Ethernet 1/3 interfaces are also designated ports and its Ethernet 1/1 interface faces the root bridge and so it is a root port (the root port is the interface with the lowest cost to get to the root bridge, each non-root switch will only have one root port).

Switch_TechSupport’s Ethernet 1/4 port faces the root bridge and so it is a root port, its Ethernet 1/1 and Ethernet 1/2 ports are designated, and its Ethernet 1/3 port is an alternate port that is currently blocking frames and preventing loops.

Switch_Operations’s Ethernet 1/3 is its root port, its Ethernet 1/1 and Ethernet 1/2 interfaces are designated ports, and its Ethernet 1/4 interface is also currently blocking frames to prevent loops and broadcast storms. While these ports are blocking, they are still listening for BPDUs (Bridge Protocol Data Units) to know of any RSTP updates. In general, RSTP can recover from outages quicker than STP.

Note, these switches in this example are running RSTP, STP would show up as “Spanning tree enabled protocol ieee” in the “show spanning-tree vlan #” result.

25. Which of these options is not an STP port role? (select three)

a. Root

b. Trunking

c. Designated

d. Nondesignated

e. Alternate

f. Spanning

Answer: b, e, f

Please note that the question asks what is NOT a Spanning Tree Protocol port role. An alternate port is a RSTP port designation (RSTP alternate ports identify an alternative path to the root bridge through another switch). Note, listening and learning are also not port roles but are transitionary states between port roles.

26. Referring back to the previous example, which command will make Switch_LA the root bridge?

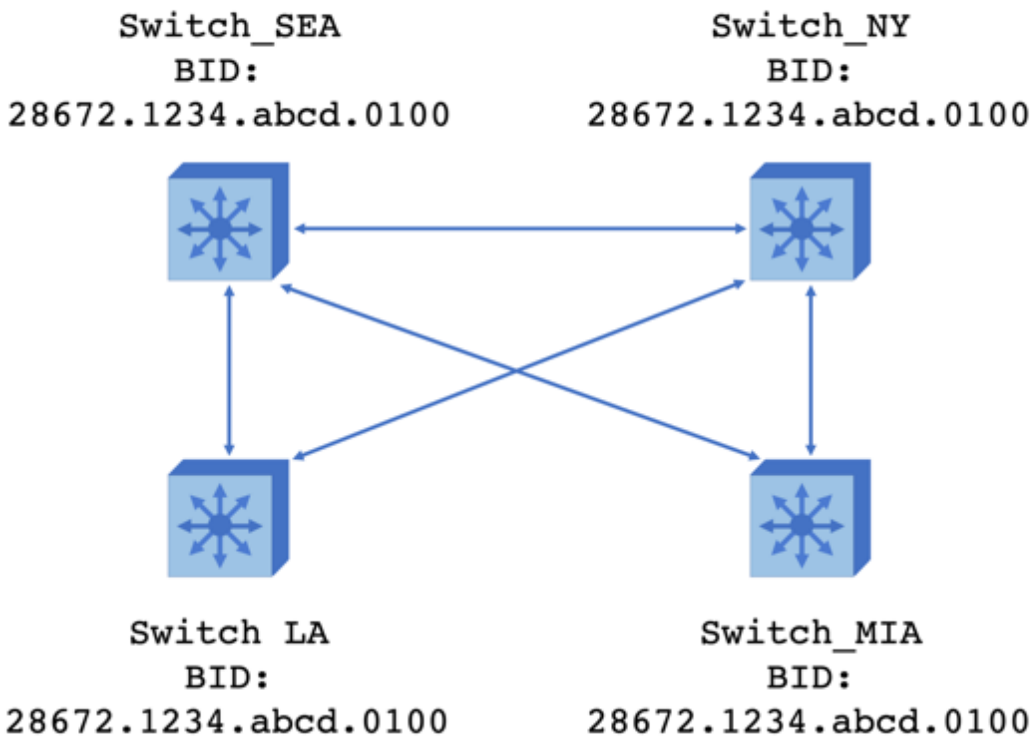

a.
```
Switch_LA# configure terminal
Switch_LA(config)# spanning-tree vlan 1-4094
priority 24576
```

b.
```
Switch_LA# configure terminal
Switch_LA(config)# spanning-tree all root priority
24576
```

c.
```
Switch_LA# configure terminal
Switch_LA(config)# rapid spanning-tree all root
primary
```

d.
```
Switch_LA# configure terminal
Switch_LA(config)# spanning-tree all primary
```

Answer: a

This command will change the STP bridge priority from 28672 to 24576 (for VLAN 1 through 4094 in this example). Once the switches share BPDUs (Bridge Protocol Data Units), they will see that Switch_LA now has a new BID of 24576.1234.abcd.1000, and they will elect Switch_LA the new root bridge.

Valid values for bridge priorities are multiples of 4096 starting at 4096 or:

1. 4096
2. 8192
3. 16484
4. 20480
5. 24576
6. 28672
7. 32768
8. 35864
9. 40960
10. 45056
11. 49152
12. 53248
13. 57344, and
14. 61440

27. How many bits are in a Bridge ID?

a. 48 bits

b. 64 bits

c. 32 bits

d. 128 bits

e. 4 bits

f. 8 bits

Answer: b

The 4 bit Priority + 12 bit VLAN ID + 48 bit MAC address = 64 bits. FYI, as for the other answer options: there are 32 bits in an IPv4 address, 48 bits in a MAC address, and 128 bits in an IPv6 address (and there are 4 bits in a nibble and 8 bits in a byte).

28. Which new emerging technology may help further increase bandwidth in our existing data centers?

a. Quantum electron fabrics (QuElf)

b. Bidirectional fiber (BiDi)

c. Tridirectional fiber (TriFi)

d. Extended range twinax (ETax)

Answer: b

Essentially two lasers at different wavelengths from different directions are used in the same physical fiber optic cable, such as a blue laser pointing upstream while a red laser points downstream. This type of QSFP+ (Quad channel Small Form Factor+) module will currently allow us to move data at 40 Gbps speeds over the same MMF (Multi-Mode Fiber) cables that we're using now for 10 Gbps connections. Tridirectional lasers would be interesting, but the answer for now is "b."

29. How many zone sets can a VSAN have?

a. 0 (none)

b. 1

c. 2

d. As many as the data center/SAN administrators want

e. As many as the switch's CPU and RAM can support

Answer: b

Each VSAN should have only one active zone set, though within each zone set there can be many zones, and devices on that VSAN can belong to more than one zone. Zone sets and zones add another layer of security and data protection to our SAN environment

30. Staff using the Staff Server below are allowed to access the budget data on the Finance Server; however, the student interns on the Interns Server are not allowed to access the budget data, which ACL will be effective?

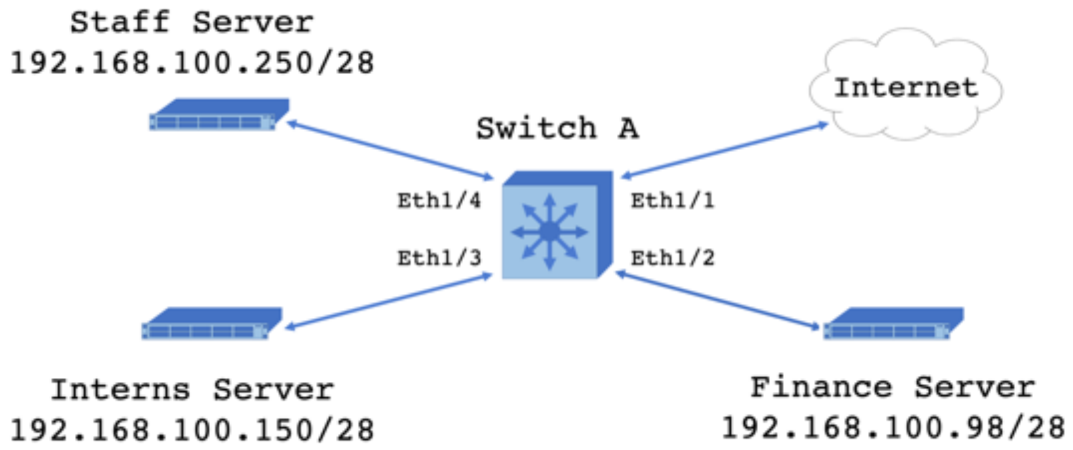

a.
```
a. Switch_A(config)# ip access-list 100 deny ip
   192.168.100.128 0.0.0.8 192.168.100.64 0.0.0.8
   Switch_A(config-acl)# config Ethernet 1/3
   Switch_A(config-if)# ip access-list 101 out
```

b.
```
b. Switch_A(config)# ip access-list 101 deny ip
   192.168.100.0 0.0.0.255 192.168.100.0 0.0.0.255
   Switch_A(config-acl)# permit ip any any
   Switch_A(config-acl)# config Ethernet 1/3
   Switch_A(config-if)# ip access-list 101 in
```

c.
```
c. Switch_A(config)# ip access-list 101 deny ip
   192.168.100.150 0.0.0.32 192.168.100.98 0.0.0.32
   Switch_A(config)# config Ethernet 1/2
   Switch_A(config-if)# ip access-group 101 out
```

d.
```
d. Switch_A(config)# ip access-list 101 deny ip
   192.168.100.144 0.0.0.15 192.168.100.96 0.0.0.15
   Switch_A(config-acl)# permit ip any any
   Switch_A(config-acl)# config Ethernet 1/3
   Switch_A(config-if)# ip access-group 101 in
```

Answer: d

Okay, I know, subnetting is more of a CCENT or a CCNA Routing and Switching exam topic, but it is still important to at least be aware of subnetting in a data center, especially in the context of ACLs. So, the first step is to take that CIDR (Classless Inter-Domain Routing) 28 and confirm what network these hosts are on. The Interns Sever with an IP address of 192.168.100.150 with a subnet mask of 255.255.255.240 belongs to subnetwork 192.168.100.144, and the Finance Server with an IP address of 192.168.100.98 with a subnet mask of 255.255.255.240 belongs to subnetwork 192.168.100.96.

That information alone eliminates three answer options, but to continue with the example, in the CLI we next create an extended ACL called 101 and the first line in it denies hosts from the 192.168.100.144 subnet mask 255.255.255.240 subnetwork from communicating via IP with any host on the 192.168.100.96 subnet mask 255.255.255.240 subnetwork. Note, the use of wildcards in 0.0.0.15. Another way to do this could have been: “ip access-list 101 deny ip 192.168.100.144/28 192.168.100.96/28”, but I wanted to use wildcards in this example. We use “permit ip any any” to allow any other network to access any other network after the ACL has filtered for 192.168.100.144 hosts trying to access 192.168.100.96 hosts. Then we apply it to the input side of the Ethernet 1/3 interface (recall that in general ACLs on interface inputs are more efficient than using ACLs on interface outputs).

Is your subnetting a little rusty like mine? If so, I like jotting down this chart (or at least part of it) on the whiteboard they give you when you are taking the exam, then I can reference it as needed:

Subnetting Reference Chart:

/CIDR	Netmask	# of Subnets			# of Host			Block Size
		Class A	Class B	Class C	Class A	Class B	Class C	
/8	255.0.0.0	1			16777214			256
/9	255.128.0.0	2			8388606			128
/10	255.192.0.0	4			4194302			64
/11	255.224.0.0	8			2097150			32
/12	255.240.0.0	16			1048574			16
/13	255.248.0.0	32			524286			8
/14	255.252.0.0	64			262142			4
/15	255.254.0.0	128			131070			2
/16	255.255.0.0	256	1		65534	65534		256
/17	255.255.128.0	512	2		32766	32766		128
/18	255.255.192.0	1024	4		16382	16382		64
/19	255.255.224.0	2048	8		8190	8190		32
/20	255.255.240.0	4096	16		4094	4094		16
/21	255.255.248.0	8192	32		2046	2046		8
/22	255.255.252.0	16384	64		1022	1022		4
/23	255.255.254.0	32768	128		510	510		2
/24	255.255.255.0	65536	256	1	254	254	254	256
/25	255.255.255.128	131072	512	2	126	126	126	128
/26	255.255.255.192	262144	1024	4	62	62	62	64
/27	255.255.255.224	524288	2048	8	30	30	30	32
/28	255.255.255.240	1048576	4096	16	14	14	14	16
/29	255.255.255.248	2097152	8192	32	6	6	6	8
/30	255.255.255.252	4194304	16384	64	2	2	2	4

(Another good subnetting reference is Hintz, Obediente, & Karakok's (2017) *Cisco Data Center DCICN 200-150*, Chapter 14, and Todd Lammle's (2016) *CCENT ICND1 Study Guide 3rd Ed.,* Chapter 4.)

31. True or false? In this diagram and ACL example, can the interns access the Internet form the Interns Server?

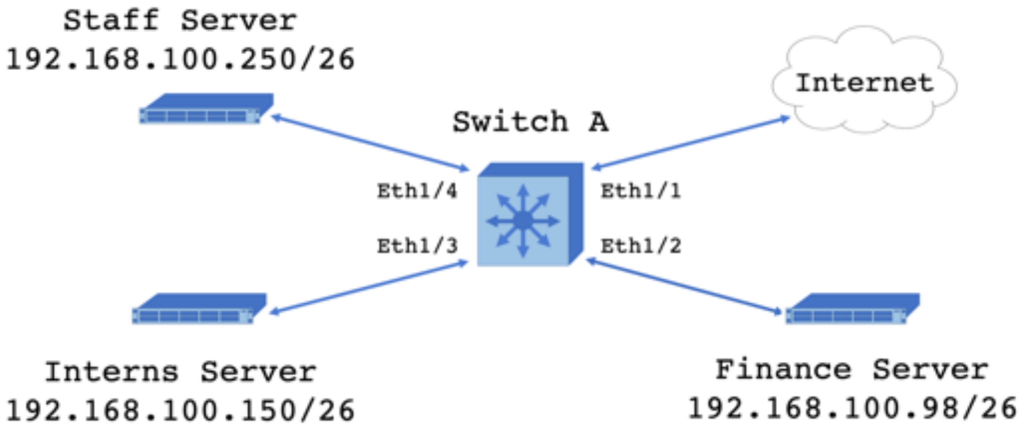

```
Switch_A(config)# ip access-list 101 permit ip
192.168.100.64 0.0.0.144 any
Switch_A(config-acl)# config Ethernet 1/1
Switch_A(config-if)# ip access-group 101 out
```

a. True

b. False

Answer: b (False)

Yes, I know, ACLs on outputs requires the switch process frames even though it is going to drop them, but it's a Nexus switch and it has horsepower for days (assume in this example that this ACL does not significantly impact the performance of the switch). In this example, a CIDR 26 is a subnet mask of 255.255.255.192 (or 0.0.0.144 as in wildcard notation), so the Interns Server is now on the 192.168.100.128 subnetwork. In the ACL syntax you put the source host or network first, then the destination host or network second ("any" in this example). The ACL is only allowing hosts from the 192.168.100.64 subnetwork to access the Internet (so the Finance Server has access to the Internet but the Staff Server does not). Also, recall that there is an implicit "deny all" at the end of the ACL, so there are two reasons that the Intern Server cannot reach the Internet in this example.

32. Be careful when configuring your Nexus 5000 data center switch, you may want to save your previous configuration before implementing one of these modes because your switch will reboot, and you'll lose all your previous settings:

a. NPV+

b. NPIV

c. NVIP+

d. NPV

e. NVP

Answer: d

When your switch is running in NPV (N-Port Virtualization) mode the LAN and SAN are separately administered and the switch does not have to process Fibre Channel (FC), it simply passes FC to the upstream NPIV (N-Port Identifier Virtualization) switch which is processing the FC frames. But, to switch the switch from normal mode to NPV mode it will have to reload and lose its original configurations.

Nexus in edge switch mode:

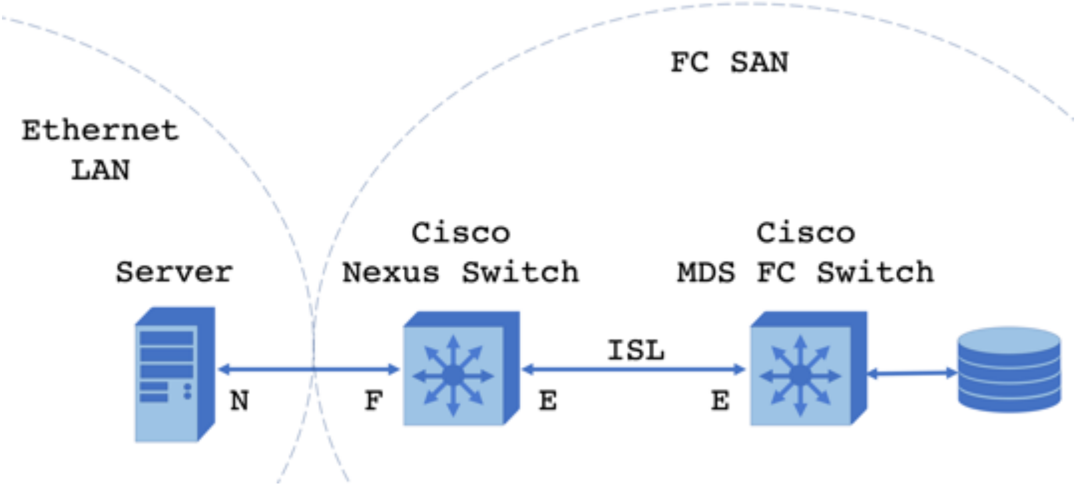

The Cisco Nexus switch is actively processing FC frames (as well as Ethernet frames), and so requires an FCID, an FC domain only has 239 FCIDs.

Nexus in NPV mode:

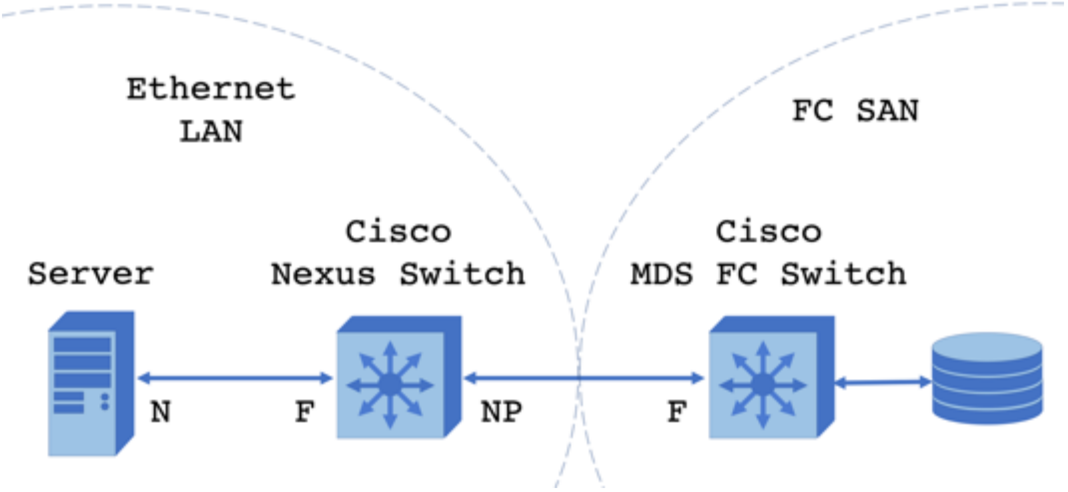

When the Cisco Nexus switch is running in NPV mode, it does not require an FCID and forwards all FC frames to the Cisco MDS switch which is running in NPIV mode.

33. Which option will allow you to view the nWWNs and pWWNs of a switch in NPV mode?

a. Switch_A# **show npv flogi-table**

b. Switch_A# **show npv status**

c. Switch_A# **show npv plogi-table**

d. Switch_A# **show npiv flogi-table**

e. Switch_A# **show npv flogi-db**

f. None of the above

Answer: a

The “show npv flogi-table” command is helpful when troubleshooting and will show the number of devices that are logged in, the internal and external connected interfaces, their VSANs, pWWNs (port World Wide Names), and nWWNs (node World Wide Names). The “show npv status” is also a helpful command to remember, it will not give you the WWNs, but you will see port status. (It may look strange, but yes the “p” and the “n” in pWWN and nWWN are each lower case in the DCICN)

34. What are the three steps in the host to Fibre Channel SAN process? (choose three)

a. FCLGI

b. FLOGI

c. FLOGO

d. PLOGI

e. PRLI

f. FCLI

Answer: b, d, e

The first step is the FLOGI, or Fabric Login, where a node gets an FCID, is added to its local FC-enabled switch’s FLOGI database and is added to the fabric’s shared FCNS (Fibre Channel Name Server) database. The second step is the PLOGI, or Port Login, where the node gets access to storage zones. The final step is the PRLI, or Process Login, which exchanges ULP (Upper Layer Protocol) SCSI handshaking, support information, and commands between Nodes. (Yes, we use the old English spelling in “Fibre” Channel)

35. Which option is a level of security where a LUN is only available to authorized users?

a. Masking

b. Mapping

c. Multiplexing

d. Modulating

e. Mirroring

f. Media Access Control Addressing

Answer: a

Logical Unit Number (LUN) masking makes the LUN visible and accessible to some hosts and servers and not to others, and is configured on the SAN or storage system. LUN mapping is the process of connecting host and servers to specific LUNs and is configured on the host or server. The other options are just other technical words that start with the letter “m.”

36. Which option is a set of standards for connecting, sending, and receiving data between computers and storage devices, and is the basis for Storage Area Network (SAN) communications in data centers? (choose the best answer)

a. IDE

b. USB

c. SAS

d. SATA

e. SCSI

f. None of the above

Answer: e

The Small Computer System Interface (SCSI) defined standardized commands for peripheral devices and was originally used to connect to local or nearby devices such as hard drives, tape drives, printers, scanners, and Zip drives (anyone remember those?). The SCSI read, write, non-data, and bidirectional commands are still used in iSCSI, Fibre Channel (FC), and FCoE (Fibre Channel over Ethernet), and are now carried over Ethernet or FC (rather than local buses) to connect devices to storage over networks.

37. On the Nexus 7000 series switch, what are the maximum number of interfaces that can be bundled together into a single port channel and how many will be active? (at least as of the current DCICN)

a. 24 total, 12 active

b. 8 total, all 8 can be active

c. 16 total, 8 active

d. 16 total, 12 active

e. 16 total, all 16 can be active

f. 32 total, 16 active

Answer: c

While 16 interfaces can be bundled together into a single port channel, only 8 will be active and the other 8 are backups. If one of the 8 active interfaces go down, one of the other 8 interfaces will take over, providing for high resiliency in your data center. (Note, the Nexus 9000 can support 32 interfaces in a port channel)

38. In SAN terminology, what is another word for storage locations? (choose the best answer)

a. initiators

b. targets

c. nodes

d. hosts

Answer: b

In data centers SANs (using iSCSI, Fibre Channel, or FCoE) a data storage location or device on the network is a *target*, while the workstation or host accessing that target over the network is an *initiator*.

39. What options below describe the results of these NX-OS commands? (choose the best answer)

```
switch# configure terminal
switch(config)# interface Ethernet 1/1
switch(config-if)# no switchport
switch(config-if)# no shutdown
```

a. The Layer 2 switch will now become a Layer 3 router

b. The Layer 2 switch will now support Layer 3 routing on interface Ethernet 1/1

c. The Ethernet 1/1 interface is now a switch port and is online

d. The Ethernet 1/1 interface is now an STP blocking port and is online

e. The Ethernet 1/1 interface is now a trunk and will process traffic for multiple VLANs

Answer: b

By default, the interfaces on a Nexus switch are switchports, the "no switchport" command will turn the switching port into a routing port (that will now need an IP address). Cisco's routing ports are off by default, so we use the "no shutdown" command to turn it on. The Layer 2 switch can now support Layer 3 routing on interface Ethernet 1/1.

40. What data transfer protocols are used for block level storage? (choose three)

a. iSCSI

b. NFS

c. CIFS

d. Fibre Channel

e. NTFS

f. FCoE

g. VMFS

Answer: a, d, f

Block level storage is used in Storage Area Networks (SANs), and SANs (in the context of the DCICN) use iSCSI, Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE) for data transfer and remote access by hosts/initiators. All three protocols rely on Small Computer System Interconnect (SCSI) commands. Blocks are independent units of storage often used in SANs, especially in support of virtual machines, and can be formatted to support file systems such as Ext4 for Linux, NTFS for Windows, or VMFS for VMware.

41. What data storage protocols are used for file level storage? (choose two)

a. iSCSI

b. NFS

c. CIFS

d. Fibre Channel

e. Frame Relay

f. FCoE

Answer: b, c

Network File Storage (NFS) was developed for UNIX operating systems and is now also used in Linux systems. Common Internet File System (CIFS) is often used in Windows server/client environments (at least from the Cisco DCICN perspective). Both file level storage standards are often used in the context of Network Attached Storage (NAS), where a host is accessing files on another device/server.

42. Can you match the DCB or related protocols to their descriptions?

Protocol:

Priority-based Flow Control (PFC)

Enhanced Transmission Selection (ETS)

Data Center Bridging Exchange (DCBX)

Link Layer Discovery Protocol (LLDP)

Quantized Congestion Notification (QCN)

Description:

IEEE 802.1Qaz:
guarantees available bandwidth to higher priority processes based on class of services (CoS)

IEEE 802.1Qbb:
Pauses class of services (CoS) based on priority, lower priority traffic is paused to allow for higher priority traffic

IEEE 802.1AB:
Carries DCBX information between devices in a specific frame

IEEE 802.1Qaz:
Allows for communicating operational parameters between DCB capable network devices

IEEE 802.1Qau:
Allows for core data center Layer 2 congestion management by, throttling traffic as close as possible to the source of that traffic

Answer:

IEEE Data Center Bridging (DCB) is a set of Ethernet features that enables newer technologies and processes such as Fibre Channel over Ethernet (FCoE) without impacting traditional Ethernet functions in a data center or network.

Priority-based Flow Control (PFC) = IEEE 802.1Qbb:
Pauses classes of services (CoS) based on priority, lower priority traffic is paused to allow for higher priority traffic.

Enhanced Transmission Selection (ETS) = IEEE 802.1Qaz:
Guarantees available bandwidth to higher priority processes based on class of services (CoS).

Data Center Bridging Exchange (DCBX) = IEEE 802.1Qaz:
Allows for communicating operational parameters between DCB-capable network devices.

Link Layer Discovery Protocol (LLDP) = IEEE 802.1AB:
A standards-based protocol that carries discovery and description information between network devices (similar to Cisco's proprietary CDP), such as DCBX information between devices.

Quantized Congestion Notification (QCN) = IEEE 802.1Qau:
Throttles Layer 2 traffic as close as possible to the source of that traffic before it impacts the overall performance of the core data center.

Drawing the lines we have:

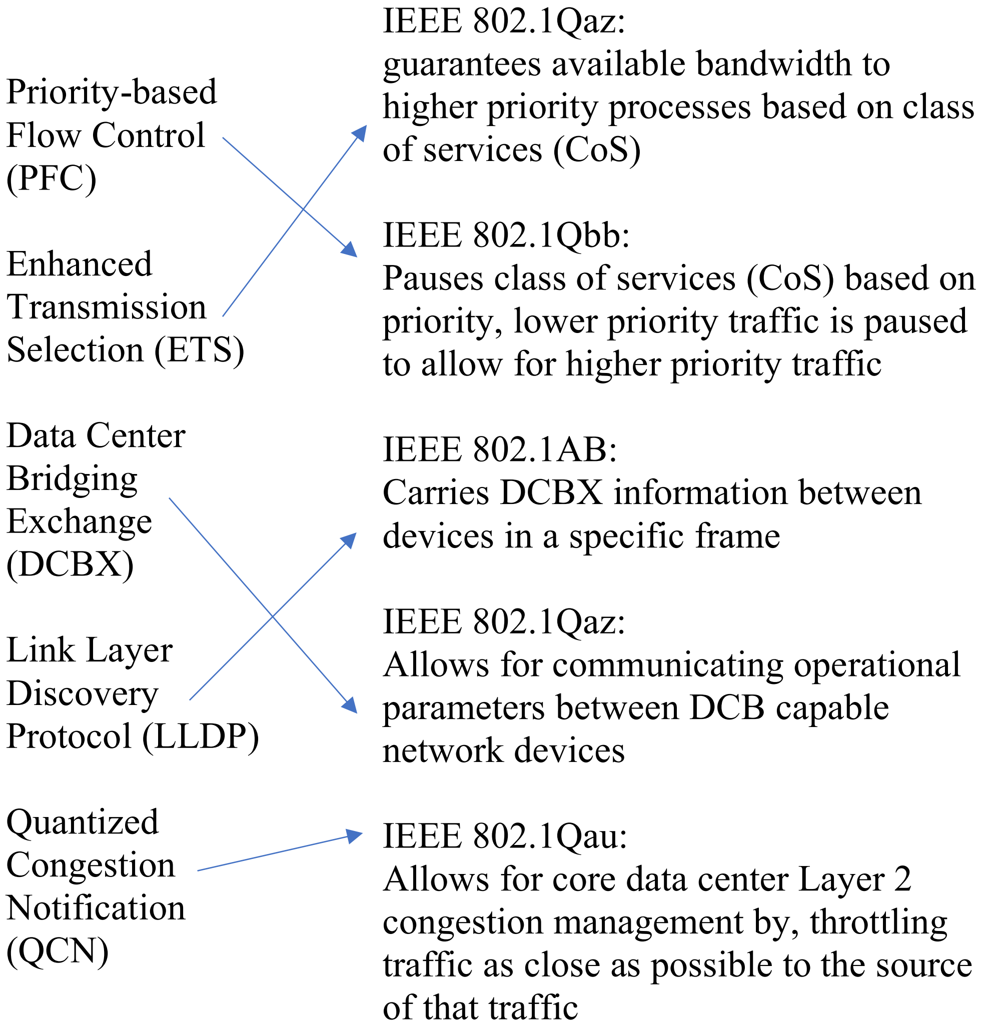

Priority-based Flow Control (PFC)

Enhanced Transmission Selection (ETS)

Data Center Bridging Exchange (DCBX)

Link Layer Discovery Protocol (LLDP)

Quantized Congestion Notification (QCN)

IEEE 802.1Qaz:
guarantees available bandwidth to higher priority processes based on class of services (CoS)

IEEE 802.1Qbb:
Pauses class of services (CoS) based on priority, lower priority traffic is paused to allow for higher priority traffic

IEEE 802.1AB:
Carries DCBX information between devices in a specific frame

IEEE 802.1Qaz:
Allows for communicating operational parameters between DCB capable network devices

IEEE 802.1Qau:
Allows for core data center Layer 2 congestion management by, throttling traffic as close as possible to the source of that traffic

43. Which of these options are not advantages of FCoE? (choose three)

a. FCoE requires both Ethernet and Fibre Channel switches

b. FCoE only requires FC switches to connect to Ethernet hosts

c. FCoE Host Bus Adapters can connect to both Ethernet and Fibre Channel switches

d. FCoE uses the same data center core as Ethernet traffic

e. All of the above

Answer: a, b, c

Option “a” is not an advantage of Fibre Channel over Ethernet (FCoE) because FCoE does not require both Ethernet and Fibre Channel switches. Option “b” is not an advantage of FCoE because FCoE does require Ethernet switches. Option “c” is not an advantage of FCoE because Host Bus Adapters are for Fibre Channel not Ethernet (in the context of the DCICN). And so that leaves “d”, and yes one of the advantages of FCoE is the ability to use your Ethernet core for both Ethernet and for FCoE, so you do not need to invest in a second core infrastructure to run your FC SAN.

44. Which option creates an ACL such that only host 172.16.10.101 and host 172.16.10.102 can access host 172.16.10.103 on Switch_B?

a.
```
Switch_B(config)# ip access-list 102 permit
ip host 172.16.10.101 host 172.16.10.103
Switch_B(config-acl)# permit ip host
172.16.10.102 172.16.10.103
```

b.
```
Switch_B(config)# ip access-list 102 permit
ip host 172.16.10.102/24 and
172.16.10.101/24 and 172.16.10.102/24
```

c.
```
Switch_B(config)# ip access-group 102 permit
ip host 172.16.10.101 172.16.10.102
Switch_B(config-acl)# permit ip
172.16.10.102 172.16.10.103
```

d.
```
Switch_B(config)# ip access-list 102 permit
ip host 172.16.10.101/24 host
172.16.10.102/24
Switch_B(config-acl)# permit ip
172.16.10.102/24 172.16.10.103/24
```

Answer: a

First, we enter terminal config mode, then create IP access list “102”, then permit IP traffic, from the source host (172.16.10.101) to the destination host (172.16.10.103), then a second line from the second host (172.16.10.102) to the destination host (172.16.10.103), and recall there is a “deny all” at the end of the ACL to block all other hosts. (Next, we’d have to apply this ACL to the appropriate interface, such as the inbound side of the interface for 172.16.10.103 on Switch_B)

45. Which of these options are not Private IP address ranges? (choose the best answers)

a. 0.0.0.0 - 127.255.255.255
b. 10.0.0.0 – 10.255.255.255
c. 128.0.0.0 – 191.255.255.255
d. 172.16.0.0 – 172.31.255.255
e. 192.0.0.0 – 223.255.255.255
f. 192.168.0.0 – 192.168.255.255
g. 224.0.0.0 – 239.255.255.255
h. 240.0.0.0 – 254.255.255.255

Answer: a, c, e, g, h

Recall that the private IPv4 private address ranges (that routers will not route on the Internet) are:

10.0.0.0 – 10.255.255.255
172.16.0.0 – 172.31.255.255 and
192.168.0.0 – 192.168.255.255

FYI, the other answer options are the IPv4 network class ranges:

Class A = 0.0.0.0 - 127.255.255.255
Class B = 128.0.0.0 – 191.255.255.255
Class C = 192.0.0.0 – 223.255.255.255
Class D = 224.0.0.0 – 239.255.255.255 (reserved for multicasting)
Class E = 240.0.0.0 – 254.255.255.255 (reserved for research and development)

46. A data center switch has an onboard power supply system that is designed to have modules that can be serviced without the switch losing power, what is this strategy commonly referred to as?

a. An UPS battery system with an inline generator

b. An N+2 power supply system

c. An N+1 power supply system

d. A 2N+1 UPS system

Answer: c

In terms of Unified Computing System server or Nexus hardware, N+1 refers to having an extra power supply module in the system chassis so that the system stays running in the event of a power supply hardware failure. The "N" in "N+1" indicates the number of power supply modules that a server or switch need to maintain minimal operations, while the "+1" means that the server or switch has the minimum number of power supply modules plus an extra module. The extra module means that any other module can fail, be replaced and serviced, and the extra module can be used to take its place without the system losing power.

47. How does the IEEE 802.1Q protocol keep traffic on difference VLANs separate?

a. A four-byte tag is inserted into a new field between the source MAC address field and the Type or Length Field in an Ethernet frame.

b. A four-byte tag is inserted into a new field between the source IP address field and the destination IP address field in an IP packet.

c. An Ethernet frame is assigned to a VLAN using a new frame header and a new cyclical redundancy check specific to that VLAN.

d. An IP packet is assigned to a VLAN using a new packet header and a new cyclical redundancy check specific to that VLAN.

Answer: a

Option “c” describes Cisco’s ISL or Inter-Switch Link functionality which is being phased out by Cisco in preference of the industry standard 802.1Q. The tags are only used to and from trunk connections between network devices, typically switches. Switching is a Layer 2 function based on source and destination MAC addresses; thus options “b” and “d” cannot be correct.

48. What is the startup process for a Nexus switch?

a. Bootloader ➔ BIOS ➔ Kickstart Image ➔ System Image

b. BIOS ➔ Kickstart Image ➔ Bootloader ➔ System Image

c. Kickstart Image ➔ Bootloader ➔ BIOS ➔ System Image

d. BIOS ➔ Bootloader ➔ Kickstart Image ➔ System Image

e. None of the above

Answer: d

On power up, the Nexus first loads its BIOS (Basic Input/Output System), then the Bootloader loads the Kickstart image into RAM, the Kickstart Image loads the system image from Flash memory into RAM where it reads the startup-configuration file from NVRAM and starts up the switch. The switch bootup sequence can be interrupted between these stages when troubleshooting.

49. What aspect of Fibre Channel N-Port Virtualization mode on a Nexus switch is not true:

a. An NPV enabled switch will expend processing and memory resources processing Fibre Channel frames.

b. A 5000 and 7000 series Nexus switch can natively process Fibre Channel frames.

c. Using an NPV enabled switch with an upstream NPIV enabled switch will optimize the use of FCIDs in a Fibre Channel SAN.

d. The LAN and the SAN are administered separately.

Answer: a

NPV (N-Port Virtualization) mode on a Nexus switch will pass Fibre Channel frames on to an upstream F-Port on an NPIV enabled switch, and the NPIV enabled switch will then process the Fibre Channel frames. The NPV switch will not process Fibre Channel, it only forwards the frames to the NPIV switch. The NPIV switch also allows multiple Fibre Channel logins through the same port, saving the limited number of available Fibre Channel logins.

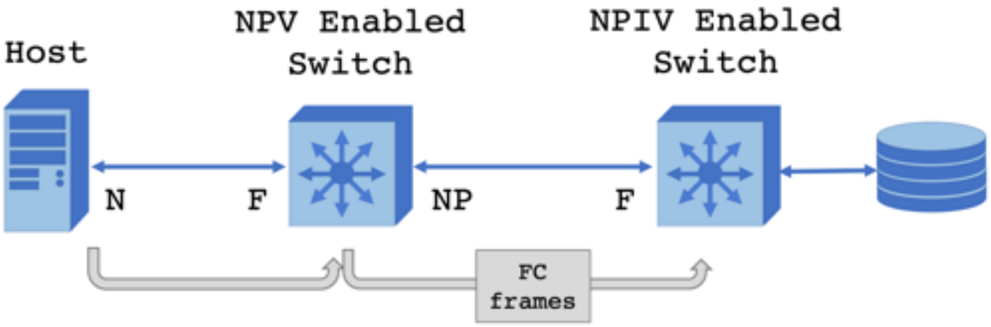

50. Which option is a way the NPIV feature of a switch optimizes available Fibre Channel logins?

a. NPIV enabled switches process Fibre Channel frames sent to it from a downstream NPV switch

b. Multiple Virtual Machines can share the same physical N-Port

c. Each Virtual Machine must have its own Fibre Channel login, its own WWPN, and its own N-Port.

d. The LAN and the SAN are administered separately.

Answer: b

NPIV (N-Port ID Virtualization) mode of switch operation will allow multiple FCIDs to be assigned to different to applications or virtual machines (VMs) behind the same physical host's HBA's N-Port:

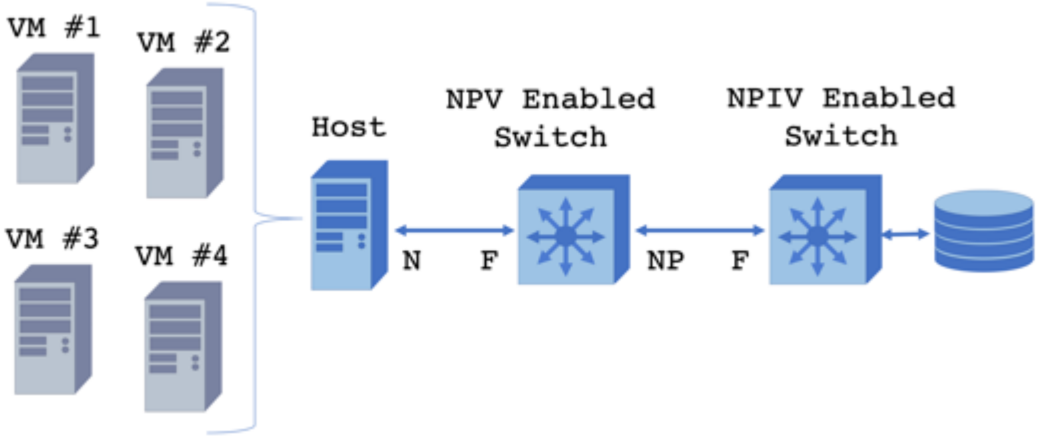

51. What iSCSI naming format was used to name the node: "<*xxx*>.2018-02.com.companyname:video.san.host1"

a. EUI-64

b. IQN

c. EUI

d. NAA

e. NCI

f. DHCP

Answer: b

The IQN (iSCSI Qualified Naming) format results in a human readable address based on a date (typically the date the following domain name was established), the domain name (your domain or other name you or your SAN administrator want to use), and then a device name (also named by you or your SAN admins). So, the complete IQN is "iqn.2018-02.com.companyname:video.san.host1". The other naming schemes are the IEEE EUI's (Extended Universal Identifier) where the device's manufacturer assigns a 64-bit hexadecimal name, and the T11's (the T11 is an ITU standards committee for FC) NAA (Network Address Authority) where the device's manufacturer also assigns a 64 or 128-bit hexadecimal name to the device.

52. Each device connected to a Fibre Channel Storage Area Network must have what two WWNs? (choose two)

a. fcidWWN

b. ipWWN

c. nWWN

d. macWWN

e. pWWN

Answer: c, e

Each device in a Fibre Channel (FC) network has a 64 or 128-bit World Wide Name (WWN) that the IEEE assigned to that device's manufacturer. The device's Node World Wide Name (nWWN) and Port World Wide Name (pWWN) are derived from the WWN. For instance, if a dual port FC Host Bus Adapter (HBA) has the WWN 01:02:03:04:ab:cd:ef:00 the nWWN would be 01:02:03:04:ab:cd:ef:00 and the pWWN of the first port on the HBA could be 01:02:03:04:ab:cd:ef:01 and the second port on the HBA could have the pWWN 01:02:03:04:ab:cd:ef:02 (note, pWWNs are also referred to as WWPNs, or World Wide Port Names, by some vendors).

53. What are three common connection topologies for Fibre Channel? (choose three)

a. point-to-point

b. point-to-multipoint

c. arbitrated loop

d. token ring loop

e. logical bus

f. switched fabric

Answer: a, c, f

In the context of Cisco data centers, the three methods to interconnect Fibre Channel devices is point-to-point (two devices directly connected to each other, so not really a ‘network’ topology), arbitrated loop, and the switched fabric. An arbitrated loop FC topology uses 8-bit Arbitrated Loop Physical Addresses (AL-PA) that allow for 126 nodes in the loop. The more popular method, especially in terms of modern data centers and SANs is the switched fabric (a multi-star topology in terms of classic LAN topologies). The switched fabric is a topology of interconnected FC capable switches that provide reliability, high-speed connections, and alternate paths in case of link failures.

54. What are the three sections of an FCID? (choose three)

a. 8-bit domain ID

b. 12-bit network ID

c. 8-bit area ID

d. 8-bit port number

e. 8-bit network ID

f. 24-bit WWN prefix

Answer: a, c, d

In a Fibre Channel Fabric, in addition to WWNs each device must also have a Fibre Channel Identifier (FCID). An FCID is 24-bits and is comprised of an 8-bit domain, an 8-bit area ID, and an 8-bit port number. Because of the 8-bit domain ID limitations (and some reserved domains), there can only be 239 total devices in a single FC SAN.

55. In the context of the current DCICN, which Cisco switches does Cisco recommend administrators typically deploy as NPV switches? (choose the best three options)

a. Nexus 3000 series switches

b. Nexus 4000 series blade switches

c. Nexus 5000 series switches

d. UCS 6200 series fabric interconnects

e. Nexus 7000 series switches

f. Nexus 9000 series switches

Answer: b, c, d

In practice, and in the context of the DCICN, you would typically use a Nexus 4000, 5000, or a UCS 6200 as an NPV switch, as these are the switches you would normally use in your Access or Leaf layer of your data center. The Nexus 7000 and 9000 would be typically used up in the Collapsed Core or Spine layer, and the Nexus 3000 is a low latency Access/Leaf layer switch typically used in big data applications (that's the best answer for now, though I'm probably going to have to update this question in the next edition of this book…)

56. Which two hosts are in the same Class B subnet as the address 10.227.82.154/16?

a. 10.227.0.0

b. 10.227.255.255

c. 10.227.255.0

d. 10.227.0.255

e. 10.0.227.227

f. none of the above

Answer: c, d

Here's another subnetting example, sorry. If it has been a while since your CCENT (no the CCENT is not required to take the DCICN, but having a working knowledge of subnetting will be helpful) 10.227.0.0 is the subnet, 10.227.255.255 is the broadcast IP addresses of the subnet 10.227.0.0/16, and 10.0.227.227 is an IP address outside of the 10.227.0.0/16 subnet. Both 10.227.255.0 and 10.227.0.255 may look odd (if you are used to using Class C subnets) but are valid IP addresses that can be assigned to hosts in the 10.227.0.0/16 subnet.

57. Which of these routing protocols is a distance vector protocol?

a. EIGRPv0

b. RIPng

c. OSPFv2

d. None of the above

e. All of the above

Answer: b

Routing Information Protocol (RIP) determines the shortest path to a network based on the number of connections, or hops, between the source router and the desired destination for its packets. RIPv1 was classful, while RIPv2 adds support for variable length subnet masks. RIPng adds support for IPv6.

58. Which of these routing protocols is a Cisco developed protocol?

a. EIGRPv1

b. RIPng

c. OSPFv3

d. None of the above

e. All of the above

Answer: a

EIGRP, or Enhanced Interior Gateway Protocol, is Cisco's response to the shortcomings of RIP and OSPF, and combines the best of both of these routing protocols. EIGRP uses the Diffusing Update Algorithm (DUAL) to determine the 'cost' of a feasible route based on distance and delay, the route with the lowest costs are used to reach destinations. EIGRP (v1) is classless and supports variable length subnet masks.

59. Which of these routing protocols is a link state routing protocol?

a. EIGRPv1

b. RIPv2

c. OSPFv2

d. None of the above

e. All of the above

Answer: c

Open Shortest Path First (OSPF) bases its routing decisions on the total ‘cost’ of a path to a destination based on the speed of the connections. The overall fastest routes are saved in a router’s routing table. OSPFv2 added enhanced prevention for IPv4 routing loops, OSPFv3 added support for IPv6, and both versions support classless subnets.

60. Can you match these adapters to their uses? (choose the best answer)

NIC

vNIC

HBA

CNA

SPF

SPF+

QSPF+

Can connect a host to both an Ethernet and a Fibre Channel network

A representation of an Ethernet interface

Connects a host to an Ethernet network

Connects a switch to another 10Gbps device via copper or fiber

Connects a host to a Fibre Channel network

Combines four 10Gbps adapters

Connects a switch to another 1Gbps device via copper or fiber

Answer:

NIC (Network Interface Card) = Typically connects a host to an Ethernet network via copper Cat5/6 cables with RJ-45 connectors.

vNIC (Virtual Network Interface Card) = A virtualized representation of an Ethernet interface typically used to connect multiple virtual machines to other devices via a single physical NIC (or multiple physical NICs).

HBA (Host Bus Adapter) = Typically connects a host to a Fibre Channel network via dual send/receive fiber optic cables.

CNA (Converged Network Adapter) = Can connect a host to both an Ethernet and a Fibre Channel network by combining the functionality of a NIC and an HBA (this card has drivers for both Ethernet and Fibre Channel).

SFP (Small Form-factor Pluggable) = A switch module that connects a switch to another 1Gbps device via copper or fiber.

SFP+ (Enhanced Small Form-factor Pluggable) = A module that connects a switch to another 10Gbps device via copper or fiber.

QSFP+ (Quad Small Form-factor Pluggable) = Uses a break out cable to connect a single 40Gbps connection to four 10Gbps adapters on another switch.

Drawing in the lines we have:

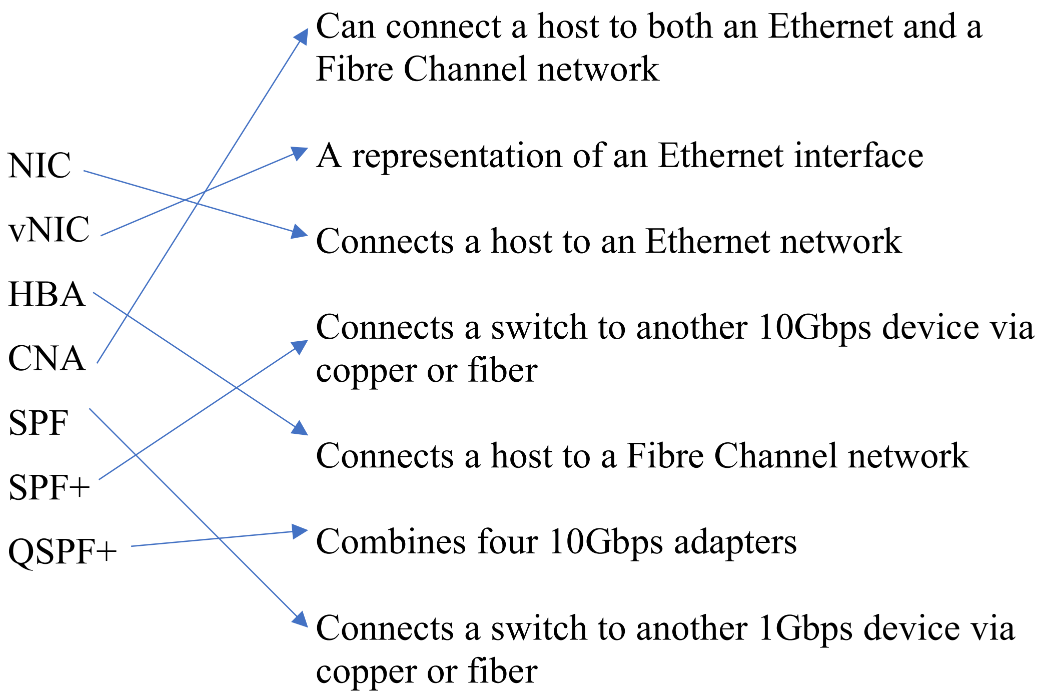

NIC

vNIC

HBA

CNA

SPF

SPF+

QSPF+

Can connect a host to both an Ethernet and a Fibre Channel network

A representation of an Ethernet interface

Connects a host to an Ethernet network

Connects a switch to another 10Gbps device via copper or fiber

Connects a host to a Fibre Channel network

Combines four 10Gbps adapters

Connects a switch to another 1Gbps device via copper or fiber

61. Can you match the RAID levels to the descriptions?

RAID 0	Mirroring across disks
RAID 1	Striping across disks
RAID 5	Mirroring with double distributed parity
RAID 6	Mirroring with distributed parity
RAID 1+0	Striping with mirroring
RAID 0+1	Mirroring with striping

Answer:

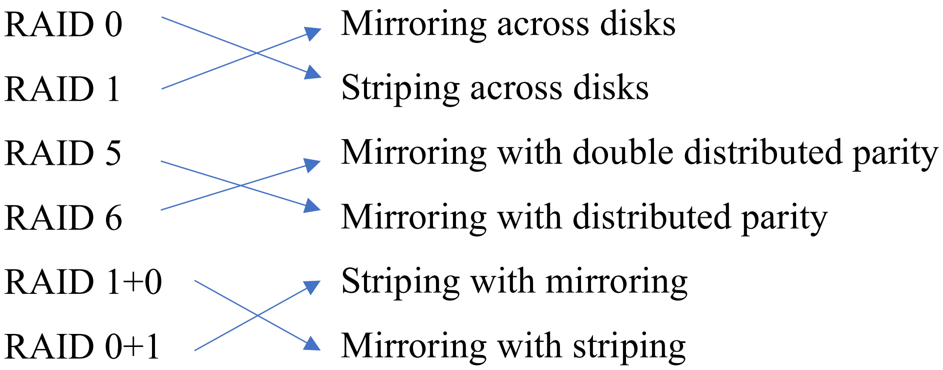

These are the most common Redundant Array of Independent Disks (RAID) configurations (RAID 2, 3, and 4 are not used often), with RAID 6 probably being the most cost effective in larger arrays. Parity is an error correction method where information on a failed hard drive can be rebuilt on a replacement drive.

RAID 0 or disk striping, also known as Just a Bunch of Disks (JBOD), combines the disks into a single, larger logical drive. Storage blocks A1 to A6 are saved across both drives/disks:

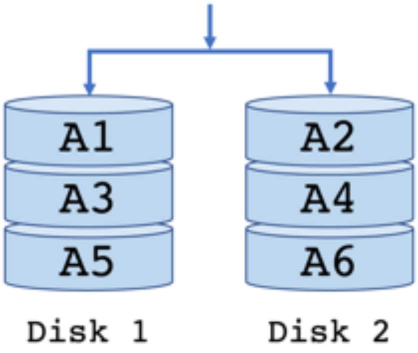

RAID 1, or disk mirroring, a disk drive can be lost without losing data, storage blocks are written to both drives/disks:

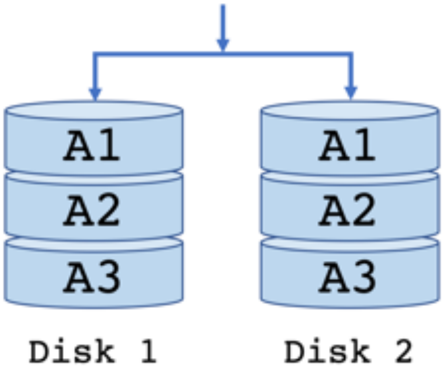

A RAID 5 configuration requires at least a third disk to store parity (for example, *AP, BP*, and *CP* in the diagram below) information for the block being stored, the storage block on a lost drive can be rebuilt based on the parity information stored on one of the other drives:

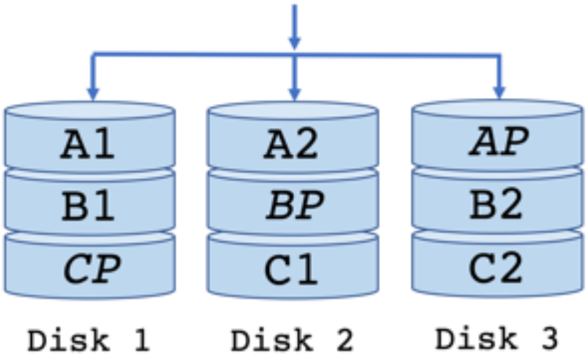

In a RAID 6 configuration, the storage block parity is stored on two drives, the array can now lose two drives and still have the storage block rebuilt based on the parity information:

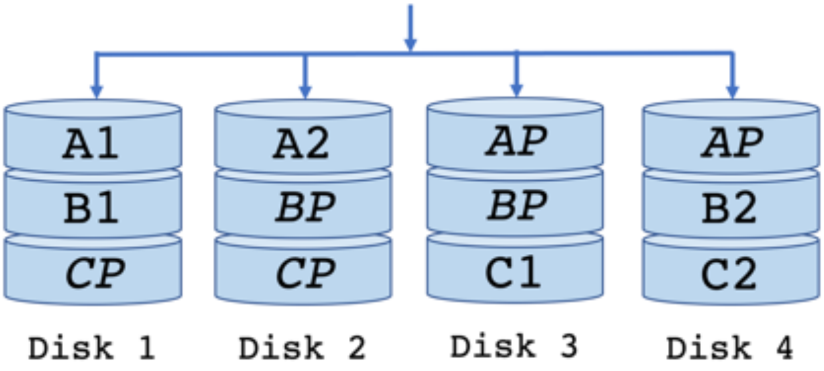

In RAID 1+0 storage configuration, storage blocks are mirrored in disk pairs, a drive can be lost without losing data:

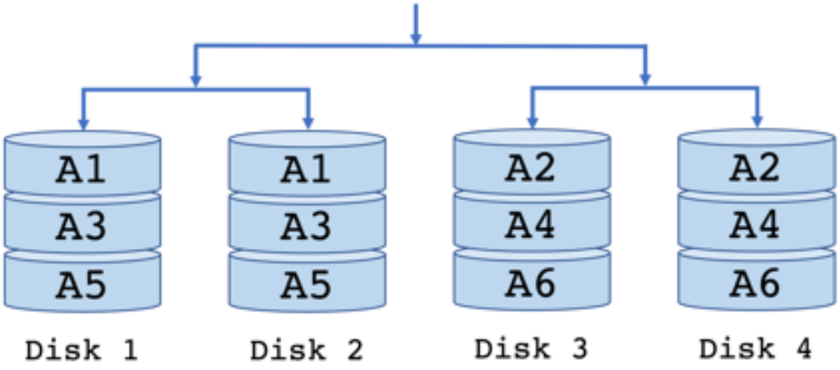

In a RAID 0+1, storage blocks are also mirrored or saved twice in disk drives groups in alternating pairs of drives:

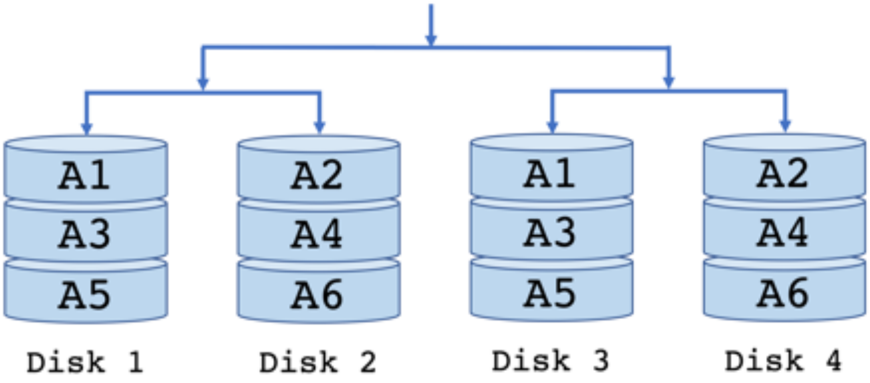

62. How will the switch respond after entering the following commands (chose one):

```
Switch_A(config)# feature interface-vlan
Switch_A(config)# interface vlan 30
Switch_A(config-if)# ip address 192.168.30.1/24
Switch_A(config-if)# no shutdown
```

a. You will get an NX-OS syntax error message

b. The switch will create VLAN 30

c. The switch will create an SVI for VLAN 30

d. The switch will create a VPN on VLAN 30

e. None of the above

Answer: c

The switch will create a Switched Virtual Interface (SVI) that will allow for routing packets between VLANs on a Nexus switch. The "feature interface-vlan" activates the SVI Layer 3 feature, "interface vlan 30" creates an SVI for VLAN 30 (note, for this question assume you have already created VLAN 30 on this switch), "ip address 192.168.30.1/24" assigns an IP address to the SVI, this will likely be the default gateway for the devices on the 192.168.30.0/24 network. The last step is to the turn on the interface with the "no shutdown" command.

63. How will the switch respond after entering the following commands:

```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/1
Switch_A(config-if)# switchport mode trunk
Switch_A(config-if)# switchport native vlan 2
Switch_A(config-if)# switchport trunk allow
vlan 2-12
Switch_A(config-if)# exit
Switch_A(config)# exit
Switch_A# copy running-config startup-config
```

a. Interface 1/1 will become a trunk

b. Interface 1/1's native vlan will become vlan 2

c. Interface 1/1 will pass frames for VLANs 2 through 12

d. Switch_A will save its current running configuration as its new startup configuration

e. All of the above

f. None of the above

Answer: e

A trunked interface will forward Layer 2 frames to and from different VLANs on the same physical connection between the switches (this way VLANs can extend across multiple physical switches). This example (on the previous page or copy and pasted below) specifically enters the global configuration mode from the privileged EXEC mode. Next, we enter the interface configuration mode for the interface at switch slot 1 and port 1. We then define this interface as a trunk. Next, we assign a new native VLAN (the default is VLAN 1). We ask the trunk to send and receive frames for VLAN 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, and VLAN 12. In this example, we then exit from the interface configuration mode to the global configuration mode, then exit from the global configuration mode to the EXEC mode (using the ‘exit’ command will put us back into the previous configuration mode). Finally, we save the new settings with the “copy running-config startup-config”. These settings will now still be active if/when we reboot the switch:

```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/1
Switch_A(config-if)# switchport mode trunk
Switch_A(config-if)# switchport native vlan 2
Switch_A(config-if)# switchport trunk allow
vlan 2-12
Switch_A(config-if)# exit
Switch_A(config)# exit
Switch_A# copy running-config startup-config
```

64. How will the switch respond after entering the following commands:

```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/2
Switch_A(config-if)# ip access group 101 in
Switch_A(config-if)# access-list 101 deny tcp
any any eq 23
Switch_A(config-if)# access-list 101 permit ip
any any
Switch_A(config-if)# exit
Switch_A(config)# exit
Switch_A# copy running-config startup-config
```

a. Interface Ethernet 1/2 will block all traffic

b. Interface Ethernet 1/2 will block inbound TFTP

c. Interface Ethernet 1/2 will block inbound TCP

d. Interface Ethernet 1/2 will block inbound Telnet

e. All of the above

f. None of the above

Answer: d

In this example, ACL 101 will block TCP Telnet (port 23) traffic, but allow other traffic to pass. Please remember the implicit 'deny all' that we do not see but is at the end of the ACL list, so we need the 'permit any', and we also save the running-configuration again.

Also, recall some of the common IP port numbers you may see, especially in the context of ACLs:

Protocol	TCP/UDP	Port Number
FTP (File Transfer Protocol)	TCP	20/21
SSH (Secure Shell)	TCP	22
Telnet	TCP	23
SMTP (Simple Mail Transfer Protocol)	TCP	25
DNS (Domain Name System)	TCP/UDP	53
DHCP (Dynamic Host Configuration Protocol)	UDP	67/68
TFTP (Trivial File Transfer Protocol)	UDP	69
HTP (Hypertext Transfer Protocol)	TCP	80
Post Office Protocol (POP)	TCP	110
NTP (Network Time Protocol)	UDP	123

65. What components are not in a C-series UCS server? (choose two)

a. Central Processing Units

b. RAM

c. Fabric Extenders

d. Redundant power supplies

e. Redundant cooling fans

f. Front and rear mezzanine slots

Answer: c, f

The USC C-series is Cisco's line of rack-mounted servers, it has sockets for CPUs, DIMM slots for RAM, redundant power supplies, and redundant cooling fans. Cisco's B-series blade servers also have sockets for multiple CPUs and DIMM slots, though these servers also have front and rear mezzanine slots for VICs (Virtual Interface Cards), additional RAM, and RAID controllers. The B-series servers do not have onboard power supplies or cooling fans, those are included in the UCS 5108 chassis. Hard drives are optional for both series of servers. Fabric Extenders are separate components.

66. How will the switch respond after entering the following commands? (choose two)

```
Switch_A# configure terminal
Switch_A(config)# zone name ProductionZone vsan
2
Switch_A(config-zone)# member pwwn
12:34:56:78:9a:bc:ef:11
Switch_A(config-zone)# exit
Switch_A(config)# exit
Switch_A# copy running-config startup-config
```

a. A new “ProductionZone” zone will be created in VSAN 2

b. A new “VSAN 2” will be created in “ProductionZone”

c. The device with the Port World Wide Name 12:34:56:78:9a:bc:ef:11 will be added to VSAN 2

d. The device with the Port World Wide Name 12:34:56:78:9a:bc:ef:11 will be added to “ProductionZone”

Answer: a, d

This series of commands will create a new “ProductionZone” zone within VSAN 2 and add the device with the pWWN (Port World Wide Name) 12:34:56:78:9a:bc:ef:11 to ProductionZone. Recall that each port on a Fibre Channel device in our SAN will have a Port World Wide Name (so a dual port HBA will have two pWWNs). A zone is used as an extra layer of security to organize devices and the data that they should have access to. In general, you want to restrict initiator access to only the data that they need, putting your initiators in zones is one way to do that.

67. How will the switch respond after entering the following commands?

```
Switch_A# configure terminal
Switch_A(config)# zone name ProductionZone vsan
2
Switch_A(config-zone)# member pwwn
12:34:56:78:9a:bc:ef:11
Switch_A(config-zone)# exit
Switch_A(config)# zoneset name
ProductionZoneset
Switch_A(config-zoneset)# zone name
ProductionZone
Switch_A(config-zoneset)# exit
Switch_A(config)# exit
Switch_A# copy running-config startup-config
```

a. A new “ProductionZone” zone will be created in VSAN 2

b. The device with the Port World Wide Name 12:34:56:78:9a:bc:ef:11 will be added to “ProductionZone”

c. The “ProductionZone” zone will be added to the “ProductionZoneset” zone set

d. We save the running configuration

e. All of the above

Answer: e

I hope you read all the answer options, and didn’t stop at option “a”, which is only partially correct. These are the steps to create a zone, add a device to that zone, and add a zone to a zone set. Recall that each VSAN can only have one active zone set. Also, devices can be added to a zone using the device’s pWWN, FCID, or switch interface.

68. Which statement about Port Channel modes are accurate:

a. An “active” mode port on one device connected to an “active” mode port on another device will form a port channel.

b. An “active” mode port on one device connected to a “passive” mode port on another device will form a port channel.

c. A “passive” mode port on one device connected to a “passive” mode port on another device will not form a port channel.

d. An “on” mode port on one device not running LACP will not form a port channel with an “active” or “passive” mode port on another device.

e. All are correct

Answer: e

Recall that interfaces on connected switches can be combined into a port channel to take advantage of physical connection redundancy and increased bandwidth. Connected ports in Link Aggregation Control Protocol (LACP) “active” mode can form port channels between themselves or with other ports in “passive” mode. Note, “feature lacp” must be active on the switch.

69. Match the device to the layers that the device operates on:

Device:	Layer:
Hub	1. Physical
Bridge	2. Data Link
Repeater	3. Network
Router	4. Transport
Switch	5. Session
Firewall	6. Presentation
UCS Server	7. Application

Answer:

A Hub operates at Layer 1, the Physical layer; it does not make frame forwarding decisions based on Data Link or Network layer protocols,

A Bridge operates at Layer 2, the Data Link layer; it can process packets between collision domains (like a switch). Yes, no one uses hubs and bridges anymore, but Cisco wants us the know the fundamentals, and how hubs and bridges evolved into modern switches is fundamental.

A Repeater operates at Layer 1, the Physical layer, it repeats the entire structure of the ones and zeros it receives,

A Router operates at Layer 3, the Network layer, it processes packets between networks,

A general purpose switch operates at Layer 2, the Data Link layer, it processes (learns, filters, forwards, and maybe drops) frames across its network (note, Nexus switches can operate on both Layer 2 and Layer 3 depending on how its interfaces are configured),

A firewall, like a Cisco APA (Adaptive Security Appliance), typically processes Layer 2 frames and Layer 3 packets, however, there are more and more firewalls that can operate at/analyze all seven layers,

A Cisco UCS (Unified Computing System) server being used as a host for virtual machines in a data center can also be operating at all seven layers of the OSI model. The operation of a modern firewall and a UCS are good examples of the need for a common framework like the OSI model to describe device and network functionality.

Technically, one can argue that all of these devices also operate on Layer 7, because they all have hardware in them that will convert/format 1's and 0's into frames and packets, and packets and frames back into 1's and 0's, but in the scope of the DCICN:

Hub = Physical
Bridge = Data Link
Repeater = Physical
Router = Network
Switch = Data Link

70. What will a switch do with a frame when it doesn't know that frame's destination MAC address?

a. The switch will drop the frame,

b. The switch will drop the frame and send out an ICMP message to the source (unless the switch has a default route, then it will send the frame out to the default route),

c. The switch will flood the frame out all ports,

d. The switch will flood the frame out all ports, except the port the frame originally arrived from,

e. None of the above

Answer: d

The switch will broadcast the frame out to all connected devices if it doesn't have the frame's destination MAC address in its address table (except it will not send the frame back out to the device that sent the frame).

Note, a switch typically doesn't ARP if it receives a frame with a destination MAC address it doesn't already know, but it will send a host's ARP request out all ports (except the host's port) if that host is asking for the MAC address of another device.

71. What will a router do with a packet when it doesn't know the location of that packet's destination IP address?

a. The router will drop the packet,

b. The router will drop the packet and send out an ICMP message to the source (unless the router has a default route, then it will send the packet out to the default route),

c. The router will flood the packet out all ports,

d. The router will flood the packet out all ports, except the port the packet originally arrived from,

e. None of the above

Answer: b

The router will not flood packets out its ports, rather it will drop the packet, unless it has a default route to use, then it will send the packets to that default route.

72. After reviewing the results of the “show interface e1/1 switchport” command below, choose the three best options below that will describe the operation of Switch_A:

```
Switch_A# show interface e1/1 switchport
Name: Ethernet 1/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
Protected: false
Appliance trust: none
```

a. Interface Ethernet 1/1 is acting as a trunk

b. packets processed out Ethernet 1/1 will be given a new 26-byte header and a new 4-byte FCS trailer

c. frames processed out Ethernet 1/1 will be given a new 26-byte header and a new 4-byte FCS trailer

d. frames for VLAN 1 will not be given a new 26-byte header and a new 4-byte FCS trailer

e. packets for VLAN 1 will not be given a new 26-byte header and a new 4-byte FCS trailer

Answer: a, d, e

I didn’t mean for this to be a trick question, but you may see some questions organized like this (please read those questions on the exam carefully!). Yes, Ethernet 1/1 is a trunk, and so it is forwarding Layer 2 frames for various VLANs. But “dot1q”, or the IEEE 802.1Q method of VLANs, uses a tag inserted in the frame header. Cisco’s ISL (Inter-Switch Link) VLAN method encapsulates the original frames between a new 26-byte header and a new 4-byte FCS trailer.

And so, frames for VLAN 1 will not be given a new 26-byte header and a new 4-byte FCS trailer (802.1Q can work across standard Ethernet connections and frame sizes, which makes it increasingly more popular that ISL). Also, both 802.1Q and ISL function at Layer 2 (frames), not Layer 3 (packets), so option “b” can’t be right and there are now two reasons “e” is right. In addition, VLAN 1 is also the native VLAN, and so in 802.1Q it does not need a VLAN tag, so there are two reasons option “d” is right.

73. Which of these options are <u>not</u> differences between Cisco's Nexus NX-OS operating system and Cisco's IOS operating system?

a. You are in EXEC mode as soon as you log into an NX-OS device

b. NX-OS uses a feature based licensing model

c. In the NX-OS, all Ethernet interfaces are called "Ethernet" interfaces

d. In the NX-OS, SSH is enabled by default

e. In the NX-OS, Telnet is disabled by default

f. NX-OS uses a kickstart image and a system image

g. All of the above

Answer: g

Options "a, b, c, d, e, & f" are all features that are in Nexus switches that are not in IOS devices like Catalyst switches (note, the capital "I" in IOS, not to be confused with Apple's iOS in iPhones and iPads). NX-OS includes and installs with all available features, but you have to buy and unlock those features with license keys. The IOS is different, you have to also buy the license to whatever feature you want, then you also have to upload and install the new features. In NX-OS you don't have to worry about which interface is 10Mbs Ethernet, 100Mbps Ethernet, 1Gbps Ethernet, 10Gbps Ethernet or whatever speeds, they are all simply called "Ethernet" interfaces. The NX-OS has a kickstart based on a Linux kernel with basic drivers and file systems and it also has a system image with the rest of the operating system. Also, for old school techs like me the "write memory" command is gone, and so get used to "copy r s" (an abbreviation of "copy running-config startup-config") to save all your hard work.

74. Referring to the results of the “show vlan” command below, why is a host on Ethernet 1/23 not able to communicate with a database server on Ethernet 1/2? Also, consider your next steps in troubleshooting if you ran into this issue in your data center: <u>(choose the best two options)</u>

```
Switch_A# show vlan
VLAN  Name      Status    Ports
---- -------  ---------   --------------
1    default     active   Eth 1/1
2    Accounting  active   Eth 1/2, Eth 1/3
                          Eth 1/4
3    TechSupport active   Eth 1/5, Eth 1/6
                          Eth 1/7, Eth 1/8
4    HR          active   Eth 1/9, Eth 1/10
5    Payroll     active   Eth 1/11, Eth 1/12
                          Eth 1/13
6    Operations  active   Eth 1/14, Eth 1/15
                          Eth 1/16, Eth 1/17
                          Eth 1/18
7    Staff       active   Eth 1/19, Eth 1/20
                          Eth 1/21, Eth 1/22
                          Eth 1/23, Eth 1/24
8    Research    active   Eth 1/25, Eth 1/26
9    Lab         active   Eth 1/27, Eth 1/28
                          Eth 1/29, Eth 1/30
                          Eth 1/31, Eth 1/32
VLAN  Type   Vlan-mode
---   -----  ---------
1     enet   CE
2     enet   CE
3     enet   CE
4     enet   CE
5     enet   CE
6     enet   CE
7     enet   CE
8     enet   CE
9     enet   CE
```

(choose the best two options):

a. The interface that the host is on needs to be added to one of the VLANs

b. The host and the server are on different VLANs

c. There also needs to be a way to route between VLANs

d. There also needs to be a way to route between the Accounting VLAN and the Staff VLAN

Answer: b, d

The host on interface Ethernet 1/23 of Switch_A is on VLAN 7 or the "Staff" VLAN, while the database server it is trying to reach is connected to the interface Ethernet 1/2 on VLAN 2, or the "Accounting" VLAN. Assuming we want to allow the staffer access to the accounting database, the host on Ethernet 1/23 and the database on Ethernet 1/2 need either be on the same VLAN or there needs to be a way to route specifically between the Staff and the Accounting VLANs (option "d" is a more specific answer than option "c"). The "show ip interface" and the "show ip route" commands would tell us if there is a route between VLANs. The "show vlan", "show ip interface" and "show ip route" are all commands you'll want to remember if you have to troubleshoot VLAN connection issues.

75. Can you match the VTP modes to the switch’s functions and abilities in that mode?

VTP mode:
Server

VTP mode:
Transparent

VTP mode:
Client

VTP mode:
Off

Create, modify, and delete local VLANs, forward advertisements, does not synchronize with other switches

Create, modify, and delete VLANs, sends and receives advertisements, synchronizes with other switches

Does not create, modify, or delete VLANs, receives and forwards advertisements, synchronizes with other switches

Create, modify, and delete local VLANs, does not forward advertisements, does not synchronize with other switches

---

Answer:

VLAN Trunking Protocol (VTP) manages VLAN changes across your network using Layer 2 advertisement messages between switches, so you as the network/data center admin don’t have to log into each switch to make the VLAN update. When a VLAN change is made on the switch in Server mode, those changes are sent down the connected trunks to other switches. Client mode switches will update themselves, Transparent mode switches will pass the changes along but will not update themselves. A switch with VTP mode Off is not participating in VTP (helpful if you have some specialized VLANs you don’t want changing).

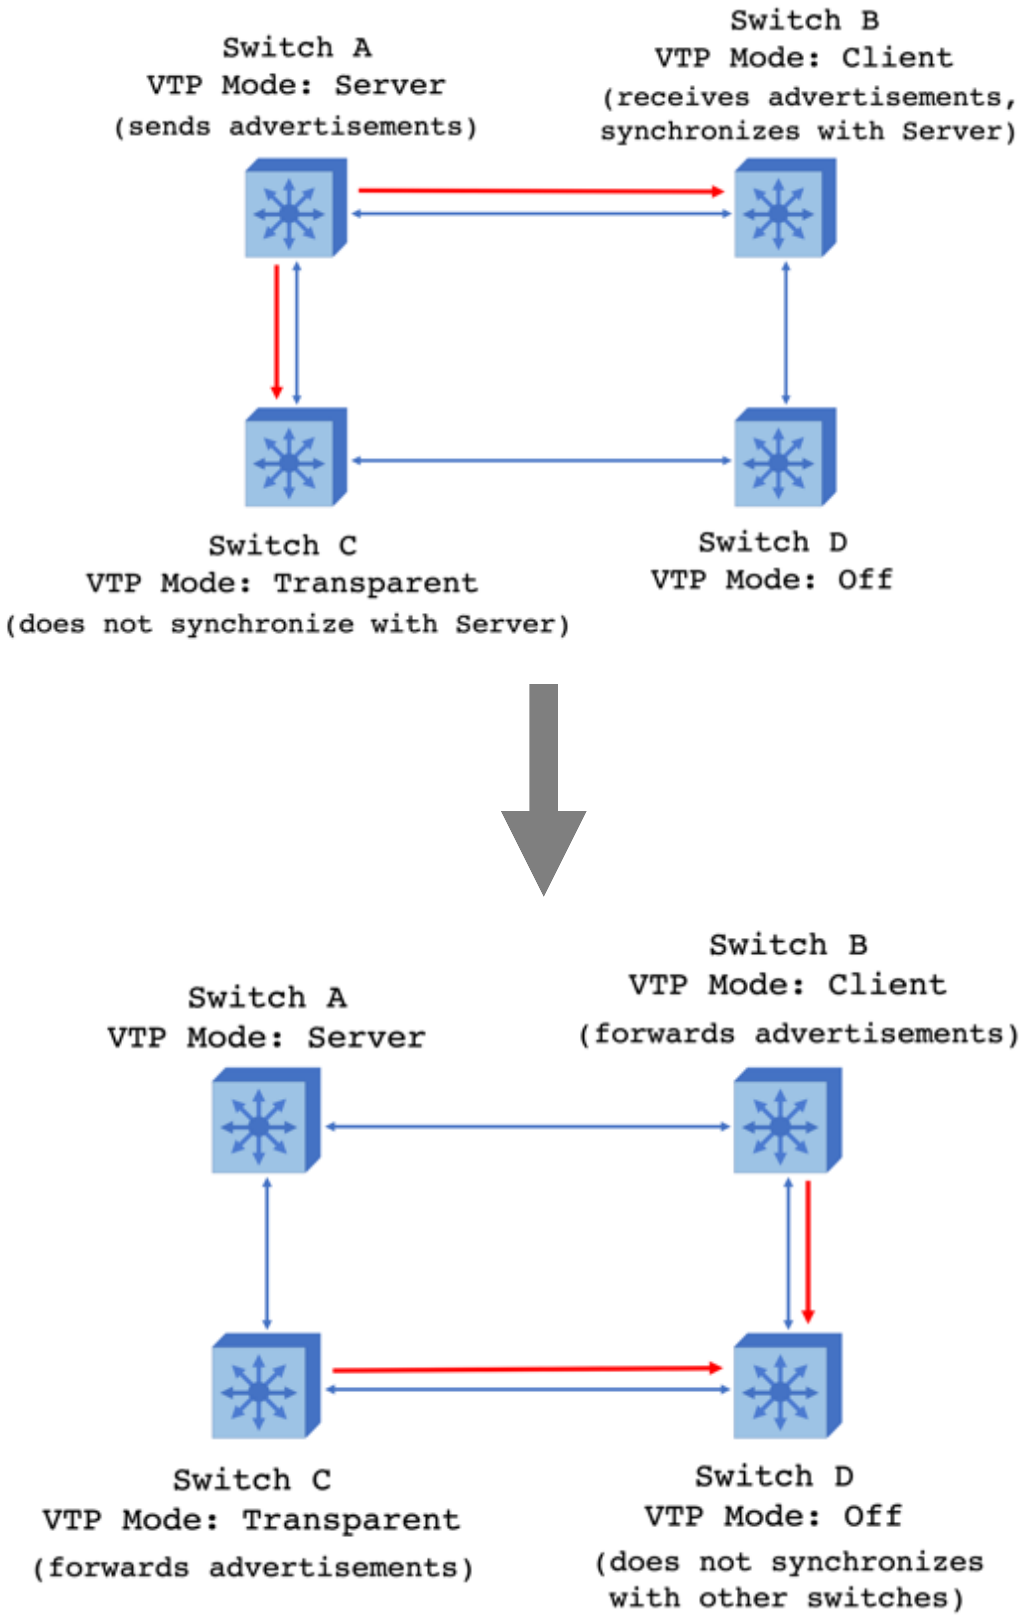
Switch A
VTP Mode: Server
(sends advertisements)
Switch B
VTP Mode: Client
(receives advertisements,
synchronizes with Server)
Switch C
VTP Mode: Transparent
(does not synchronize with Server)
Switch D
VTP Mode: Off
Switch A
VTP Mode: Server
Switch B
VTP Mode: Client
(forwards advertisements)
Switch C
VTP Mode: Transparent
(forwards advertisements)
Switch D
VTP Mode: Off
(does not synchronizes
with other switches)

And so, matching the VTP modes to the switch's functions and abilities in that mode:

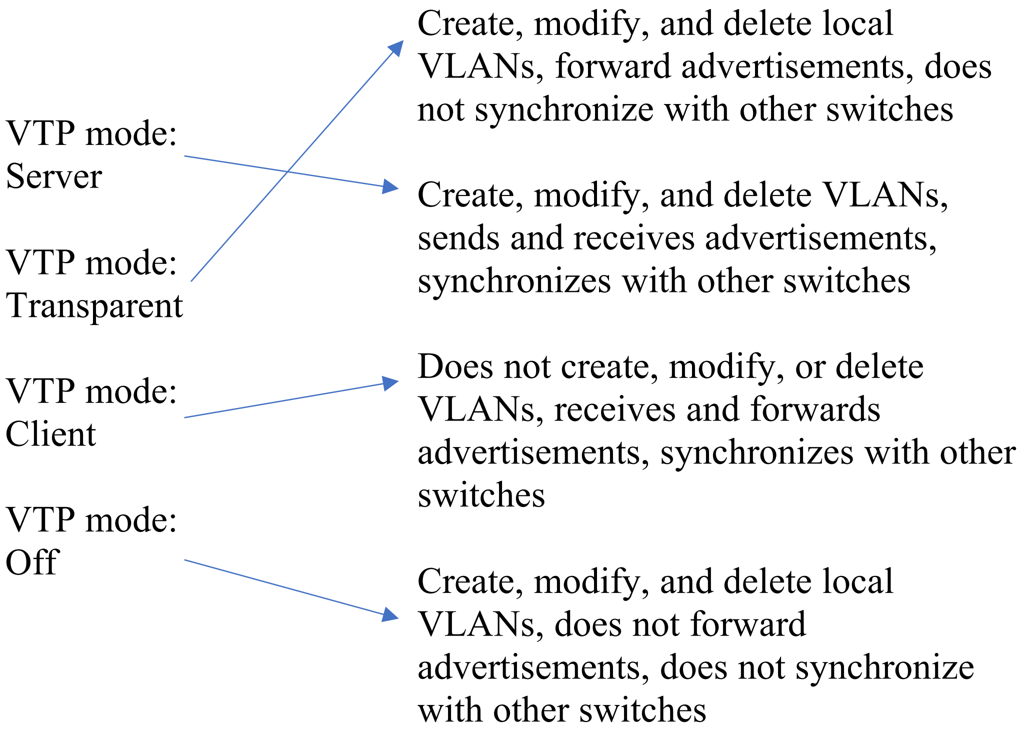

VTP mode: Server

VTP mode: Transparent

VTP mode: Client

VTP mode: Off

Create, modify, and delete local VLANs, forward advertisements, does not synchronize with other switches

Create, modify, and delete VLANs, sends and receives advertisements, synchronizes with other switches

Does not create, modify, or delete VLANs, receives and forwards advertisements, synchronizes with other switches

Create, modify, and delete local VLANs, does not forward advertisements, does not synchronize with other switches

76. What two options are sublayers of the Data Link OSI layer?

a. The LLC Sublayer
b. The MAC Sublayer
c. The TCP/IP Sublayer
d. The IEEE 802.2 Sublayer
e. The Ethernet II Sublayer

Answer: a, b

The LLC (Logical Link Control) is the higher sublayer of the Data Link Layer and includes encapsulation and communication with the Network Layer. The MAC (Media Access Control) is the lower sublayer of the Data Link layer and includes physical access, hardware addresses, and flow and media control technologies (for instance, CMSA/CD traffic control for Ethernet traffic).

77. What option is the correct network and host section of the given Class B IPv4 address?

a. 129.56.235.101, where 129.56 is the network and 235.101 is the host.

b. 192.199.235.101, where 192.199.235 is the network and 101 is the host.

c. 10.10.10.10, where 10 is the network, and 10.10.10 is the host.

d. None of the above

Answer: a

We know that the "129.56…" is the network section of the IP address and "…235.101" is the host section because we are told that this is a Class B address. But note, we are not given the subnet mask, so we do not specifically know how many hosts or how many networks we have…

78. Which option is <u>not</u> an advantage of a static route?

a. A static route is typically considered more secure and reliable.

b. A static route requires much less switch processing resources.

c. Static routes are easier to configure and manage on smaller networks.

d. “Switch_A(config)# ip route 0.0.0.0 0.0.0.0 10.10.100.1” will configure a static default route that will be used when the router (or Layer 3 switch) does not have a route in its routing table.

e. Managing static routes is very efficient in large networks.

Answer: e

The first four options basically describe the features and functions of a static route, though static routes in large networks can quickly become resource intensive as they all have to be managed manually on each switch/router. Static routes can also be used to manage specific traffic patterns that you do not want to be changed by dynamic routing updates.

79. If a router receives a RIP (Routing Information Protocol) route update, an OSPF (Open Shortest Path First) route update, and an EIGRP (Enhanced Interior Gateway Routing Protocol) route all to the same destination, which route will it prioritize in its routing table?

a. The RIP route

b. The OSPF route

c. The EIGRP route

d. None of the above

Answer: c

The EIGRP route has the shortest Administrative Distance (90) as compared to OSPF (110) and RIP (120), and so it will be the preferred route used by the router. Other Administrative Distance's include:

Routing Process/Protocol	**Administrative Distance**
Directly connected	0
Static	1
EIGRP (summary route)	5
eBGP	20
EIGRP (internal)	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EIGRP (external)	170
iBGP	200

80. An IP packet is sent by a sender to a destination starting with a Time to Live value of 1. The first hop in the path to the destination reduces this Time to Live field by 1 and sends an ICMP Time Exceeded message back to the sender. The sender will next increment the Time to Live packet by 1 to 2 and sends the packet again. The first and the second hop each reduce the Time to Live field by 1 and the second router will send back an ICMP Time Exceeded message. This process continues with the last router sending back an ICMP Time Exceeded message until the destination is reached. The sender becomes aware of each hop along the path to the destination based on these ICMP Time Exceeded messages. What is the Cisco NX-OS command to implement this process from Switch_A to 192.168.200.100? (assume the use of the default VRF on the switch)

a. Switch_A# ping 192.168.200.100

b. Switch_A# tracert 192.168.200.100

c. Switch_A# tracepath 192.168.200.100

d. Switch_A# traceroute 192.168.200.100

Answer: d

The above scenario describes how the IP packet header's Time to Live field and ICMP (Internet Control Message Protocol) Time Exceeded responses are used to trace the route or track the hostnames and IP addresses along a path from a source to its destination. (Also, "tracert" = Windows, "tracepath" = UNIX/Linus, and "traceroute" = Cisco NX-OS CLI, all different commands on different systems for essentially the same process)

81. What OSI Layers does a Ping command test?

a. Layer 1

b. Layer 1 & 2

c. Layer 1, 2, & 3

d. Layer 1, 2, 3, & 4

e. Layer 1, 2, 3, 4, & 5

f. Layer 1, 2, 3, 4, 5, & 6

g. Layer 1, 2, 3, 4, 5, 6 & 7

h. None of the above

Answer: c

A Ping, or Packet Internet Groper, uses ICMP (Internet Control Message Protocol) to reach out to a network destination with an ICMP echo request. If the destination received the request, it will respond back with an ICMP echo reply. However, this process is independent of whatever applications are running on the source or destination devices. Pinging a destination can only confirm that the Physical layer, Data Link layer, and Network layer to that destination should be working.

82. Which three options below best describe TCP segment flow between a sender and a receiver device? (choose three)

a. The sending device sends a set of segments, starts a timer, and waits for an acknowledgment from the receiving device before sending more segments.

b. The sending device sends a set of segments, and keeps sending segments until the communication is over.

c. The receiving device receives segments from the sender and sends an acknowledgment back to the sender.

d. The receiving device receives segments from the sender.

e. The sending device sends a set of segments, starts a timer, if the timer expires without it receiving an acknowledgement from the receiving device, it will resend the segments.

Answer: a, c, e

When using TCP (Transmission Control Protocol) to communicate, the sender and the receiver devices stay in sync with each other by using Acknowledgements (ACKs), the sender must receive Acknowledgments from the receiver before it will send additional segments.

For instance, with a TCP window size of one, the sender sends a segment, the receiver receives it and sends an Acknowledgment (ACK), the sender receives the Acknowledgment and will now send the second segment. The receiver receives the second segment and sends an Acknowledgment, the sender receives the Acknowledgment and will now send the third segment. The receiver receives it and sends an Acknowledgment, the sender receives the Acknowledgment and will next send the fourth segment. UDP datagram flow (options “b” and “d” above) does not rely on first establishing a connection before sending data.

Example: Fixed TCP Window Size = 1

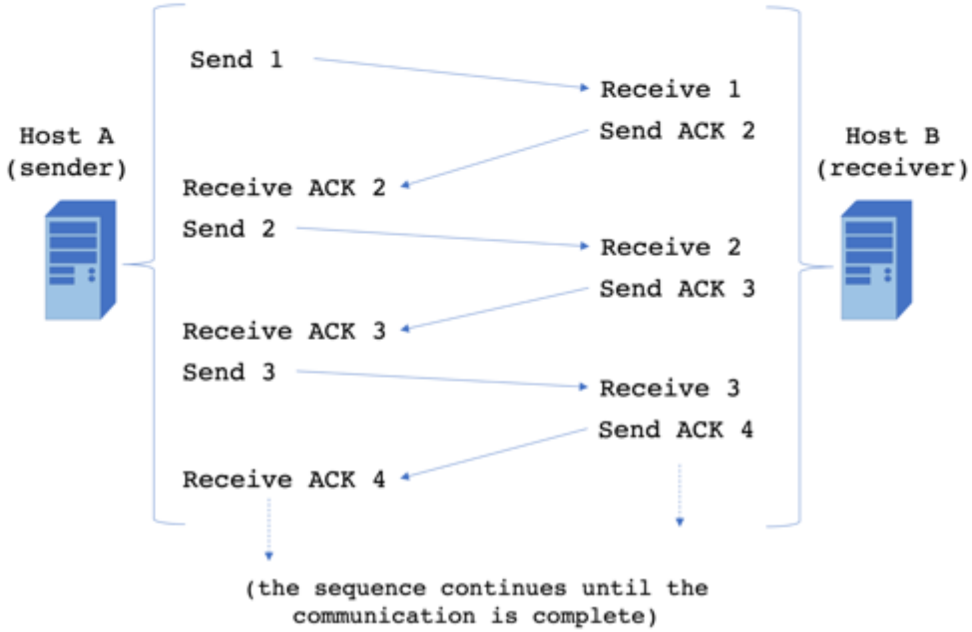

We can also illustrate the process with a fixed window size of three, the sender will send three segments and wait for an Acknowledgment (ACK) before it sends another three:

Example: Fixed TCP Window Size = 3

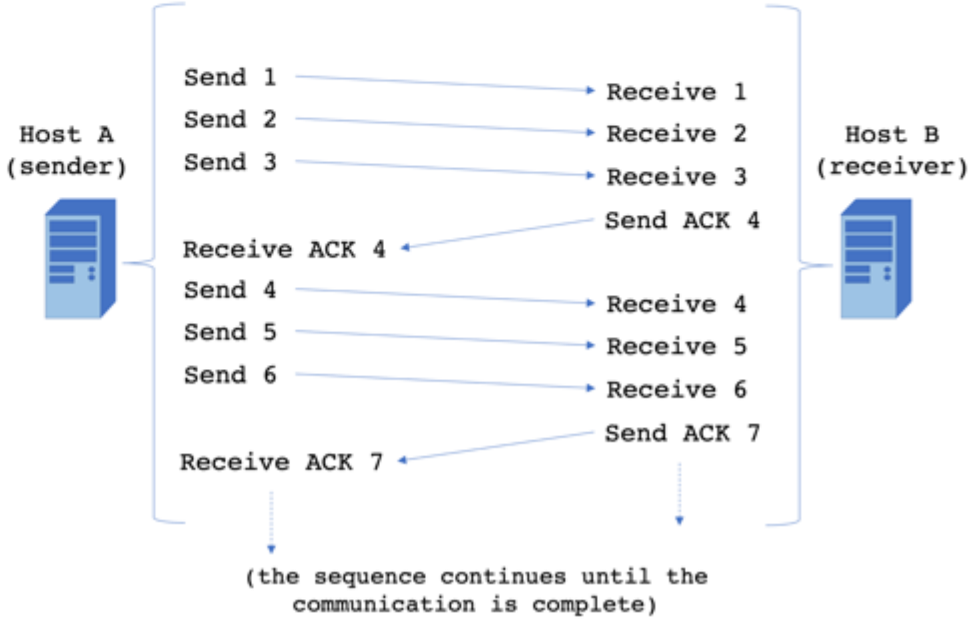

When using a dynamic or sliding TCP window, the sender will adjust the window size (or the number of segments it sends at once) based on window size requests from the receiver. For instance, if the receiver experiences bandwidth issues that cause it to miss one of the three segments that the sender sent, it can request a smaller window size, such as a window size of two. Next, the sender will resend the missing segment along with new segments using the smaller window:

Example: Sliding TCP Window

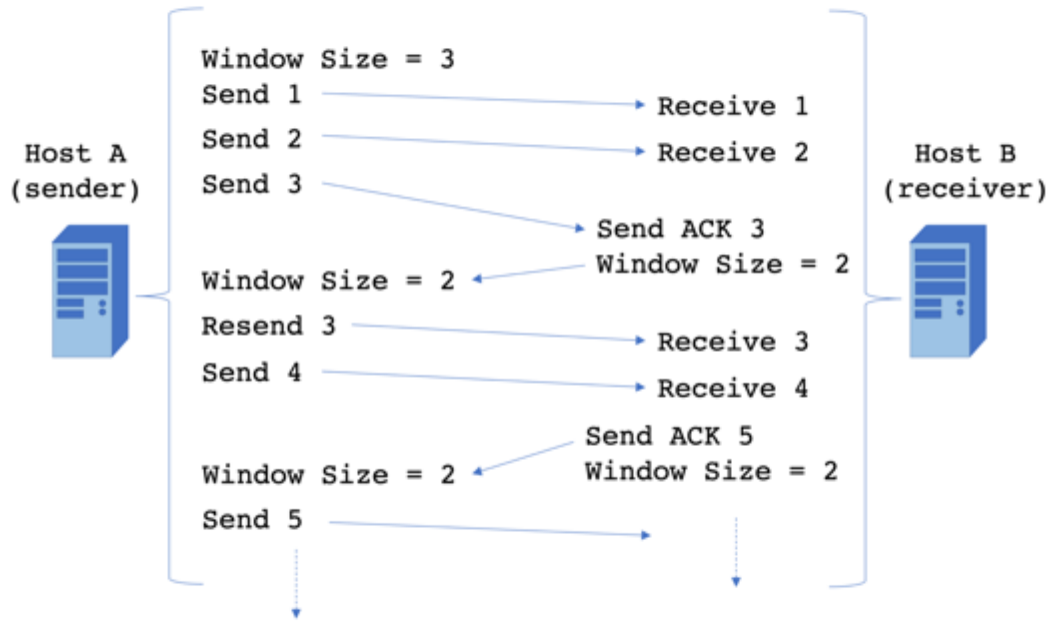

(the sequence continues until the communication is complete)

Note, the receiver's Acknowledgments (ACKs) sent to the sender are numbered based the segment that the receiver wants to receive next.

83. Can you match the SAN Storage Tier Level to the characteristics of that level?

Tier 0

Tier 1

Tier 2

Tier 3

Tier 4

A mixed SAN of disk and tape, typically used for data backup.

A mixed SAN of disk and tape, typically used for unstructured data, such as systems used for archival or data at rest.

A SAN of 10K to 15K RPM disks, may be hundreds of TBs of storage, structured and unstructured data, such as systems used for transactional data.

A large-scale SAN, able to replicate itself, 10K+ IOPS, mainly structured data, such as system used by online retailers.

A SAN of solid state drives, 100s of GBs to 10s of TBs in capacity, 100K+ IOPS, such as systems used in currency trading.

Answer:

In practice, our SANs and data centers may fall in between many of these "tiers", but this is how Cisco (at least as of the DCICN 200-150 and based on the standards of the Storage Networking Industry Association or SNIA) generalizes the scale, complexity, and cost of different tiers of SANs:

Tier 0 = A SAN of solid-state flash drives, 100GBs+ to 10 TBs+ in capacity, with over 100K+ IOPS, such as systems used in currency trading for critical data that must change quickly (considered an "ultra-high" performance system).

Tier 1 = A large-scale SAN, with enough capacity to replicate itself, with over 10K+ IOPS, mainly structured data, such as systems used by online retailers for critical data (considered a "high performance" performance system).

Tier 2 = A SAN of 10K to 15K RPM disks, may be hundreds of TBs of storage, structured and unstructured data, such as systems used for noncritical, operational/transactional data (considered a "medium" performance system).

Tier 3 = A mixed SAN of disk and tape, typically used for unstructured data, such as systems used for archival or data at rest (considered a "low" performance system).

Tier 4 = A mixed SAN of disk and tape, typically used for data backup,

(Note, IOPS = Input/Output Operations Per Second, it is a generalized means of measuring the throughput capabilities of a SAN environment)

For those drawing the lines, the correct tier to tier characteristics look like:

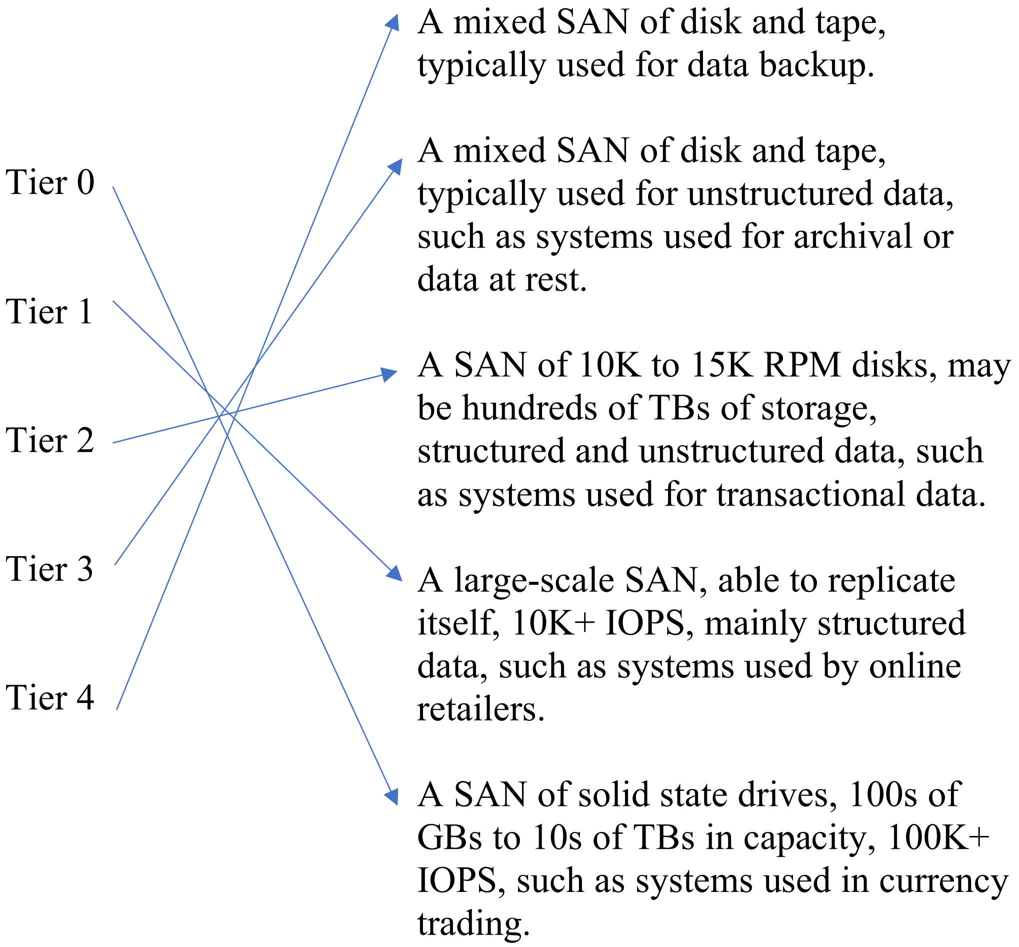

Tier 0

Tier 1

Tier 2

Tier 3

Tier 4

A mixed SAN of disk and tape, typically used for data backup.

A mixed SAN of disk and tape, typically used for unstructured data, such as systems used for archival or data at rest.

A SAN of 10K to 15K RPM disks, may be hundreds of TBs of storage, structured and unstructured data, such as systems used for transactional data.

A large-scale SAN, able to replicate itself, 10K+ IOPS, mainly structured data, such as systems used by online retailers.

A SAN of solid state drives, 100s of GBs to 10s of TBs in capacity, 100K+ IOPS, such as systems used in currency trading.

84. Which options below are “out-of-band” methods to connect to and manage a Nexus switch? (choose three)

a. RS-422 using the Console port

b. RS-232 using the Console port

c. SSH or Telnet using the Connectivity Management Port (CMP)

d. SHH or Telnet using the “mgmt0” Management port

e. SHH or Telnet on a production network connected to the switch

---

Answer: b, c, d

While “e” is also a means to connect to the CLI (Command Line Interface) of the switch, this is considered an “in band” method to connect to the switch because you are using the production network to reach the switch. Setting up a separate management network or VLAN for the sole purpose of managing your routers and switches would be considered an “out of band” method to connect to the CLI because you are not using the production network to reach the switch. Similarly, a local connection using a laptop or server to the Console port is also considered “out of band” because the production network is not used (and so any issues on the production network will not affect your ability to connect to the switch). Also, in practice you’ll want to use SSH (Secure Shell) not unsecured Telnet to connect to your switches. Note, the CMP on a Nexus has its own processor and gets power from the switch’s aux bus, so it can be used to rescue the switch if the console port is not working.

85. Which switchport mode below is typically used to connect a host to a switch? (choose the best answer)

a. Ethernet

b. Dynamic auto

c. Dynamic desirable

d. Access

e. Trunk

Answer: d

There are two modes of operation for a switchport interface: access and trunk. The access mode is used to connect a host server or workstation to a switch for network connectivity. The trunk mode is used to connect switches to other switches or to routers and to process frames for different VLANs over the same network connection.

However, there are different ways to create a trunk or access port; "dynamic auto" will auto-negotiate with the end device and become either an access port or a trunk. A "dynamic desirable" interface will configure itself to be a trunk as long as the interface on the other device is requesting to become a trunk. The specific "trunk" command, as in option "e", will manually create a trunk if "dynamic auto", "dynamic desirable", or "trunk" mode is used on the other interface in the connection.

86. Which option best describes Gateway Load Balancing Protocol (GLBP)?

a. A Cisco proprietary standard that uses a Forwarding Router and Standby Routers together to form a virtual router, one of the Standby Routers will take over if the Forwarding Router goes down.

b. An industry standard where a Master Router and Backup Routers work together to form a virtual router that uses the Master Router's IP address, a Backup Router will take over if anything happens to the Master Router.

c. A Cisco proprietary protocol that uses one router as an Active Virtual Gateway with a single IP address but multiple virtual MAC addresses for the other Active Virtual Forwarder routers, routing load is shared among all routers in the group.

d. An industry standard that uses one router as an Active Virtual Gateway with a single IP address but multiple virtual MAC addresses for the other Active Virtual Forwarder routers, routing load is shared among all routers in the group.

e. None of the above

Answer: c

Cisco's Gateway Load Balancing Protocol (GLBP) uses an Active Virtual Gateway (AVG) that coordinates multiple Active Virtual Forwarder (AVFs) routers, and routing load is shared among all routers in the group. This method is a more optimal method of utilizing router resources usage than VRRP or HSRP.

One router (or Layer 2 switch) is active, the other switch in this example is listening, however, both devices are working together to share the overall workload:

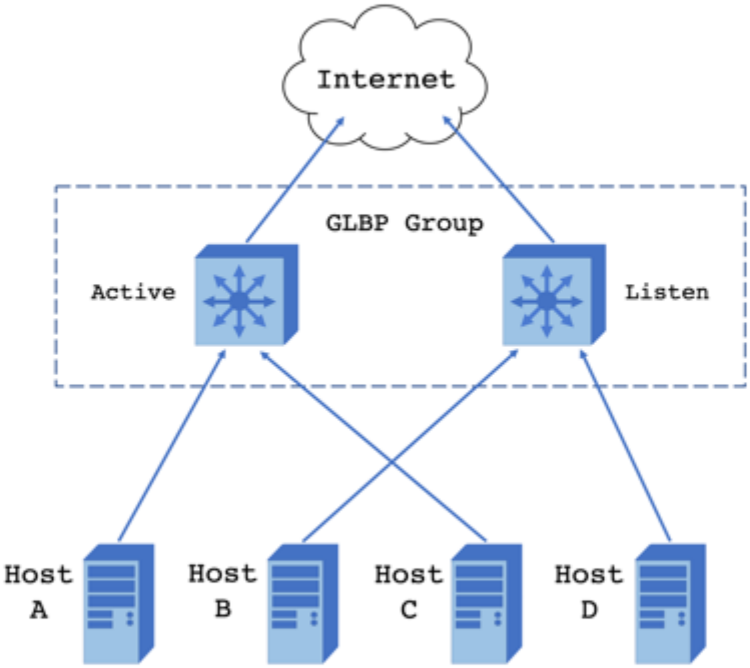

If anything happens to the active device, the listening device will take over:

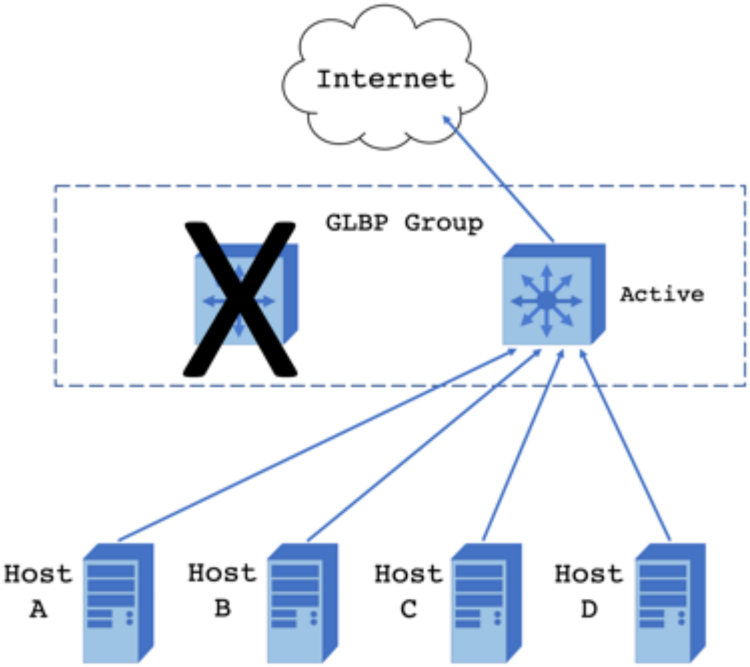

87. Which option best describes Virtual Router Redundancy Protocol (VRRP)?

a. A Cisco proprietary standard that uses a Forwarding Router and Standby Routers together to form a virtual router, one of the Standby Routers will take over if the Forwarding Router goes down.

b. An industry standard where a Master Router and Backup Routers work together to form a virtual router that uses the Master Router's IP address, a Backup Router will take over if anything happens to the Master Router.

c. A Cisco proprietary protocol that uses one router as an Active Virtual Gateway with a single IP address but multiple virtual MAC addresses for the other Active Virtual Forwarder routers, routing load is shared among all routers in the group.

d. An industry standard that uses one router as an Active Virtual Gateway with a single IP address but multiple virtual MAC addresses for the other Active Virtual Forwarder routers, routing load is shared among all routers in the group.

e. None of the above

Answer: b

Virtual Router Redundancy Protocol (VRRP) organizes a group of routers with one router assigned as the Master Router and the other routers assigned Virtual Backup Router status. The Master Router sends heartbeat messages to the Backups at specific intervals (every 1 second by default), a Backup will take over for the Master if three messages are missed.

88. Which option best describes Hot Standby Router Protocol (HSRP)?

a. A Cisco proprietary standard that uses a Forwarding Router and Standby Routers together to form a virtual router, one of the Standby Routers will take over if the Forwarding Router goes down.

b. An industry standard where a Master Router and Backup Routers work together to form a virtual router that uses the Master Router's IP address, a Backup Router will take over if anything happens to the Master Router.

c. A Cisco proprietary protocol that uses one router as an Active Virtual Gateway with a single IP address but multiple virtual MAC addresses for the other Active Virtual Forwarder routers, routing load is shared among all routers in the group.

d. An industry standard that uses one router as an Active Virtual Gateway with a single IP address but multiple virtual MAC addresses for the other Active Virtual Forwarder routers, routing load is shared among all routers in the group.

e. None of the above

Answer: a

Cisco's Hot Standby Router Protocol, or (HSRP) predates both VRRP and GLBP, and also used hello or heartbeat packets between the Forwarding Router and the Standby Routers (a Standby will take over for the Forwarding Router). Routers are given a default priority of 100. A router priority set to 90 (*the lower the number, the lower the HSRP priority*) will make it a Standby Router (and a router with a priority of 100 will become the Forwarding Router), similarly a router with its priority set higher to 110 will make it the active Forwarding Router (with other routers with the default priority of 100 becoming Standby Routers).

89. In a network running Spanning Tree Protocol (STP) there are three switches, each are given the default Bridge Priority of “32768”, Switch_A’s MAC address is 1234.56AB.CDEF,
Switch_B’s MAC address is 6543.21AB.CDEF,
Switch_C’s MAC address is ABCD.EF12.3456,
which switch will become the root bridge?

a. Switch_A

b. Switch_B

c. Switch_C

d. None of the above (active failover)

e. All of the above (load balancing)

Answer: a

A switch’s STP Bridge ID is a combination of its given Bridge Priority and its MAC address, so:

Switch_A’s Bridge ID becomes 32768.1234.56AB.CDEF,
Switch_B’s Bridge ID becomes 32768.6543.21AB.CDEF, and
Switch_C’s Bridge ID becomes 32768.ABCD.EF12.3456,

Switch A has the lowest alphanumerical Bridge ID, so it takes priority and become that VLAN’s STP root bridge (the root bridge is selected based on lowest bridge ID, so *a low BID = higher priority*). Possible STP Bridge IDs are 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, to 61440 in increments of 4096. So, if we wanted Switch_B or Switch_C to become the root bridge, we could give one of them a new bridge priority of 28672 or less.

90. What are four advantages of RSPT (802.1W) over CST/regular STP (802.1D)? (choose four)

a. VLAN/network changes communicated via BPDUs could take 60+ seconds to take effect on connected network switches.

b. VLAN/network changes communicated via BPDUs could take 6+ seconds to take effect on connected network switches.

c. VLAN/network changes communicated via BPDUs take could 1+ milliseconds to take effect on connected network switches.

d. BPDUs are sent from the root bridge to all other switches in that VLAN.

e. BPDUs are sent from any switch that recognizes a network change to all other switches in that VLAN.

f. An Alternate port is the start of the next best path to the root bridge and will take over when there is a failure to the root port.

g. A Backup port is the next best port and will take over when there is a failure to the designated port.

Answer: b, e, f, g

Rapid Spanning Tree Protocol (RSTP as defined by IEEE 802.1W), builds on the advantages of Common Spanning Tree (CST) or Spanning Tree Protocol (STP as defined by IEEE 802.1D), by allowing all switches to send BPDUs (Bridge Protocol Data Units) to each other. STP's listening, blocking, and disabled ports become RSTP's discarding ports (a port that does not forward BPDUs), and an alternate and a backup port designation is added. Options "a, b, and c" assume a 2 second default Hello Timer. Options "a" and "d" describe STP. Note, an RSTP alternate port will establish a path to the root bridge using another switch, and a RSTP backup port is the second best path on that switch to the root bridge.

91. Which options do not describe Multiple Spanning Tree Protocol, or MSTP (IEEE 802.1s)? (choose two)

a. Rather than an instance of STP or RSPT running for every VLAN on the network, MSTP will assign VLANs to specific MSTP instances.

b. STP and RSTP processes are automatically combined by root bridges into a single MSTP process.

c. MSTP is less switch CPU intensive

d. MSTP is more switch CPU intensive

Answer: b, d

A network or data center administrator has to manually create MSTP “instances”, VLANs are then assigned to these instances. MSTP can be helpful when STP or RSTP processes are similar across multiple VLANs on the network (e.g. when there are only a couple ways to reach resources or the Core routers, then the STP processes in a network may be very similar across multiple VLANs). So, switches can run a few MSTP instances for multiple VLANs, rather than many STP or RSTP instances (in large networks), one for every VLAN. This simplification strategy also means less work for the switches’ CPUs.

92. Which of these commands will assign the range of Ethernet interfaces 1/1 to 1/6 to VLAN 2 and Ethernet interfaces 1/7 to 1/12 to VLAN 3?

a.

```
Switch_A# configure terminal
Switch_A (config)# interface Ethernet 1/1 – 6
Switch_A (config-if-range)# access vlan 2
Switch_A (config-if-range)# interface Ethernet
1/7 – 12
Switch_A (config-if-range)# access vlan 3
```

b.

```
Switch_A# configure terminal
Switch_A (config)# interface Ethernet 1/1 – 6
Switch_A (config-if-range)# switchport trunk
vlan 2
Switch_A (config-if-range)# interface Ethernet
1/7 – 12
Switch_A (config-if-range)# switchport trunk
vlan 3
```

c.

```
Switch_A# configure terminal
Switch_A (config)# interface Ethernet 1/1 – 6
Switch_A (config-if-range)# switchport vlan 2
Switch_A (config-if-range)# interface Ethernet
1/7 – 12
Switch_A (config-if-range)# switchport vlan 3
```

d.

```
Switch_A# configure terminal
Switch_A (config)# interface Ethernet 1/1 – 6
Switch_A (config-if-range)# switchport access
vlan 2
Switch_A (config-if-range)# interface Ethernet
1/7 – 12
Switch_A (config-if-range)# switchport access
vlan 3
```

Answer: d

The correct syntax to add switch interfaces to a VLAN is:

```
Switch_A# configure terminal
Switch_A (config)# interface Ethernet [your
interface module/port or range of ports]
Switch_A (config-if-range)# switchport access
vlan [your vlan #]
Switch_A (config-if-range)# interface Ethernet
[your other interface module/port or range of
ports]
Switch_A (config-if-range)# switchport access
vlan [your other vlan #]
```

Note, in this example we're working with two separate VLANs (and don't forget to save your running-config when you are done)

93. In the diagram below, what is the purpose of the physical L1 and L2 connections? (choose two)

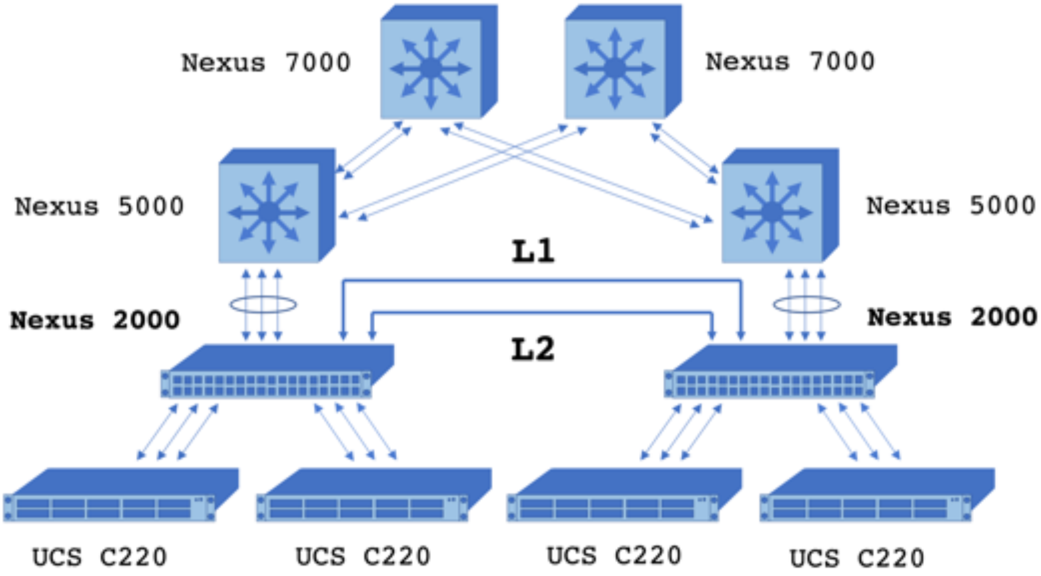

a. L1 is the Layer 1 connection between the two Nexus 2000s.

b. L2 is the Layer 2 connection between the two Nexus 2000s.

c. L1 is the primary cluster connection between the two Nexus 2000s

d. L2 is the secondary cluster connection between the two Nexus 2000s

e. L1 and L2 create a dual switchport port channel between the two Nexus 2000s

f. L1 and L2 create a dual ISL link between the two Nexus 2000s

Answer: c, d

The L1 and L2 ports typically found on Nexus 2000-series Fabric Extenders (FEXs) are not switchport interfaces (so options “e” and “f” are not correct). Also, Nexus Fabric Interconnects only extend the interfaces of its parent Nexus 3000, 5000, 7000, or 9000 series switch, they do not do any switching themselves. L1 and L2 allow for a primary and a secondary connection between the Nexus 2000s that provide for high availability and synchronization between each other.

Recall that the UCS Managers are installed on these, and so there is a primary and a secondary instance of the UCS Manager that is running on these FEXs. So, if the FEX with the active UCS Manager dies, the UCS Manager on the other FEX will take over. Also, please be sure that if you build a cluster of dual FEXs, that you use the same model Nexus 2000-series FEX that have the same number of ports on each FEX (for example, use two 48-ports or two 96-ports FEXs, don’t mix and match these in production environments!).

94. Which option best describes the FIP process?

a. First the FCoE VLAN Discovery, then the FCF Discovery, then the FLOGI,

b. First the FLOGI, then the FCoE VLAN Discovery, then the FCF Discovery,

c. First the CNA initiate, then the FC Command, then the FC Command Response,

d. First the initiator Sends, then the target ACKs, then the target Sends, then the initiator ACKs,

---

Answer: a

The Fibre Channel over Ethernet Initialization Protocol (FIP), begins with the with the Fibre Channel over Ethernet (FCoE) Discovery process where an initiator sends a multicast Request requesting the VLAN information from Fiber Channel Forwarders (FCF), or the Nexus Switch in the diagram. The FCF (the Nexus Switch) responds with VLAN information.

In the FCF Discovery phase, the initiator next asks, via a Solicitation multicast, the FCFs for the MAC address of the specific FCF that it wants to log into. The FCF responds with its MAC address and available VF-Ports.

The initiator next begins the FLOGI (Fabric Login) process, similar to the standard Fibre Channel FLOGI, and the FCF accepts the login. Now that the link is established, the initiator will send Fibre Channel commands and the FCF (which is connected to the Fibre Channel storage/target) will respond.

Example: FIP Virtual Link Establishment

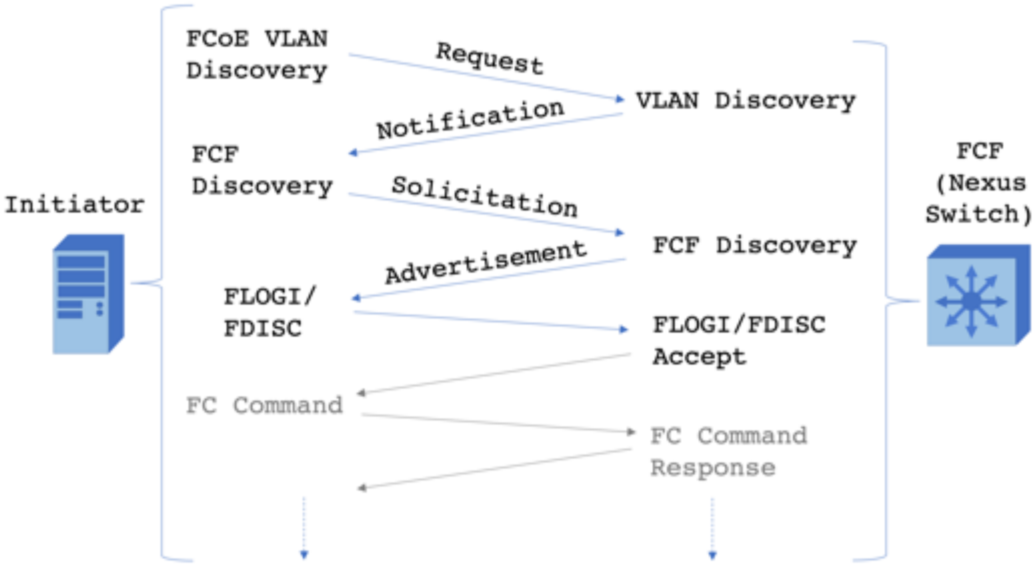

(the sequence continues until the FC communication is complete)

95. What is the relationship between DCBX and LLDP? (choose the best answer)

a. LLDP parameters are transported between CDP compliant devices within DCBX packets

b. DCBX parameters are transported between CDP compliant devices within LLDP packets

c. LLDP parameters are transported between DCB compliant devices within DCBX packets

d. DCBX parameters are transported between DCB compliant devices within LLDP packets

e. None of the above

Answer: d

There are a number of question examples in this book where I purposely drop in a lot of acronyms, this is done on purpose so that you get additional opportunities to become familiar/re-familiar with them. Sorry, this is one of those questions. So LLDP is Link Layer Discovery Protocol, CDP is Cisco Discovery Protocol, DCBX is Data Center Bridging Exchange Protocol, and DCB is Data Center Bridging.

Data Center Bridging Exchange (DCBX) Protocol exchanges Data Center Bridging (enhancements to Ethernet for data centers) parameters and capabilities between network devices. The DCBX information is delivered within Link Layer Discovery Protocol (LLDP) packets. Recall that Link Layer Discovery Protocol is an industry standard, while Cisco Discovery Protocol also communicates information about device capabilities, though only among Cisco devices.

DCBX information shared between devices could include peer discovery, Enhanced Transmission Selection (ETS) priority, Priority-based Flow Control (PFC), and Congestion Notification.

96. What is the difference between PFC and LFC? (choose two)

a. PFC parameters are transported within LFC packets

b. LFC parameters are transported within PFC packets

c. PFC will pause specific types of traffic

d. PFC will pause all traffic

e. LFC will pause specific types of traffic

f. LFC will pause all traffic

Answer: c, f

Priority Flow Control, or PFC (IEEE 802.1bb), is class of services-based; when there is network congestion, it will pause specific types or classes of traffic. Link-level Flow Control, or LFC (IEEE 802.3x), can only pause all traffic during periods of congestion. Priority Flow Control is one of the Ethernet enhancements included in Data Center Bridging and is used to make data centers more effective and efficient. With PFC, data centers will prioritize important traffic and pause lower priority traffic when there are bandwidth issues.

97. Which of the below options are not included in an STP Hello BPDU? (choose the best two answers)

a. Timer values

b. Sender's root cost

c. Sender BID

d. Root BID

e. Root MAC Address

f. Sender MAC Address

Answer: e, f

Switches send STP (Spanning Tree Protocol) Hello BPDUs (Bridge Protocol Data Units) via multicast, and so a sender's or a root bridge's MAC addresses are not sent as a specific field in a BPDU frame (however the sender's or root's MAC addresses can be derived from the sender's or root's BID). A BPDU has specific fields for Hello, Max Age, and Forward Delay timers, Sender's root cost (the STP cost between the switch sending the BPDU and the root bridge), sender BID (the Bridge ID of the switch sending the BPDU), and the root BID (the Bridge ID of the switch that the switch sending the BPDU thinks is the root bridge). Switches running STP/RSTP in a VLAN communicate with themselves and elect a root bridge based on these BPDU exchanges.

98. In a non-virtualized, Fibre Channel Storage Area Network, which port descriptions are not true? (choose two)

a. An N-Port is a connection from a Fibre Channel Node

b. An F-Port is a connection to a Fibre Channel Switch

c. An L-Port is a connection from a SAN to a LAN

d. An E-Port is an expansion port that connects Fibre Channel switches to each other

e. An S-Port is a connection to Fibre Channel Storage

Answer: c, e

There are no such things as L-Ports and S-Ports in a Fibre Channel (FC) SAN, but there are N, F, and E-Ports. N-Ports are on the Node side (host/initiator or storage/target) of the connection, F-Ports are on the FC switch side of that connection. E-Ports connect FC switches (via ISLs, or Inter-Switch Links):

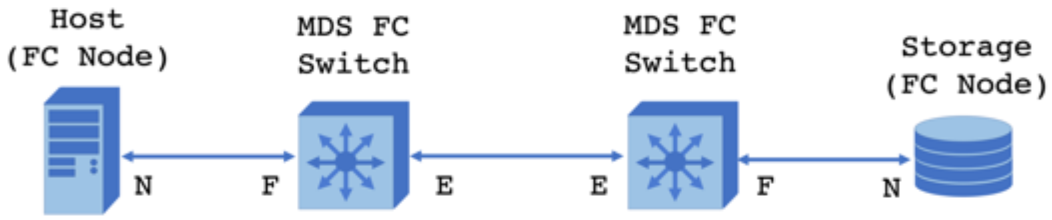

99. How does FSPF determine the best path between Fibre Channel switches?

a. The optimal path is based on the speed of each ISL over the path

b. The optimal path is based on the number of hops in the path

c. The optimal path is based on the speed of each ISL over the path and the number of hops in the path

d. None of the above

Answer: c

Fabric Shortest Path First (FSPF) is a routing protocol for Fibre Channel Storage Area Networks (and also for VSANs in that SAN). FSPF calculates the optimal path from one node to another based on the combined ISL connection speed of the switches between those nodes and the number of hops to the destination node.

100. Which option best describes the Fibre Channel write process?

a. The initiator sends a Write Request IU, the target sends a Ready IU, the initiator sends data IUs that are saved by the target, the target responds with an Acknowledgement IU that the write is complete,

b. The target sends a Write Request IU, the initiator sends a Ready IU, the target sends data IUs that are saved by the initiator, the initiator responds with an Acknowledgement IU that the write is complete,

c. The initiator sends a set of frames, starts a timer, and waits for an Acknowledgement from the receiving device before sending more frames,

d. The target sends a set of frames, starts a timer, and waits for an Acknowledgement from the receiving device before sending more frames,

---

Answer: a

In general, Fibre Channel initiators are workstations, virtual machines, or host servers who want to read from or write to the Fibre Channel target, or the storage device. Assuming the initiator has successfully FLOGI'd (a successful Fabric Login) and PLOGI'd (a successful Port Login), then it sends the Write Request Information Unit (IU), the target responds that it is ready, the initiators sends data, and the target sends an Acknowledgement when it has saved the data:

Fibre Channel Writes:

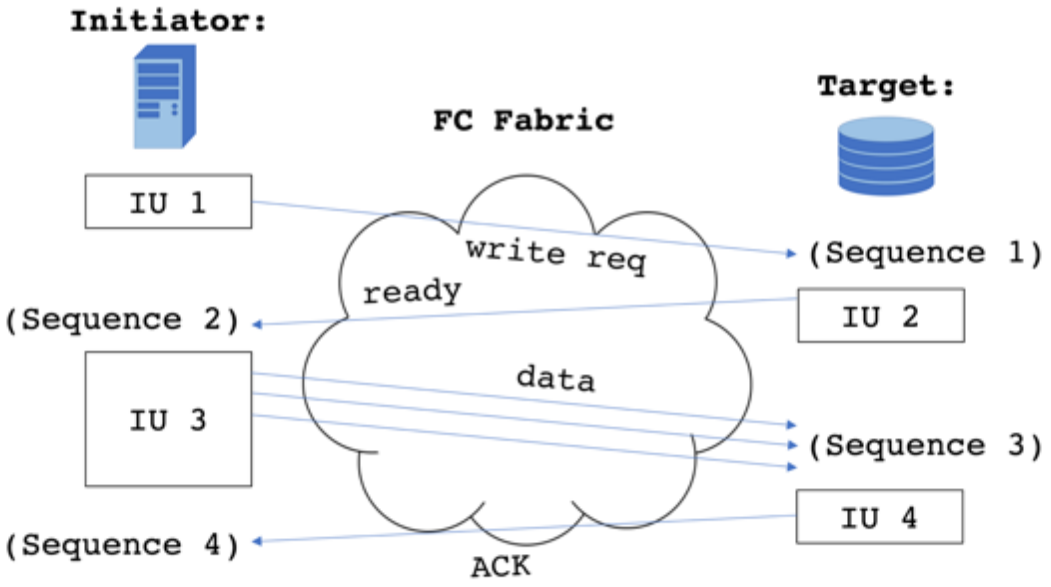

The Fibre Channel read process works similarly, with the initiator beginning the process with a Read Request IU, the target responds by sending data, and then lets the initiator know when the sending of data is complete:

Fibre Channel Reads:

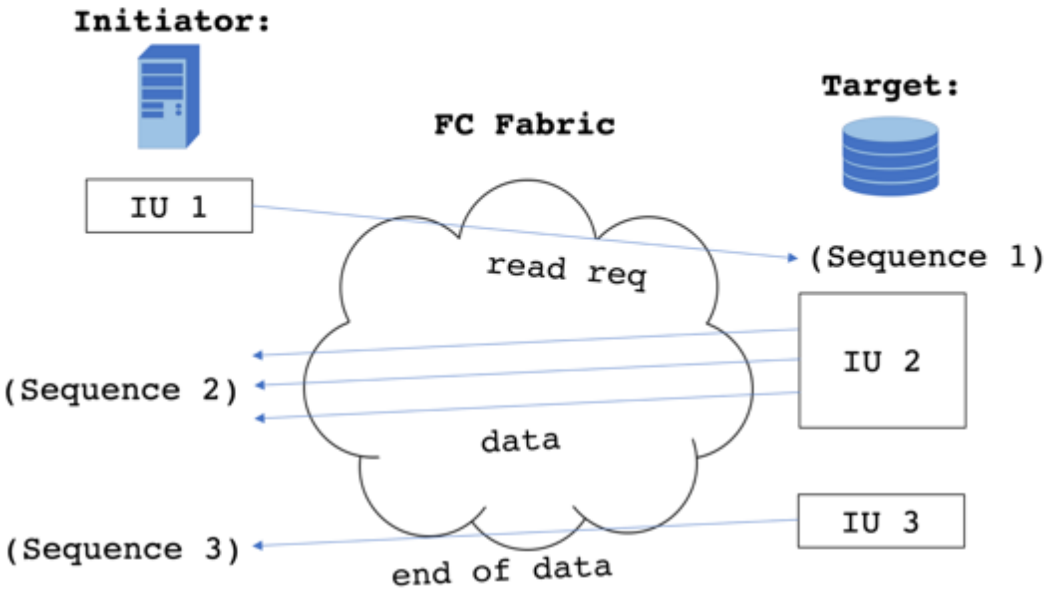

101. What are the default metrics used by EIGRP to determine optimal routes for packets?

a. Load

b. Delay

c. Reliability

d. Minimum Bandwidth

e. Maximum Transmission Unit

f. All of the above

Answer: b, d

Careful, while all of these options are Enhanced Interior Gateway Routing Protocol (EIGRP) metrics, by default only delay and minimum bandwidth are used to determine the best routes. Minimum bandwidth is measured along the path from the source router to the destination router (the link along the path with the lowest bandwidth is used in the calculation) while the delay is the time taken by the packet to reach its destination (the total of each delay across each link in the path).

102. Which options describe the operation of the switch after these commands are complete? (choose two)

```
Switch_A# configure terminal
Switch_A (config)# slot 1
Switch_A (config-slot)# port 1 type fc
Port type is changed.  Please reload the switch.
Switch_A (config-slot)# copy running-config
startup-config
Switch_A (config-slot)# reload
WARNING: This command will reboot the system
Do you want to Continue? (y/n) [n] y
<output omitted>
Switch_A# configure terminal
Switch_A (config)# interface fc 1/1
Switch_A (config-if)# switchport trunk allowed
VSAN 10
Switch_A (config-if)# exit
Switch_A (config)# exit
Switch_A# copy running-config startup-config
```

a. Only VSAN 10 is allowed on Ethernet interface 1/1

b. Only VSAN 10 is allowed on Fibre Channel interface 1/1

c. The startup configuration is saved as the new running configuration

d. The running configuration is saved as the new startup configuration

Answer: b, d

If you are reading the answer options too quickly without re-reading them, you might select options like “a” and “c” which are subtlety different from the real answers. Please be careful and when you go on to take the exam, be sure to read the question and your answers twice before moving to the next question. By default, all VSANs are allowed on a VSAN trunk port, the “switchport trunk allowed VSAN 10” command changes that and says that only VSAN 10 is allowed. You can confirm what interfaces allow what VSANs on your switch by using the “show vsan membership” command.

103. If the TE-Port (or EISL) mode on Switch_A is set to ON. To create a trunk with Switch_B, what should the TE-Port mode be set to? (choose two)

a. On

b. Access

c. Auto

d. Off

e. Dynamic

Answer: a, c

The administrative trunk mode of at least one switch in the pair needs to be On. A Trunking Expansion Port (TE-Port), or an Enhanced Inter-Switch Link (EISL), will be created between two switches if the administrative trunk mode of the other switch in the pair is set to Auto or is also set to On. In the specific context of Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) networks; TE-Ports allow traffic from multiple VSANs to cross between a single physical connection between FC or FCoE switches.

104. True or false? Referring to the results of the “show running-config aclmgr” below, will the staff member on a workstation connected to interface Ethernet 1/1 be able to access their favorite social media sites?

```
Switch_A# show running-config aclmgr
!Command: show running-config aclmgr
!Time: Tue Feb 16 23:49:59 2019
version 7.7
ip access-list no-more-telnet
     10 deny tcp any any eq 23
     20 permit ip any any
ip access-list no-websites
     10 deny tcp any any eq 80
     20 permit ip any any
ip access-list restrict-telnet
     10 deny tcp 192.168.200.0/24 any eq 23
     20 permit ip any any
interface Ethernet 1/1
     ip access-group no-websites in
interface Ethernet 1/2
     ip access-group restrict-telnet in
Interface Ethernet 1/3
     ip access-group no-more-telnet in
```

a. True

b. False

Answer: b (False)

In this example, the Switch_A interface Ethernet 1/1 has the inbound ACL “no-websites” applied on it. This ACL blocks TCP port 80 from any source to any destination, and so it will block Hypertext Transfer Protocol (HTTP) or most web traffic. The ACL has a hidden “deny all” at the end, so the “permit any any” was added at sequence line 20 to allow all other traffic not using TCP port 80 to pass through the interface.

105. Which command will update the ACL “restrict-telnet” on Switch_A to add a new rule as sequence number 5 to only allow telnet traffic from hosts on 192.168.100.0/24?

a.
```
Switch_A(config)# ip access-list
restrict-telnet
Switch_A(config-acl)# 5 allow tcp
192.168.100.0/24 any eq 23
```

b.
```
Switch_A(config)# ip access-group
restrict-telnet
Switch_A(config-acl)# 5 allow tcp
192.168.100.0/24 any eq 23
```

c.
```
Switch_A(config)# ip access-list
restrict-telnet
Switch_A(config-acl)# permit tcp
192.168.100.0/24 any eq 23 5
```

d.
```
Switch_A(config)# ip access-list
restrict-telnet
Switch_A(config-acl)# 5 permit tcp
192.168.100.0/24 any eq 23
```

Answer: d

In the Nexus Operating System (NX-OS), you can insert new Access Control List (ACL) rules before and after existing sequence numbers (no need for editing your ACLs in Notepad like the old IOS days), just start the new rule with the new sequence number. The correct syntax to allow port 23 traffic is “ip access-list [ACL name]” and “[sequence #] permit tcp 192.168.100.0/24 any eq 23”.

106. Which option does not describe the characteristics of ACLs?

a. Execution of an ACL starts with the first sequence number and continues down the list.

b. Every ACL needs at least one "permit" statement.

c. Outbound ACLs tend to be more efficient than Inbound ACLs.

d. The last statement automatically added to an ACL is "deny all".

e. The NX-OS only permits one ACL per interface, per protocol, per direction.

Answer: c

It is best practice to apply our ACLs as close to the source of the traffic that we want to block as we can. In general, a switch will have to process a packet from an inbound interface to an outbound interface before the packet is processed by that outbound interface's ACL. Thus, the packet makes it all the way through the switch before it is dropped. Inbound ACLs generally require less processing because the processing occurs at the inbound interface, if the packets fails any ACL rule it is discarded and does not require further processing. Thus, the packet is dropped before it is processed all the way through the switch.

Access Control Lists (ACLs) start processing at the start of the list, a packet is analyzed line by line, it is discarded as soon as it fails a line in the ACL. There is an implicit “deny all” statement at the end of all ACLs, and so if any ACL does not include at least one “permit” statement, then that ACL will deny all traffic and render the interface it is applied to unusable. Also, we can only apply one ACL per interface, per protocol (for instance, IPv4 or IPv6), per direction (inbound or outbound).

107. Can you match the Nexus AAA security definitions to the terms? (not all terms will be used)

The verification of a Nexus user or host server process

The predetermination and granting of user rights and privileges once Nexus users successfully sign-on

Recording and documenting access attempts on the Nexus switch

Auditing

Accounting

Authentication

Authorization

Administration

Answer:

The verification of a user or server process = Authentication

The setup and determination of user rights and privileges once they successfully sign-on = Authorization

Recording and documenting access attempts = Accounting

The AAA in AAA is Authentication (authorized users can sign in, others cannot), Authorization (once users sign in, they next get the privileges that the administrator has given them), and Accounting (if needed, an administrator has documentation that can help trace what user has made what changes).

And so connecting the lines we have:

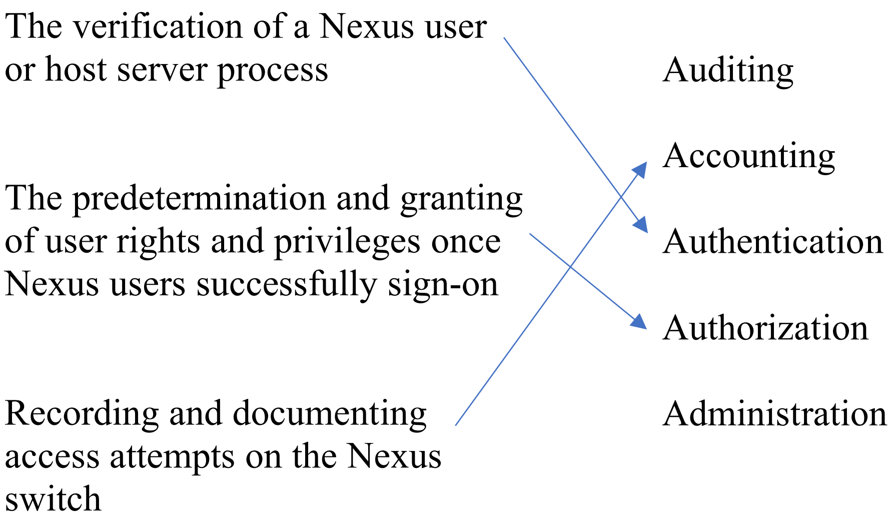

The verification of a Nexus user or host server process

The predetermination and granting of user rights and privileges once Nexus users successfully sign-on

Recording and documenting access attempts on the Nexus switch

Auditing

Accounting

Authentication

Authorization

Administration

108. What are the default NX-OS user roles on a Nexus 7000 series switch? (choose four)

a. Super-admin
b. Network-admin
c. Network-operator
d. Network-user
e. VDC-admin
f. VDC-operator
g. all of the above

Answer: b, c, e, f

256 user accounts can be defined, users are given roles, users can have multiple roles, and rules define what roles can do. Network-admins can create custom roles and assign up to 256 rules to each role. A VDC or, Virtual Device Context, is a means of dividing a physical switch into virtual switches. As such, each VDC on the switch can also have VDC-admin and VDC-operator roles. Note, for a Nexus 5000, the answer would be “b” & “c”, the 5000 series does not offer VDC at the time of this writing.

109. What protocols can be used for AAA services on NX-OS devices? (choose three)

a. RADON

b. RADIUS

c. TACACS+

d. PAP

e. LDAP

f. EAP

Answer: b, c, e

Options “d” and “f” are also authentication protocols, but they are not specifically supported by NX-OS switches. Your options in terms of external server clusters that can be used to manage the Nexus servers in your data center are RADIUS, TACACS+, and LDAP.

110. Switch_ A is currently running the default SSH configuration, which set of NX-OS CLI commands will reset the SSH key from its default 1024-bit encryption to a more secure 2048-bit encryption key?

a.
```
Switch_A# configure terminal
Switch_A(config)# feature ssh
Switch_A(config)# ssh key rsa 2048 force
Switch_A(config)# exit
```

b.
```
Switch_A# configure terminal
Switch_A(config)# no option ssh
Switch_A(config)# ssh key rsa 2048 force
Switch_A(config)# option ssh
```

c.
```
Switch_A# feature ssh
Switch_A# ssh key rsa 2048 force
Switch_A# copy running-config starting-config
```

d.
```
Switch_A# configure terminal
Switch_A(config)# no feature ssh
Switch_A(config)# ssh key rsa 2048 force
Switch_A(config)# feature ssh
```

Answer: d

This question assumes that the SSH (Secure Shell) feature is already enabled, and so first it has to be disabled from global configuration mode. Next, the SSH key is reset to 2048 (the strongest encryption available on NX-OS as of this writing), by “force” (otherwise you’ll get an error message that the key is already set to 1024), and finally you re-enable the feature SSH. The encryption being used is RSA, or Rivest, Shamir, and Adelman, the NX-OS can also run DSA or Digital System Algorithm encryption.

111. In a GLBP router group, what happens when a host ARPs the virtual router?

a. The Virtual GLBP router will respond to ARP requests

b. The listening router will respond to ARP requests

c. ARP requests sent to the router will be dropped

d. The active router will respond to ARP requests

e. The router that is the least busy (load balancing is in place) will respond to ARP requests

f. None of the above

Answer: d

The router in the active state (the Active Virtual Gateway, or AVG) in the GLBP (Gateway Load Balancing Protocol) group will respond to the ARP (Address Resolution Protocol) request from connected devices. If anything happens to take the AVG router offline, one of the remaining routers in the listen state (the Active Virtual Forwarding routers) will take over.

112. In the OSPF protocol, what are the purpose of areas?

a. To provide a minimal level of network security

b. To interconnect Leaf and Spine switches

c. To subdivide networks to restrict access to specific network resources

d. To subdivide storage resources to restrict access to certain files

e. To reduce the size of its link-state database

f. None of the above

Answer: e

In large OSPF networks, the link-state databases of each router can become very large and take up storage and processing resources. Breaking up large networks into OSPF areas helps optimize resources. Routers know how to route to other OSPF areas but they only have detailed knowledge about their own network areas, this reduces the size of their databases. Note, OSPF areas operate independently from VLANs and VSANs.

113. After reviewing the NX-OS “show ip route” command below, what options are true? (choose two)

```
Switch_A# show ip route
IP Route Table for VRF “default”
‘*’ denotes best ucast next-hop
‘**’ denotes best mcast next-hop
‘[x,y]’ denotes [preference/metric]
‘%<string>’ In via output denotes VRF <string>
<output omitted>
10.16.223.0/24, ubest/mbest: 1/0,
    *via 10.16.225.1, Ethernet 1/2, [110/50],
    504:57:30, ospf-2, intra
<output omitted>
```

a. The administrative distance to 10.16.223.0 is 110

b. The administrative distance to 10.16.223.0 is 50

c. The OSPF process ID is “2”

d. The OSPF area name is “2”

Answer: a, c

In this example, the administrative distance from the Layer 3 Switch_A to the 10.16.223.0 network is 110 (the administrative distance or priority of OSPF is 110). The metric in this example is 50 (a value calculated by the Dijkstra shortest path first algorithm based on the bandwidth and connections between Switch_A and the router connected to 10.16.223.0). The process ID, named during the implementation of OSPF on this switch, is “2” (this gets into the CCNP and CCIE realm, but for now assume the switch may be running multiple OSPF processes).

114. Which routing protocols are examples of Interior Gateway Protocols? (choose three)

a. RIP

b. BGP

c. OSPF

d. EIGRP

e. IDRP

f. EGP

Answer: a, c, d

RIP (Routing Information Protocol), OSPF (Open Shortest Path First), and EIGRP (Enhanced Interior Gateway Routing Protocol) are all Interior Gateway Protocols, or are routing protocols that operate within autonomous networks (or networks with a common administrator, like universities, companies, or agencies). Meanwhile, Exterior Gateway Protocols, such as BGP (Border Gateway Protocol), IDRP (Inter-Domain Routing Protocol), and EGP (Exterior Gateway Protocol) are used out on the Internet by Internet Service Providers (for routes between autonomous networks).

115. How do routers, or Layer 3 switches, running OSPF update each other?

a. Keep-alive packets sent on 224.255.255.255

b. Keep-alive frames sent on 223.255.255.255

c. LLDP packets sent on 224.0.0.1

d. Hello frames sent on 224.0.0.6

e. Hello packets sent on 224.0.0.5

f. None of the above

Answer: e

Routers running the Open Shortest Path First (OSPF) protocol send hello packets that discover other OSPF networks in the area, confirm configurations, and monitor the status of other routers to recalculate routes as needed. Hello packets are only sent and received on OSPF routers in the same area (also, note OSPF's use of a multicast address for its hello packets).

116. By default, when will a router running RIP send route information to other RIP routers? (only choose one)

a. Every 5 minutes it will send route changes and updates

b. Every 5 minutes it will send its full routing table

c. Every 30 seconds it will send route changes and updates

d. Every 30 seconds it will send its full routing table

e. Whenever there is a route change to its routing table

f. Periodic hello messages are not sent on a preset time interval

Answer: d

Routers (or Layer 3 Nexus switches) running RIP will send their full routing table out to its neighboring RIP routers about every 30 seconds (in practice they update each other every 25 to 35 seconds based on an algorithm such that all your routers are not synchronized to all update each other exactly every 30 seconds and thus congest your network every 30 seconds). Unlike other newer protocols, RIP will send its complete routing table not just incremental changes. For instance, EIGRP sends update messages only as needed when there are router topology changes and uses much smaller hello messages to let each other know that they are still online.

117. How does a router (or a Layer 3 Nexus switch) running OSPF send updates to other OSPF network devices?

a. It sends LSAs

b. It sends LSDBs

c. It sends its routing table

d. It sends LLDPs

e. None of the above

Answer: a

OSPF routers will send Link State Announcements (LSAs) as needed to update themselves when there are network topology and link state (status and speed) changes. LSAs are based on information from the router's Link State Database (LSDB), but it does not send its entire database the way RIP would send its entire routing table.

118. Router_ A is running RIP, OSPF, and EIGRP. There is an existing connected route, a new RIP update, a new OSPF LSA, and a new EIGRP hello packet announcing a route to Router_B. Which route will Router_A add to its routing table? (Assume that both Router_A and Router_B have been operating in production uninterrupted at least for several days)

a. The RIP route

b. The OSPF route

c. The EIGRP route

d. The connected route

e. None of the above

Answer: e

Administrative distance (AD) will define which routes a router will save to its routing table, the lower the distance the higher the priority of the route. For instance, another router directly connected to Router_A will be given an administrative distance of 0 and be added to Router_A's routing table (also, static routes manually added to a routing table are given the administrative distance of 1). An EIGRP route has an AD of 90, an OSPF route has an AD of 110, and a RIP route has an AD of 120. In this example, if Router_A has been in operation for some amount of time and received these router updates, it would know it is already physically connected to Router_B, and it will not update its routing table. So, this question may be a bit tricky but option ”e” is the best answer because Switch_A and Switch_B are directly connected to each other and this info has already been in the routing table.

119. What is commonly referred to as the Split-Horizon technique?

a. When switch information is prevented from leaving the interface where that information was received

b. When switch updates are sent with unreachable hop counts

c. When router updates are sent with unreachable hop counts

d. When router information is prevented from leaving an interface where that route information was received

e. None of the above

Answer: d

In the case of distance vector protocols, like RIPv1 and RIPv2, there are two techniques that network or data center admins can use to prevent router loops. The Split-Horizon route advertisement technique was first proposed by Torsten Cegrell in 1974, and tells a router not to advertise routes back down the interfaces where it originally received those route advertisements. So, if a router received an advertisement about Network X from port N, it would not send an advertisement about Network X back down port N.

Another distance vector technique that can be used to prevent routing loops is Route Poisoning. Route Poisoning is when a network administrator manually sets a hop count to a network to be more than the router’s maximum hop count threshold, and so the router will treat that network as unreachable (option “c” above). These two techniques will help prevent route loops.

Split-Horizon Example:
We start with three Layer 3 switches (so these switches are also routing between networks). Switch A is just added, so B has sent a RIP update (its full table) to Switch C.

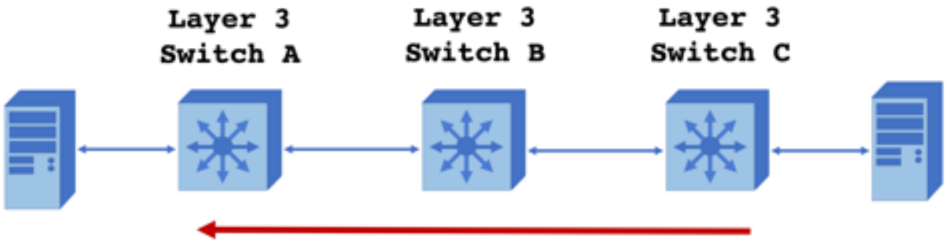

1. Switch C now knows that to reach switch A it has to first send packets to switch B, and so Switch C tells Switch B about a route from Switch C to reach Switch A. So Switch B now thinks it has two routes to get to A: 1) from B to A, and 2) from C to A

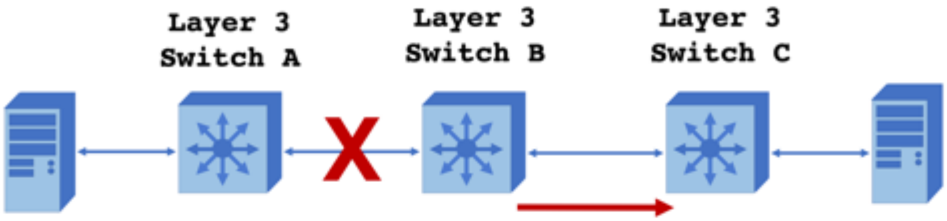

2. If anything happens to the connection between B and A, and if B needs to send packets to A, it will try to use its second route to get to A which is to go to C…

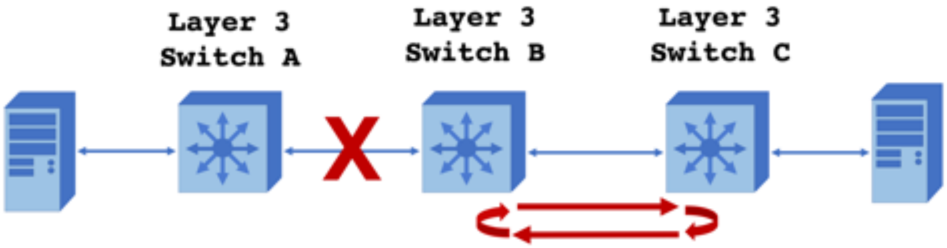

3. Next, C will see that it has packets to get to A, it knows to reach A it has to go to B first, so its sends those packets right back to B, creating a loop.

To avoid all of this, we tell C not to tell B about C's route to A, or we tell the Layer 3 switches not to send route updates to the same interface it received them.

120. Which option will monitor for broadcast traffic and drop broadcast traffic if that traffic occupies more than 20% of an interface's available bandwidth?

a.
```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/1
Switch_A(config-if)# storm-control
broadcast 20
```

b.
```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/1
Switch_A(config-if)# storm-control
broadcast 80
```

c.
```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/1
Switch_A(config-if)# traffic-storm
broadcast 20
```

d.
```
Switch_A# configure terminal
Switch_A(config)# interface Ethernet 1/1
Switch_A(config-if)# traffic-suppression
broadcast 80
```

Answer: a

To help prevent traffic storms, the "storm-control" option on a switch interface can monitor broadcast, multicast, or unicast traffic on that interface in 10 millisecond (ms) intervals and will drop traffic when that traffic gets above a pre-determined threshold. In this example, if broadcast frames take up more than 20% of the interface's available bandwidth over a period of 10 ms, then the switch will begin dropping broadcast frames and re-evaluate traffic after another 10 ms.

121.

121. What is the abbreviated version of the IPv6 address "2001:0db8:0000:0000:0000:abcd:1234:0001" (what is the best answer)?

a. 21:db8::::abcd:1234:1

b. 2001:0db8:0000::abcd:1234:0001

c. 2001:fe:ff:abcd:1234:1

d. 2001:db8::abcd:1234:1

e. 21:db8::abcd:1234:1

Answer: d

IPv6 addresses are long, while machines can deal with them just fine, for us humans it helps to shorten them a bit. And so, we can replace single groups of consecutive zeros with a double colon "::", but we can only do that once per IPv6 address. We can also eliminate leading zeros in a block, but not zeros in the middle of a block (so we can't do option "a" or "e" above).

122. What is the EUI-64 IPv6 autoconfiguration IP address for a device with the MAC address 1234.abcd.0001?

a. 0000.0000.0000.0000.0000.1234.abcd.0001

b. 1234.abcd.0001.0000.0000.0000.0000.0000

c. 2002:0db8:0000:0001:1034:abcd:efgh:0001

d. 2001:0db8:0000:0001:1234:acff:fecd:0001

e. 2002:0db8:0000:0000:0000:1234:abcd:0001

f. 2001:0db8:0000:0001:1034:abff:fecd:0001

Answer: f

This is another area where the DCICN overlaps with the CCENT. A system using EUI-64 (IEEE's 64-bit Extended Unique Identifier) for autoconfiguration of an IP version 6 (IPv6) address will start with the MAC address of 1234.abcd.0001, then…

Step 1: 1234:ab|cd:0001, where 1234:ab = the MAC OUI (Organizationally Unique Identifier) and cd:0001 = unique address,

Step 2: Insert “ff:fe” in between the OUI and the unique address = 1234:abff:fecd:0001

Step 3: Next, look at the 2nd octet: 1**2**34:abff:fecd:0001

Step 4: The “2“ as an octet in bits is “00000**1**0”, and after inverting the 7th bit in the address that “1” becomes “00000**0**0” or = “0”,

Finally, the IPv6 according to EUI-64 stateless autoconfiguration is: 2001:0bd8:0000:0001:1**0**34:abff:fecd:0001

123. Can you match the IP version 6 ranges to the IP addresses?

Global unicast range	0:0:0:0:0:0:192.168.100.100
Unique local range	2001:DB8::/32
Link-Local range	2000::/3
Multicast range	FF00::/8
An example of an IPv4 address converted to IPV6	FC00::/7
Reserved for examples	FE80::/10

Answer:

Global (publically routable IPv6 addresses) unicast range = 2000::/3 (typically you'll see these start with "2002:…" and eventually "2003:…")

Unique local (analogous to IPv4 private addresses) range = FC00::/7

Link-local (automatic private IPv6 addresses) unicast range = FE80::/10

Multicast range (analogous to IPv4 multicast addresses) = FF00::/8

An example of an IPv4 address (192.168.100.100) converted into an IPv6 address = 0:0:0:0:0:0:192.168.100.100

Reserved for documentation and training examples = 2001:DB8::/32 (3FFF:FFFF::/32 is also reserved for documentation and training examples)

And so, for those drawing in the lines we have:

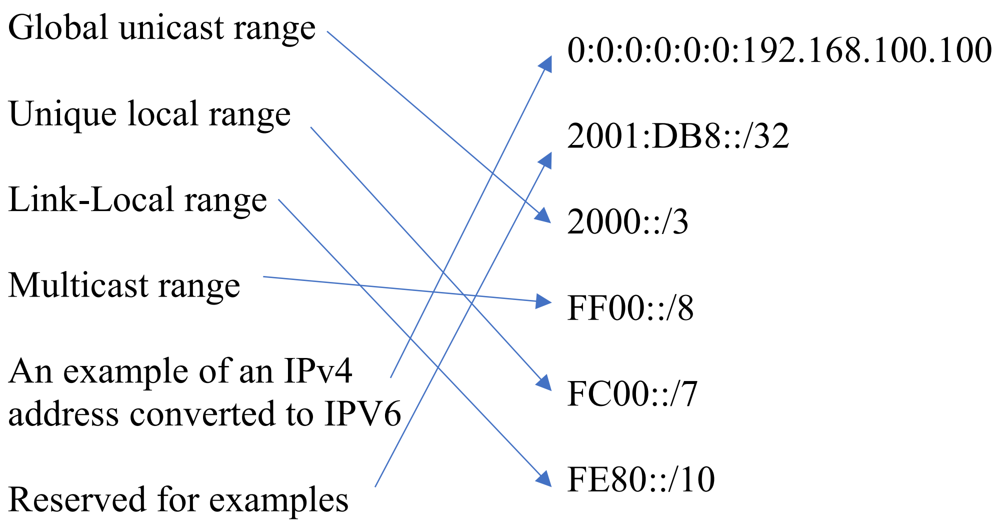

Global unicast range

Unique local range

Link-Local range

Multicast range

An example of an IPv4
address converted to IPV6

Reserved for examples

0:0:0:0:0:0:192.168.100.100

2001:DB8::/32

2000::/3

FF00::/8

FC00::/7

FE80::/10

124. What method does Fibre Channel use to ensure frames are not lost across links?

a. TCP Resend requests

b. Buffer to buffer credits

c. The bandwidth of a typical FC link exceeds those of Ethernet, eliminating the need for forward error correction

d. Data Center Congestion Notification

e. Enhanced Transmission Selection

f. None of the above

Answer: b

The buffer credit is used to define how much data the sending device can send at one time. Hop by hop, Fibre Channel devices communicate the size and status of their receive buffers; the sending device tracks this information as “credits”. The sending device knows the status of the receiving device’s buffer credits, and it will not send additional frames until it confirms that the receiver is ready. The use of buffer credits prevents FC devices from sending too many frames at once and from sending frames before the receiving device is ready.

125. What option best describes the IPv6 DHCP process?

a. Client sends a Discover ➔ DHCP server sends an Offer ➔ Client sends a Request ➔ DHCP server sends an Acknowledgment

b. Client sends a Solicit ➔ DHCP server sends an Advertise ➔ Client sends a Request ➔ DHCP server sends a Reply

c. Client sends a Request ➔ DHCP server sends an Discover ➔ Client sends an Accept ➔ DHCP server sends an Acknowledgment

d. Client sends an Advertise ➔ DHCP server sends an Discover ➔ Client sends an Accept ➔ DHCP server sends an Acknowledgment

e. None of the above

Answer: b

The stateful Dynamic Host Configuration Protocol (DHCP) process that an IPv6 client and an IPv6 DHCP server use to lease/give the client an IP address is very similar to the IPv4 process (which is option "a" above). For IPv4 DHCP remember: Discover, Offer, Request, and ACK. For IPv6 DHCP remember: Solicit, Advertise, Request, and Reply.

126. When connecting your laptop to a Nexus switch’s console port and using a Terminal Emulator program, what are the DCICN preferred serial settings you should use?

a. Speed: 96000 bps
   Data bits: 8
   Parity: none
   Stop bits: 1
   Flow control: none

b. Speed: 19200 bps
   Data bits: 8
   Parity: none
   Stop bits: none
   Flow control: none

c. Speed: 4800 bps
   Data bits: 8
   Parity: none
   Stop bits: none
   Flow control: none

d. Speed: 9600 bps
   Data bits: 8
   Parity: none
   Stop bits: 1
   Flow control: none

e. None of the above

Answer: d

Assume you have a USB to 9-pin converter and the blue rollover cable with the DB9 to RJ45 connector (that should come with your switch) to connect your laptop to the switch’s console port. Then use a Terminal Emulator program (such as free and open source PuTTY), the serial settings in option “d”, and select an available Com port on your laptop to open the connection.

127. We’re down to the last question! What options below are Fibre Channel layers?

a. Physical

b. Transmission Protocol

c. Signaling Protocol

d. Common Services

e. Network Services

f. Upper-layer Protocol Mapping

g. Application

Answer: a, b, c, d, f

Similar to the TCP/IP and the OSI models, a multi-layer process or model can be used to describe the operation of Fibre Channel. The Physical layer is similar to the same layer in the OSI model, this is the physical connection the network and the sending of FC frames via fiber or copper cables. The encoding and decoding of signals occurs at the Transmission Protocol layer, this process involves the inclusion of start of frame and end of frame information. The Signaling Protocol layer indicates where the port to port communication and flow control (using buffer to buffer credits or BB_credits) takes place. The work of World Wide Names (WWNs) and Fibre Channel IDs (FCIDs) for device identification happens at the Generic or Common Services layer. Finally, the Upper-Layer Protocol Mapping layer processes application specific procedures such as the mapping or conversion of Fibre Channel Protocols to and from SCSI commands:

The Five Layer Fibre Channel Model:

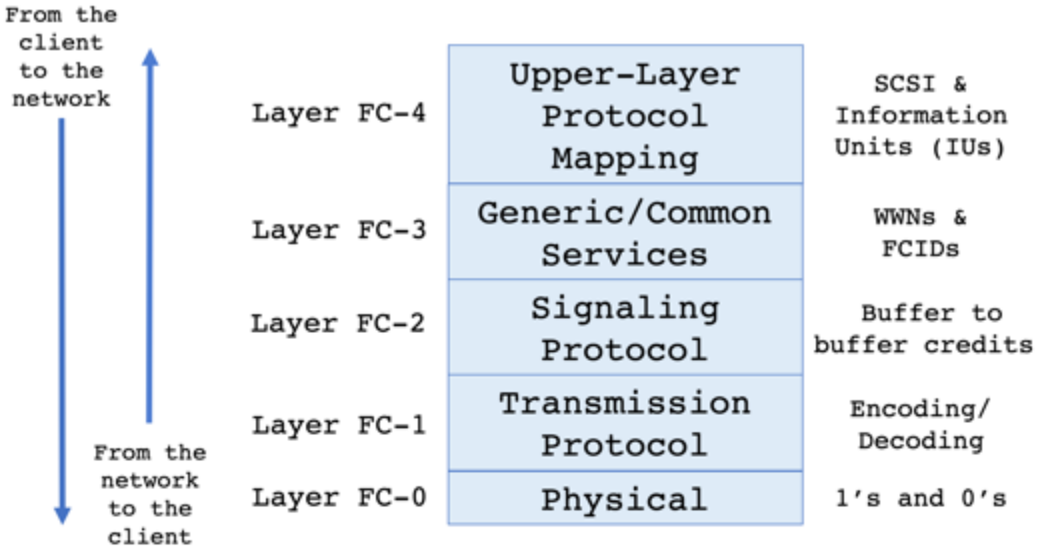

From the client to the network
From the network to the client

Layer FC-4	Upper-Layer Protocol Mapping	SCSI & Information Units (IUs)
Layer FC-3	Generic/Common Services	WWNs & FCIDs
Layer FC-2	Signaling Protocol	Buffer to buffer credits
Layer FC-1	Transmission Protocol	Encoding/ Decoding
Layer FC-0	Physical	1's and 0's

Next Steps:

Congratulations for making it all the way through this book, I hope you've found it helpful! This guide was not meant to include everything you will see on the Data Center Networking DCICN exam, so I also hope you've taken the time to write notes in the margins (or type notes if you are on the e-book version), and that you now have a good idea of what topics to go back and study or research further. Again, the official Cisco guides and online practice texts are great next steps. If you are working towards your CCNA Data Center certification, then the Data Center Technologies DCICT (200-155) exam is hopefully another next step. Keep working hard, keep studying, and never stop learning…

The official Cisco DCICN book and practice exams are great resources, but this is not an easy exam. This study guide is a companion to those resources and summarizes the subject areas into additional review questions with an answer description for each item. This book is not a "braindump" and it is not bootleg screenshots of the actual exam. Instead, this book provides additional context and examples, serves to complement other study guides, and provides additional examples. If you are getting ready to take the exam for the first time, I hope that this guide provides the extra help to pass! If you are up for re-certification, I hope that this guide serves as a refresher and reminder! Keep working hard, keep studying, and never stop learning...