Aruba OS CX Switching Fundamentals

TRAINING MANUAL

EDU-CXF-ILT-v20.21

Aruba Education Services

Aruba OS CX Switching Fundamentals

April 2020

Aruba OS CX Switching Fundamentals Student Guide

Copyright
© 2020 Aruba Networks, Inc. AirWave®, Aruba Networks®, Aruba Mobility Management System®, Bluescanner, For Wireless That Works®, Mobile Edge Architecture, People Move. Networks Must Follow., RFProtect, The All Wireless Workplace Is Now Open For Business, and The Mobile Edge Company® are trademarks of Aruba Networks, Inc. All rights reserved. All other trademarks are the property of their respective owners.

Open Source Code
Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License ("GPL"), GNU Lesser General Public License ("LGPL"), or other Open Source Licenses. The Open Source code used can be found at this site:
http://www.arubanetworks.com/open_source

Legal Notice
The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Warranty
This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.
Altering this device (such as painting it) voids the warranty.

SKU: EDU-CXF-ILT-
v20.21 April 2020

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

ArubaOS-CX Switching Fundamentals

Table of Contents

Module 0: Course Introduction 1
Introductions 2
Logistics 3
Course Objectives 4
Course Objectives 5
Agenda – Day 1 6
Agenda – Day 2 7
Agenda – Day 3 8
Agenda – Day 4 9
Agenda – Day 5 10
How to Use this Guide 11
Module 1: Networking Fundamentals 12
Objectives 13
Overview 14
Basic Network Concepts 16
What is a Computing Network? 17
Network Classifications 18
What is a Protocol? 19
OSI Reference Model 20
Layer 1: Physical Layer 21
Layer 2: Data Link Layer 22
Layer 3: Network Layer 23
Layer 4: Transport Layer 24
Layer 5: Session Layer 25
Layer 6: Presentation Layer 26
Layer 7: Application Layer 27
Question #1 28
Question #2 29
Encapsulation 30
Layer Headers 31
Protocol Data Units (PDU) 32
Physical Media 33

Rev 20.21 | © Copyright 2020 Hewlett Packard Enterprise Development LP | Confidential – For Training Purposes Only

Physical Media – Copper 34
Physical Media – Fiber Optic 35
Physical Media – Fiber Optic (Continue) 36
Physical Media – Air - Radio Frequency 37
Full Duplex and Half Duplex 38
Question #3 39
Types of Traffic 40
Types of Traffic: Unicast, Multicast, Broadcast 41
Numerical Systems 43
Binary Numerical System 44
Binary Numerical System 45
Converting Binary to Decimal 46
Converting Binary to Decimal 47
Converting Decimal to Binary 48
Converting Decimal to Binary 49
Converting Decimal to Binary – Alternate method 50
Converting Decimal to Binary – Compare and Add 51
Converting Decimal to Binary – Patterns for Speed 52
Hexadecimal Numerical System 54
Converting Binary to Hexadecimal 55
Converting Binary to Hexadecimal 56
Converting Hexadecimal to Binary 57
Question #4 58
Lab Activity 59
Additional Content 60
Decimal to Hexadecimal Conversion 61
Converting Decimal to Hexadecimal 62
Converting Decimal to Hexadecimal 63
Converting Hexadecimal to Decimal 64
Converting Hexadecimal to Decimal 65
Converting Hexadecimal to Decimal 66
Summary 67
Module 2: TCP-IP 68
Objectives 69
Overview 70
TCP/IP Stack 71
TCP/IP Stack 72
Ethernet 74
Ethernet Frame 75
IPv4 Header 77
TCP Header – Three-way Handshake 78
TCP Header – Sequence Numbers 80
TCP Header – Port Numbers 81

TCP Header 82
UDP Header 84
Question #1 85
Networking Devices 86
Networking Devices 87
Networking Devices – Switch 88
Networking Devices - Router 89
Networking Devices – Multi-Layer Switch 91
Networking Devices – Wireless Access Points 92
Networking Devices - Firewalls 93
Networking Devices - Servers 94
Operational Planes – Control, Management, and Data 95
Question #2 96
Question #3 97
Common Networking Services 98
Networking Services - DHCP 99
Networking Services - DNS 100
Networking Services - HTTP 102
Networking Services – Telnet and SSH 103
Networking Services - FTP 104
Question #4 106
Lab Activity 107
Additional Content 108
Wi-Fi Frames 109
Wi-Fi Frame 110
Understanding the Wi-Fi Header 111
Summary 112
Module 3: Basic Networking with Aruba Solutions 113
Objectives 114
Overview 115
Network Design 116
Hierarchical Model 117
2-Tier Hierarchy 118
3-Tier Hierarchy 119
Question #1 120
Aruba Switching Portfolio 121
Modern Switching Requirements 122
Switching portfolio: AOS-S 124
Switching portfolio: AOS-CX 127
AOS-CX – Software Architecture 131
Aruba Network Analytics Engine (NAE) Overview 132
AOS-CX Feature set 133
Question #2 135

AOS-CX 136
Accessing AOS-CX CLI – Console port 137
AOS-CX: Port numbering 139
AOS-CX: Prompt modes 140
AOS-CX: Context sensitive help 142
AOS-CX: Command abbreviation and completion 143
Question #3 144
Getting switch information 145
AOS-CX – The value of show commands 146
AOS-CX – Show version 147
AOS-CX – Show system 148
AOS-CX – Show running-config 149
AOS-CX – Top CPU 150
AOS-CX – Show events 151
AOS-CX – Show interfaces brief 152
Question #4 153
Basic Configuration 154
Configuring Hostname 155
Enabling an Interface 156
Verifying an Interface status 157
Discovering the network 158
Link Layer Discovery Protocol 159
ICMP 161
Ping 162
Traceroute 163
Power over Ethernet 164
Question #5 165
Lab Activity 166
Additional Content 167
AOS-CX – Hardware Architecture 168
AOS-CX – Software Architecture 169
AOS-CX – Software Architecture 170
Aruba Network Analytics Engine (NAE) Components 171
AOS-CX – Top memory 172
AOS-CX – Show interfaces transceiver detail 173
Summary 174
Module 4: VLANs 175
Objectives 176
Overview 177
Domains 179
Collision Domains 180
Collision Domains and Performance 182
Broadcast Domains 184

Question #1 186
Virtual LAN 187
LANs and VLANs 188
VLAN Creation 190
Access Ports 191
Question #2 192
802.1Q 193
Extending VLAN Across Multiple Switches 194
Solution: 802.1Q Tagging 195
802.1Q Tagging 196
Configure VLAN Trunks – Allowed VLANs 197
Configure VLAN Trunks – Native VLAN 198
Question #3 199
Forwarding tables 200
MAC Address Table 201
Address Resolution Protocol (ARP) 203
Question #4 205
Frame Delivery 206
Frame Delivery Overview 207
Frame Delivery – Initiate FTP session 208
Frame Delivery – ARP request 209
Frame Delivery – Flood ARP Broadcast 210
Frame Delivery – dot1q Tagging 211
Frame Delivery – Forward Untagged Frame 212
Frame Delivery – ARP Reply 213
Frame Delivery – PC Receives ARP Reply 214
Question #5 215
Lab Activity 216
Summary 217
Module 5: Spanning Tree Protocol 218
Objectives 219
Overview 220
Redundant Network 221
Redundant Network 222
Layer 2 loop 223
Layer 2 loop – Broadcast storms 224
Layer 2 loop – Multiple frame copies 225
Spanning Tree Protocol 226
Spanning Tree Protocol – Operation Overview 227
Spanning Tree Protocol 228
Overview of Spanning Tree Protocol 229
Rapid Spanning Tree Protocol Elements 231
Rapid Spanning Tree Elements Overview 232

RSTP – Bridge identifier 233
RSTP - Bridge Protocol Data Unit 234
RSTP – Port States 235
RSTP – Path cost 236
RSTP - Port Roles (Designated and Root) 237
RSTP - Port Roles (Alternate and Backup) 238
RSTP Operation 239
Rapid Spanning Tree – Operational Overview 240
RSTP – Root Switch Election 241
RSTP – Select Root ports in non-root switches 242
RSTP – Select Designated ports and Alternate ports 244
RSTP Edge Ports and Link Types 246
RSTP – Edge ports 247
RSTP - Topology Change Mechanism 248
RSTP – Topology Change mechanism 249
RSTP – Topology Change mechanism 250
Common Spanning Tree 251
Common Spanning Tree 252
Multiple Spanning Tree 253
Multiple Spanning Tree – 802.1s 254
MSTP – Region configuration 255
MSTP – Region configuration 256
Lab Activity 258
Additional Content 259
RSTP Proposal and Agreement Sequence 260
RSTP – Proposal/Agreement Sequence 261
RSTP – Proposal/Agreement Sequence 262
RSTP – Proposal/Agreement Sequence 263
RSTP – Proposal/Agreement Sequence 264
Rate Limit 265
Loop protection 266
Summary 268
Module 6: Link Aggregation 269
Objectives 270
Overview 271
Overview and Interface Requirements 272
Link Aggregation Overview 273
Link Aggregation Overview 274
Link Aggregation – Interface Requirements 275
Static and Dynamic LAG 276
Static LAG 277
Configuring Layer 2- Static Link Aggregation Group 278
802.3ad - Dynamic LAG 279

802.3ad Operational Modes 280
Configuring Layer 2- Dynamic Aggregation Group 281
Verifying LAG Interface 282
Load Sharing 283
Load Balancing Algorithm 284
Link Aggregation Load Balancing 285
Lab Activity 286
LACP Data Unit Flags 287
Summary 288
Module 7: IPv4 Routing Part 1 289
Objectives 290
Overview 291
Routing Introduction 292
Routing 293
Routing – Layer 3 analysis 294
IP Addressing 295
IPv4 address 296
Network Mask 297
Question #1 299
IP Routes and Default Gateway 300
IP route 301
Default Gateway 302
Inter-VLAN Routing 304
Multi-Layer Interface Types 305
DHCP Helper address 307
Inter-VLAN Routing 309
IP Routing Table 310
Question #2 312
Packet Delivery 313
Packet Delivery Scenario 314
Packet Delivery- Endpoint to Access Switch 315
Packet Delivery – Access to Multilayer switch 316
Packet Delivery – Routing Process 317
Packet Delivery – Multilayer to Access Switch 318
Packet Delivery – Access Switch to Endpoint 319
Question #3 320
Virtual Routing and Forwarding 321
Virtual Routing and Forwarding 322
Lab Activity 323
Summary 324
Module 8: VRRP 325
Objectives 326
Overview 327

VRRP 328
Need for Layer 3 redundancy 329
First Hop Redundancy Protocol 330
Virtual Router Redundancy Protocol 331
VRRP Instances 332
VRRP Instances – Validate Capacity 333
Master Election 334
Virtual IP address 335
VRRP Failover Operation 336
VRRP Preemption 337
VRRP and MSTP Coordination 338
Question #1 339
Question #1 340
Lab Activity 341
Summary 342
Module 9: IP Routing Part 2 343
Objectives 344
Overview 345
Subnetting 346
IPv4 Address Classes 347
IPv4 Address Classes and Default Masks 348
Reserved Addresses 349
Private and Public IPv4 Addressing 351
Classful Network Disadvantage 352
Subnetting 353
Subnet Mask – Using Binary 354
Subnet Mask – Class B Example 355
Subnet Mask – Class C Example 356
Finding the Network, Subnet and Host part 357
Subnetting and Network Design 358
Subnetting Design Example 359
Subnetting Tasks 360
Subnet Analysis Introduction 361
Subnet Analysis 362
Subnet Analysis 363
Subnet Analysis – Identify First Host IP 364
Subnet Analysis – Identify Broadcast Address 365
Subnet Analysis - Identify the Last Host Address 366
Listing the Available Subnets 367
Listing the Available Subnets 368
Determining the Best Mask 369
Determining the Best Mask 370
Determining the Best Mask 371

Question #1 372
Question #2 373
VLSM and CIDR 374
Fixed vs Variable Length Subnet Mask 375
VSLM Example 377
VSLM Example 378
VSLM Example 379
VSLM Example 380
VSLM Example 381
VLSM Example 382
VLSM Example 383
Classless Inter-Domain Routing Introduction 384
CIDR Example 385
Question #3 386
Lab Activity 387
Summary 388
Module 10: IP Routing Part 3 389
Objectives 390
Overview 391
Route Types and Administrative Distance 392
Route Types in ArubaOS-CX 393
Static Routes 395
Administrative Distance 396
Floating Static Routes 397
Scalability Issues 399
Question #1 400
Routing Protocols 401
Interior and Exterior Gateway Protocols 402
Distance Vector Routing Protocols 403
Link-State Routing Protocol 405
Question #2 407
Lab Activity 408
Summary 409
Module 11: OSPFv2 Single Area 410
Objectives 411
Overview 412
OSPFv2 413
OSPF Introduction 414
Router ID Overview and Selection Criteria 415
OSPF General Operation Overview 416
OSPF General Operation - Neighboring 417
OSPF General Operation – Build a List of Every Link 418
OSPF General Operation – Build a List of Best Paths 419

OSPF General Operation – Build a List of Best Paths 420
OSPFv2 421
Hello Messages 422
OSPF Neighbor States 423
OSPF Neighbor States 425
OSPF Neighbor States 426
OSPFv2 427
OSPF Network Types 428
Broadcast Network scalable problem 430
Designated Router and Backup Designated Router 431
Designated Router Election 432
OSPF Area 433
OSPF LSA Type 1 (LSA 1) 434
OSPF LSA Type 1 (LSA 1) 436
OSPF LSA Type 2 (LSA 2) 438
Path Selection 439
OSPF Convergence 440
Passive interfaces 441
Lab Activity 442
Additional Content 443
OSPF Scalability 444
Using Cost to Manipulate Routes 445
Configuring Cost Value 446
Configuring OSPF in AOS-CX 447
Summary 449
Module 12: Stacking 450
Objectives 451
Overview 452
Stacking Technologies 453
Operational Planes – Control, Management, and Data 454
Introduction to stacking technologies 455
Distributed Data Plane and Distributed Link Aggregation 456
Aruba VSF Stacking Solution and Platforms 457
VSF Member Roles and Links 458
VSF Open Virtual Switch Database 459
VSF Topologies 460
VSF Requirements 461
VSF member ID and Port numbers 462
VSF Configuration Example 463
VSF Pre-Provisioning 464
Tracing Layer 2 traffic: Unicast 465
Tracing Layer 2 traffic: Unicast 466
VSF Failover and OSPF Graceful-Restart 467

VSF Failover and OSPF Graceful-Restart 468
VSF Link failure 469
Multi-active Detection – Split Detection 470
VSX Introduction 471
Question #1 472
Question #2 473
Question #3 474
Lab Activity 475
Summary 476
Module 13: Secure Management and Maintenance 477
Objectives 478
Overview 479
Management and Maintenance 480
Out-Of-Band Management Port 481
Management VRF 482
Ping and Traceroute in the Management VRF 483
SSH for ArubaOS-CX 484
HTTPS for ArubaOS-CX 485
Web Interface 486
Authentication, Authorization and Accounting 487
Role Based Access Control - RBAC 488
RBAC Configuration 489
RADIUS-based Management Authentication 490
Simple Network Management Protocol - SNMP 492
SNMP Versions 494
Configuration File Management 495
Checkpoint Overview 497
Checkpoint Configuration 498
Checkpoint Auto Mode 499
Operating System Image Management Introduction 500
Operating System Image Management Access 501
Password Recovery Process 502
Factory Default 503
Question #1 504
Question #2 505
Question #3 506
Question #4 507
Lab Activity 508
Summary 509
Module 14: Management Tools 510
Objectives 511
Overview 512
Management Tools 513

Introduction to Aruba NetEdit 514
NetEdit Installation 515
NetEdit Initial Configuration – Set Password 516
NetEdit Initial Configuration - Licensing 517
NetEdit Initial Configuration – Network Configuration 518
NetEdit GUI Login 519
NetEdit GUI Login – Change Password 520
NetEdit Device Discovery 521
Configure Switches to Support NetEdit 522
NetEdit – Device Discovery Process 523
NetEdit – Device Discovery Process 524
NetEdit – Device Discovery Process 525
NetEdit Device Monitoring 526
NetEdit – Device Details 527
NetEdit – Device Details 528
NetEdit – Device Details 529
NetEdit Device Management 530
NetEdit – Configuration Plan 531
NetEdit – Create a new VLAN 532
Aruba CX Mobile App Overview 533
Aruba CX Mobile App Features 534
Question #1 535
Question #2 536
Question #3 537
Lab Activity 538
Summary 539
Glossary of Terms 540

aruba
a Hewlett Packard Enterprise company
Course Introduction
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 0 – 1

Introductions

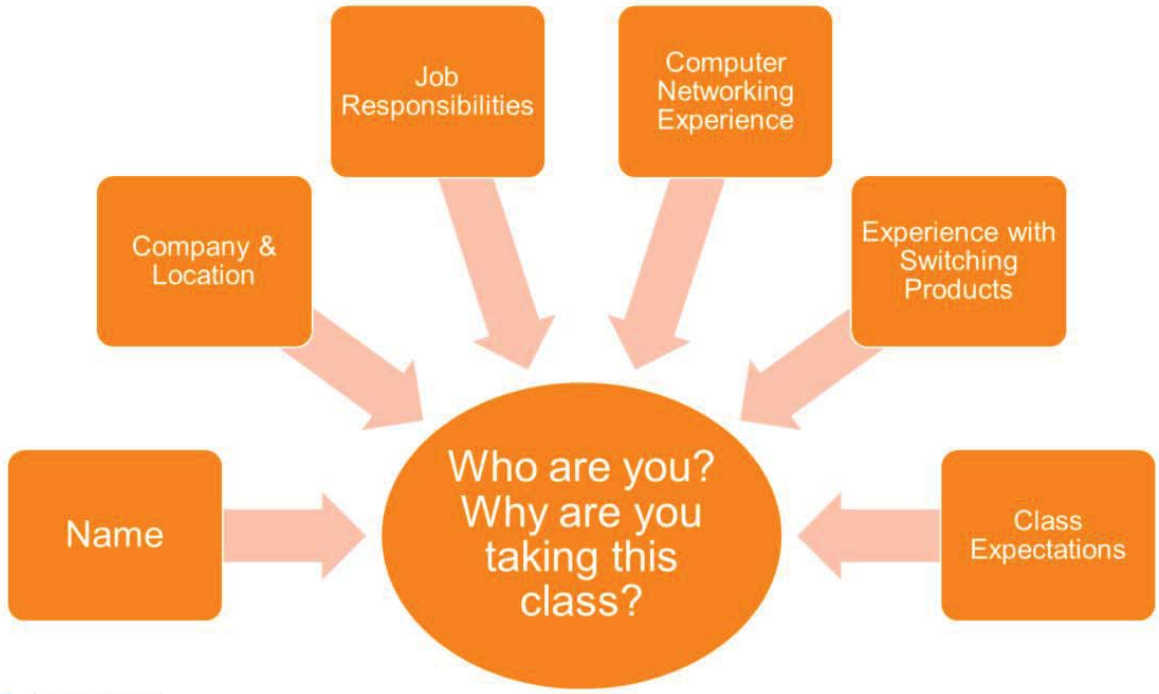

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Logistics

Introduction

- Breaks
- Lunch
- Restrooms

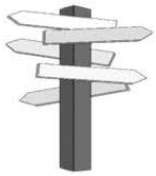

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Course Objectives

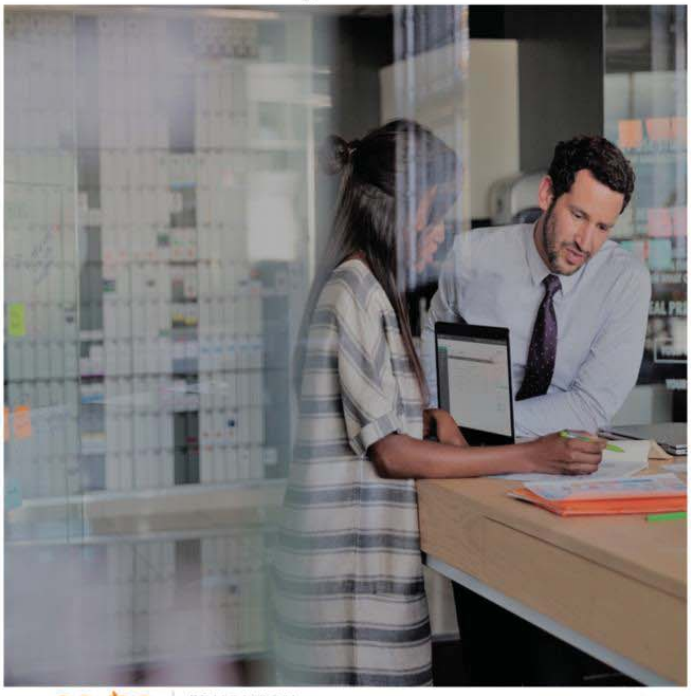

- Network Fundamentals
- Review Aruba Switching portfolio
- Aruba OS-CX Network Operating System
- VLANs
- Spanning-Tree Protocol
- VRRP
- Link Aggregation
- IP Routing
- Subnetting

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 0 – 4

Course Objectives

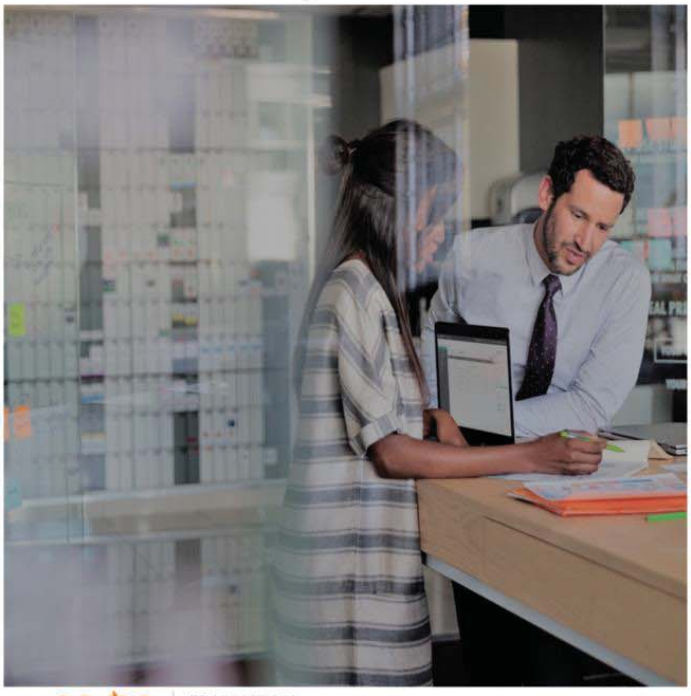

- OSPFv2 – Single Area
- Stacking using VSF
- Secure Management and Maintenance
- Aruba NetEdit

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 0 – 5

Agenda – Day 1

1 – Networking Fundamentals

2 - TCP/IP Stack

3 – Basic networking with Aruba Solutions

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Agenda – Day 2

- 4 – VLANs
- 5 - Spanning-Tree Protocol

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Agenda – Day 3

- 6 - Link Aggregation
- 7 – IPv4 Routing - Part 1
- 8 - VRRP

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Agenda – Day 4

- 9 - IP Routing - Part 2
- 10 – IP Routing- Part 3
- 11-OSPFv2

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Agenda – Day 5

- 12 - Stacking
- 13 – Secure Management and Maintenance
- 14 – Management Tools

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

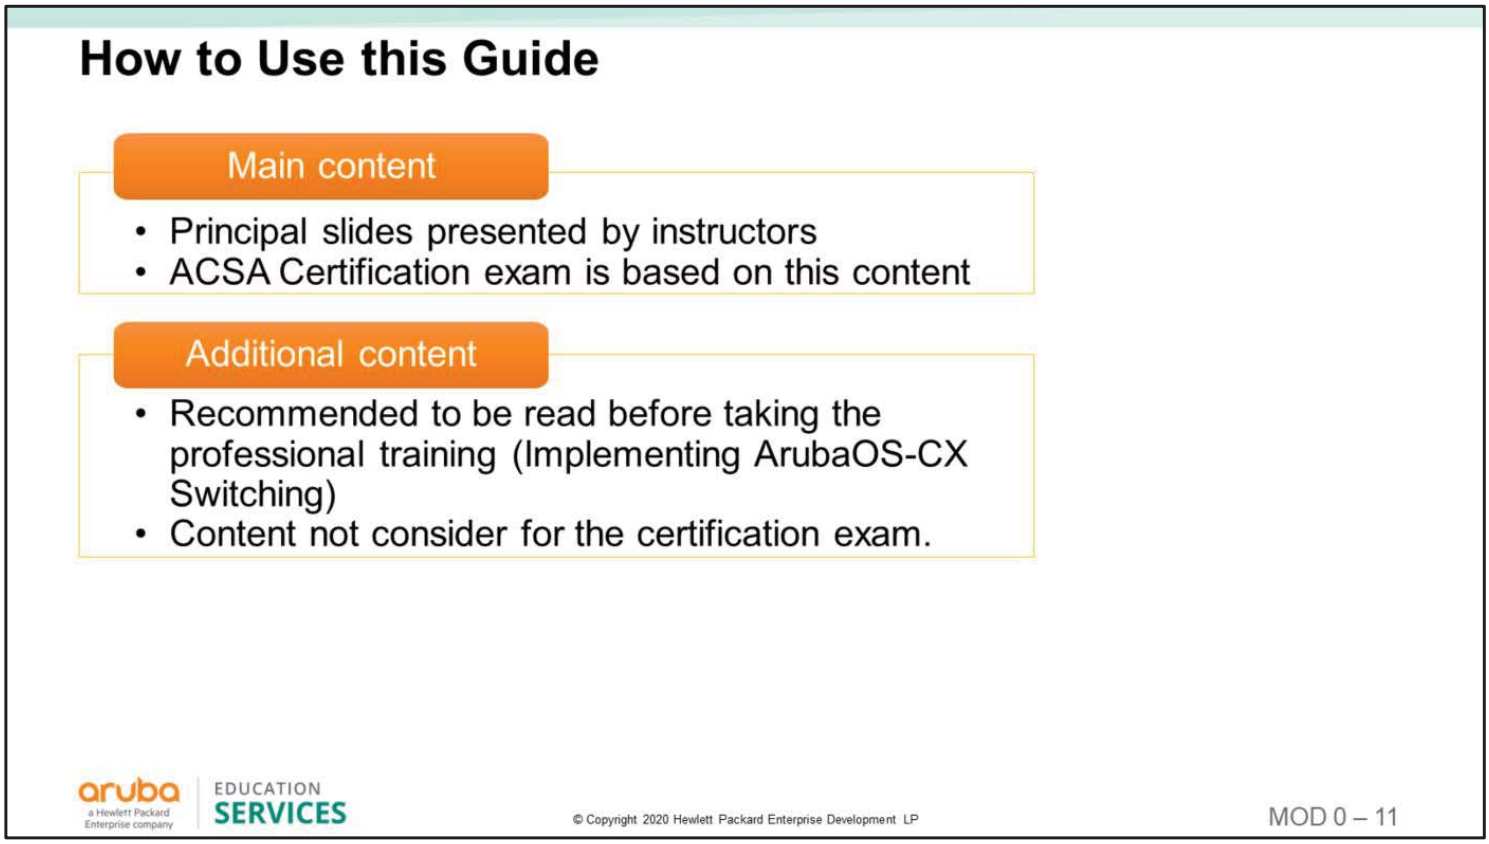

The main content on this guide will be presented by the Aruba instructors and it will be considered in the Certification exam to for the Aruba Certified Switching Associate.

After each module you will find **Additional Content** which Aruba recommends you ready before taking the professional switching course: Implementing Aruba OS-CX Switching.

aruba
a Hewlett Packard Enterprise company
Networking Fundamentals
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP

Objectives

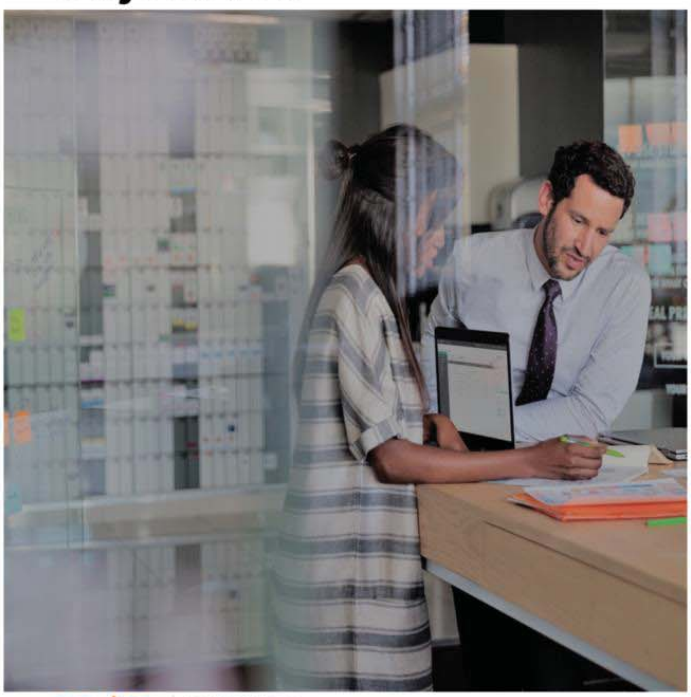

–Learning Objectives

- –Describe computing networks
- –Describe protocols and the OSI model
- –Explain encapsulation and headers
- –Convert numbering systems: decimal, binary, and hexadecimal
- –Describe the TCP/IP protocol stack
- –Compare unicast, multicast, and broadcast

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 2

After completing this module, you will be familiar with the fundamental concepts that serve as a foundation for mastering computer network technology.

|You will be able to define computer networking and describe fundamental concepts including Local Area and Wide Area Networks, along with their typical use cases and components.

| You will be able to describe how devices communicate over these networks using protocols, and how this communication is organized into a seven-layer Open Systems Interconnection (OSI) model.

| You will be able to explain how each layer of the OSI model has different functions, and adds certain header information to encapsulate data for transmission, and decapsulate it upon receipt.
| Then you will learn about the decimal, binary, and hexadecimal numbering systems, enabling you to convert numbers from one system to another.

This theoretical knowledge starts to take shape as you learn about actual protocols used in corporations and the Internet for communications.
| You will be able to describe the TCP/IP protocol stack and how it maps to the OSI model. You will be able to explain how each layer in the stack adds certain header information – from the application, down to the TCP or UDP protocol, the IP protocol, and on down to the Ethernet or Wi-Fi Network Interface Card (NIC).

|Finally, you will be able to describe three types of traffic used in network communications – unicast, multicast, and broadcast.

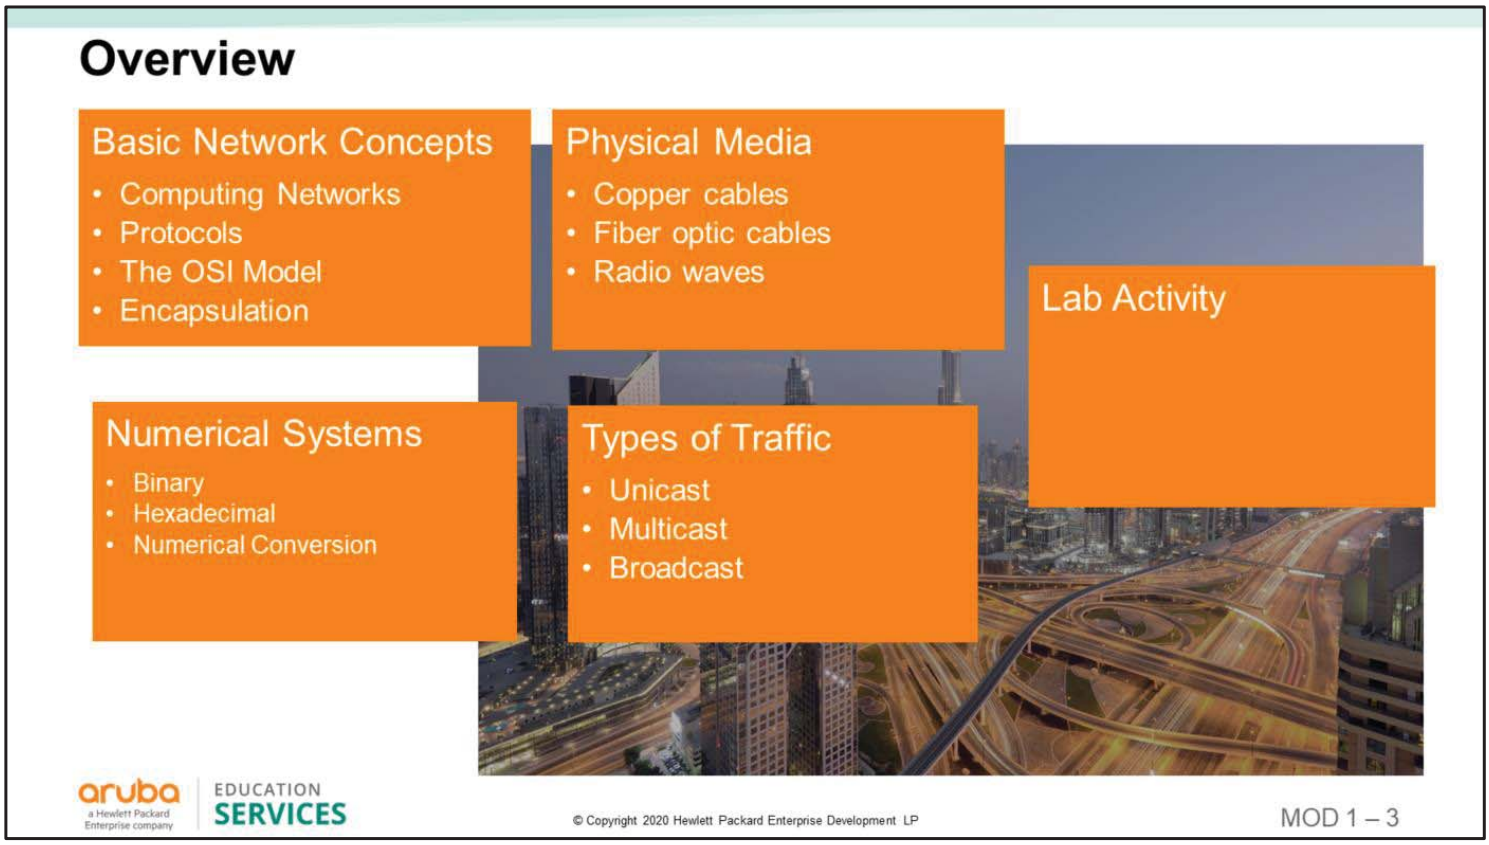

The course lays the foundational knowledge upon which the rest of your training, and perhaps a significant portion of your career is based. Mastery of this material will greatly ease future learning about more advanced concepts.

So, let us start at the beginning. You learn what a basic computer network is. You compare local and wide area networks, use cases, and implementation types. Communication over these networks relies on protocols, which are organized into a standardized, seven-layer OSI model. Once you learn about each of these seven layers, you will be well on your way to understanding network communications.

You learn which protocols map to each layer of the OSI model, and how each protocol encapsulates the data with header information to create segments, packets, and frames. Then this data is transmitted over some physical transmission medium – copper or fiber optic cables, or radio waves in the case of Wi-Fi. You learn how each media type has a unique method of representing the 1’s and 0’s of digital data, which leads to an exploration of binary numbers.

You learn about the binary number system, and how to convert from binary to decimal, and from decimal to binary. This is a vital skill with practical, real-world application – especially as you move into more advanced concepts, where you learn about network addressing, advanced routing concepts, and certain security concepts. You will also learn about the hexadecimal numerical system and related conversions to and from decimal and binary. This is important for working with certain Local Area Network (LAN) addresses, as well as working with the Internet Protocol version 6 addresses (IPv6).

Now you learn what is arguably the most important set of communication rules ever devised

by man. The TCP/IP protocol stack is used in every corporation in the world, and across the internet. Your knowledge of this protocol will unlock a host of valuable skills – including network design, implementation, and troubleshooting. Detailed knowledge of this protocol stack is especially vital for those interested in network security and mitigation malicious attacks on valuable corporate resources. You will get a high-level overview of this protocol suite and explore Ethernet and Wi-Fi frame formats and concepts.

Finally, you learn about three types of network traffic – unicast, multicast, and broadcast, before exploring the lab environment in a hands-on activity.

Basic Network Concepts

Computing Network, Protocol and OSI model

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

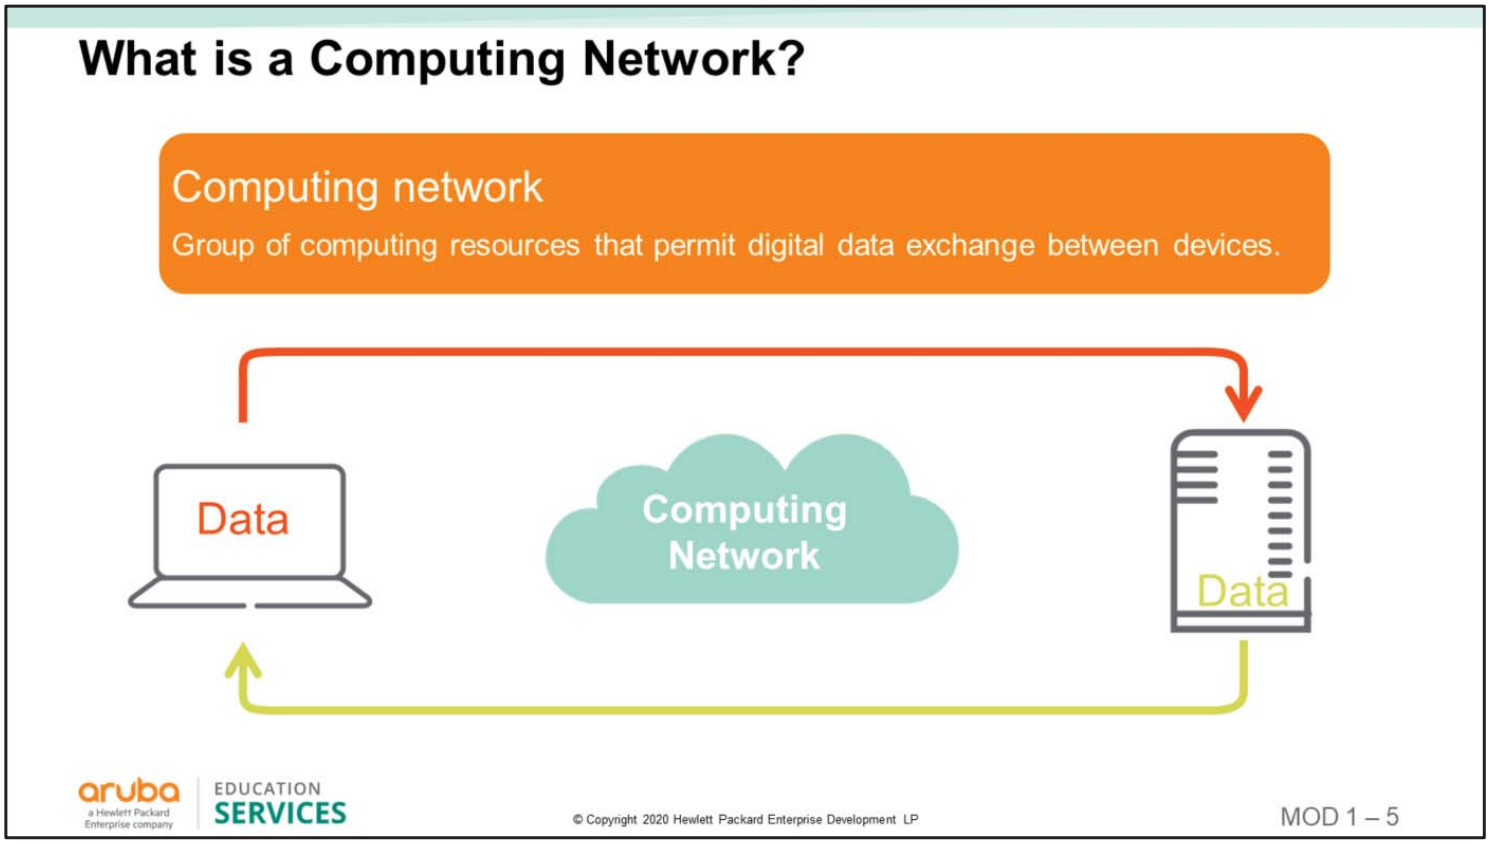

A computing network is defined as a group of computing resources that permit digital data exchange between computer devices - regardless of the type or vendor.

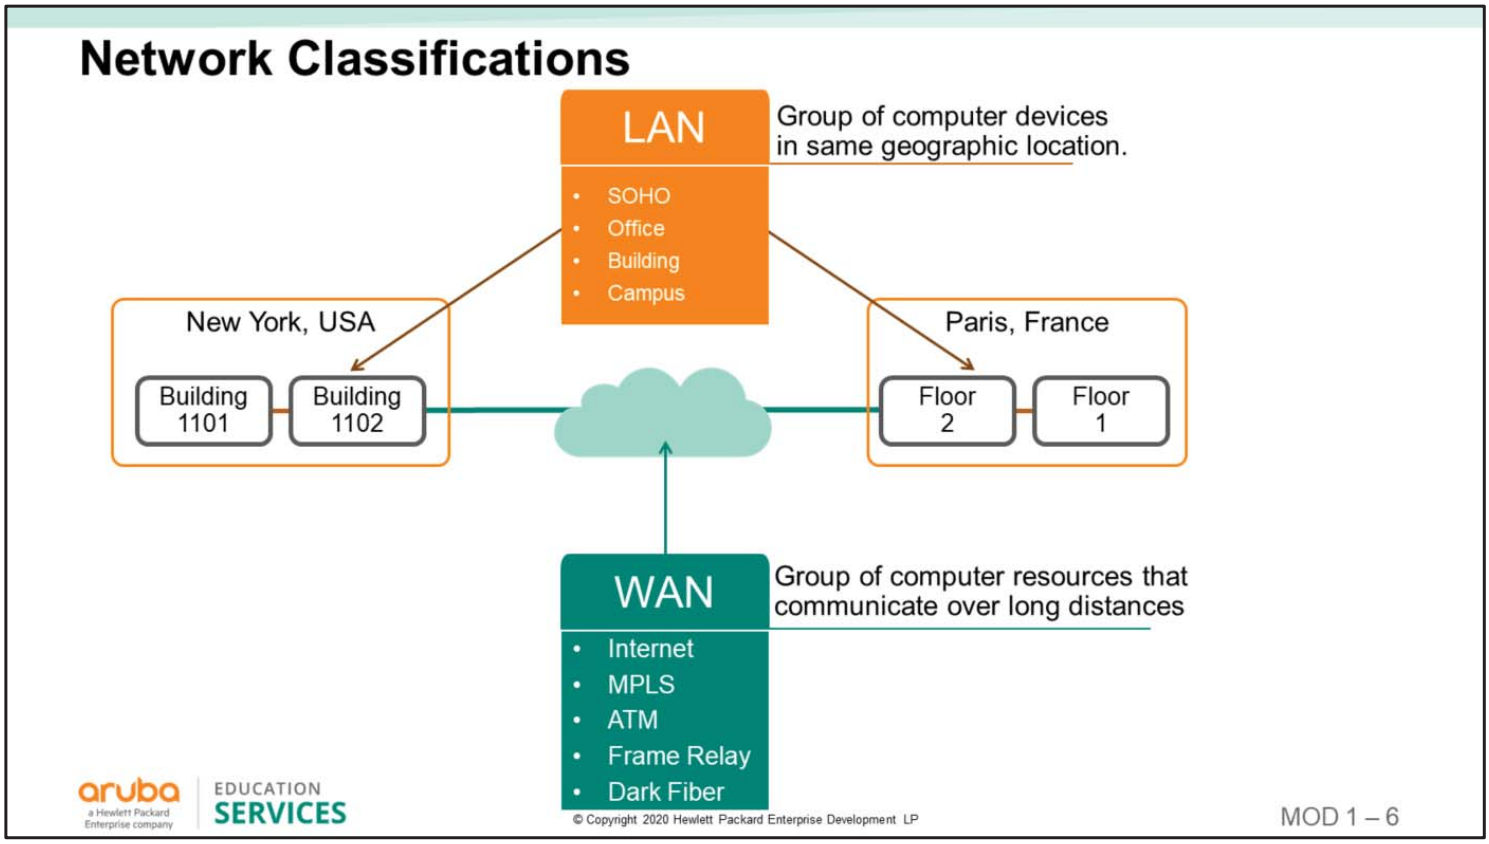

Based on the geographical coverage a computing network can be categorized as a Local Area Network (LAN) or a Wide Area Network (WAN).

A LAN is a group of computer devices that are geographically in the same place. For example, a group of devices within a building can be considered a LAN.
|LANs are used in several settings:

- Small Office/Home Office (SOHO)
- Office LANs
- Building LANs
- Campus LANs

|A WAN in the other hand is a group of computer resources that can communicate over large geographical distances – typically a few kilometers or miles, and perhaps thousands of miles, such as the Internet. The Internet is considered a WAN since it permits communication across countries and continents.

|Typically, WANs are deployed by Internet Service Providers (ISP) since those companies have the economic resources to interconnect long distances. Examples of WAN technologies include the following:

- Internet
- Multi-Protocol Label Switching (MPLS)
- Asynchronous Transfer Mode (ATM)
Frame Relay (largely obsolete)
Dark Fiber

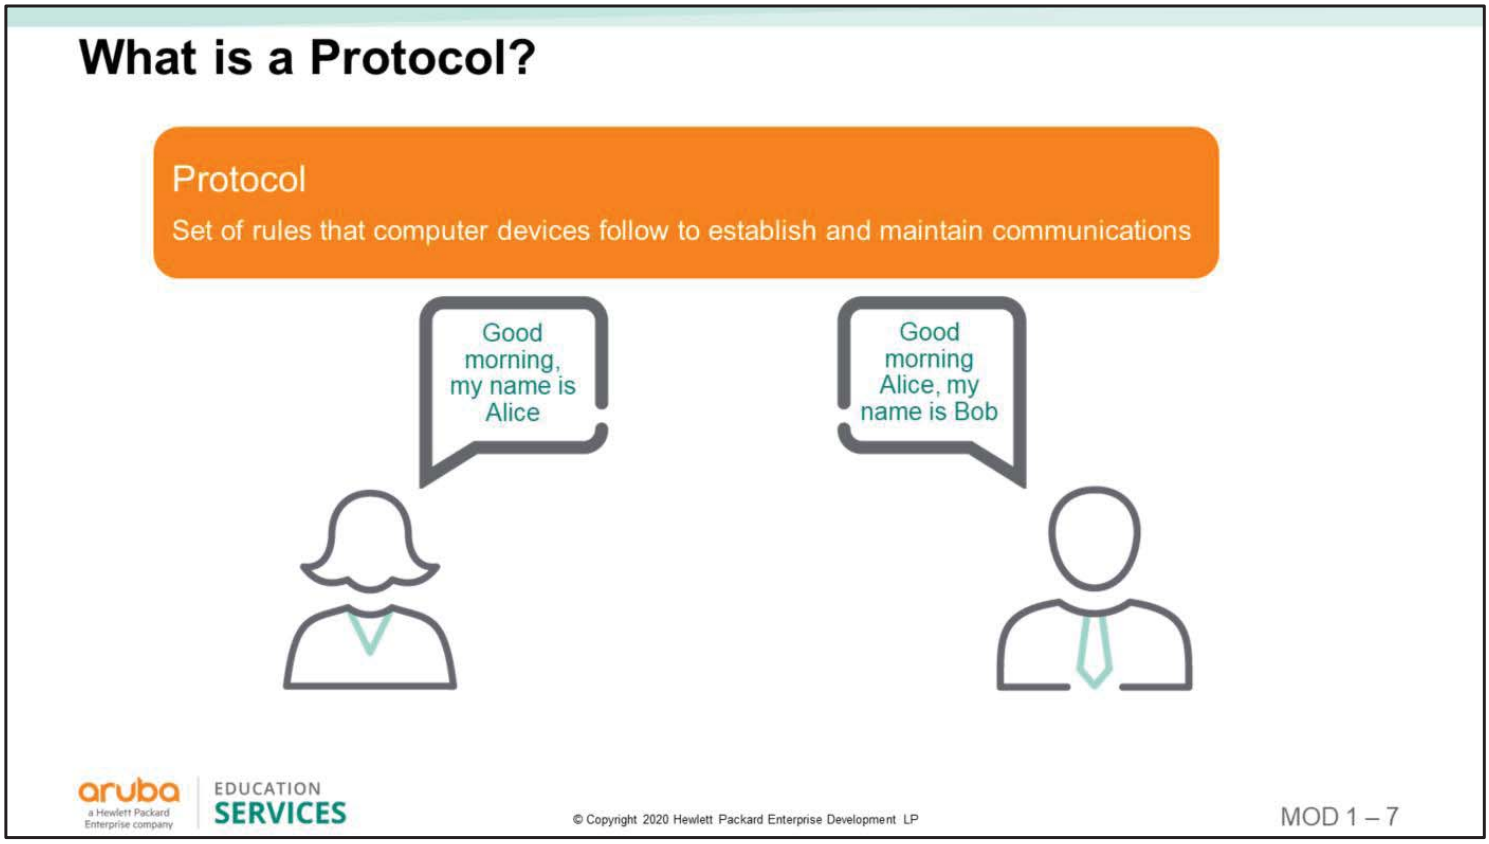

Communication is the main purpose of a computing network, and this communication is enabled using protocols.

A protocol is a group of rules or instructions that computer devices follow in order to properly establish and maintain communications.

Consider this analogy:

Alice meets Bob for first time, and starts the conversation by saying “Good morning, my name is Alice”.
Bob replies, “Good Morning Alice, my name is Bob,”

This brief conversation is actually a procedure. Notice that Alice starts the communication with a greeting, and then she identifies herself. Bob’s reply is also a procedure. He acknowledges Alice, and then he identifies himself by name.

The implicit rules in this conversation help to establish and maintain a conversation. Likewise, computing devices exchange messages in a specific order, following specific rules.

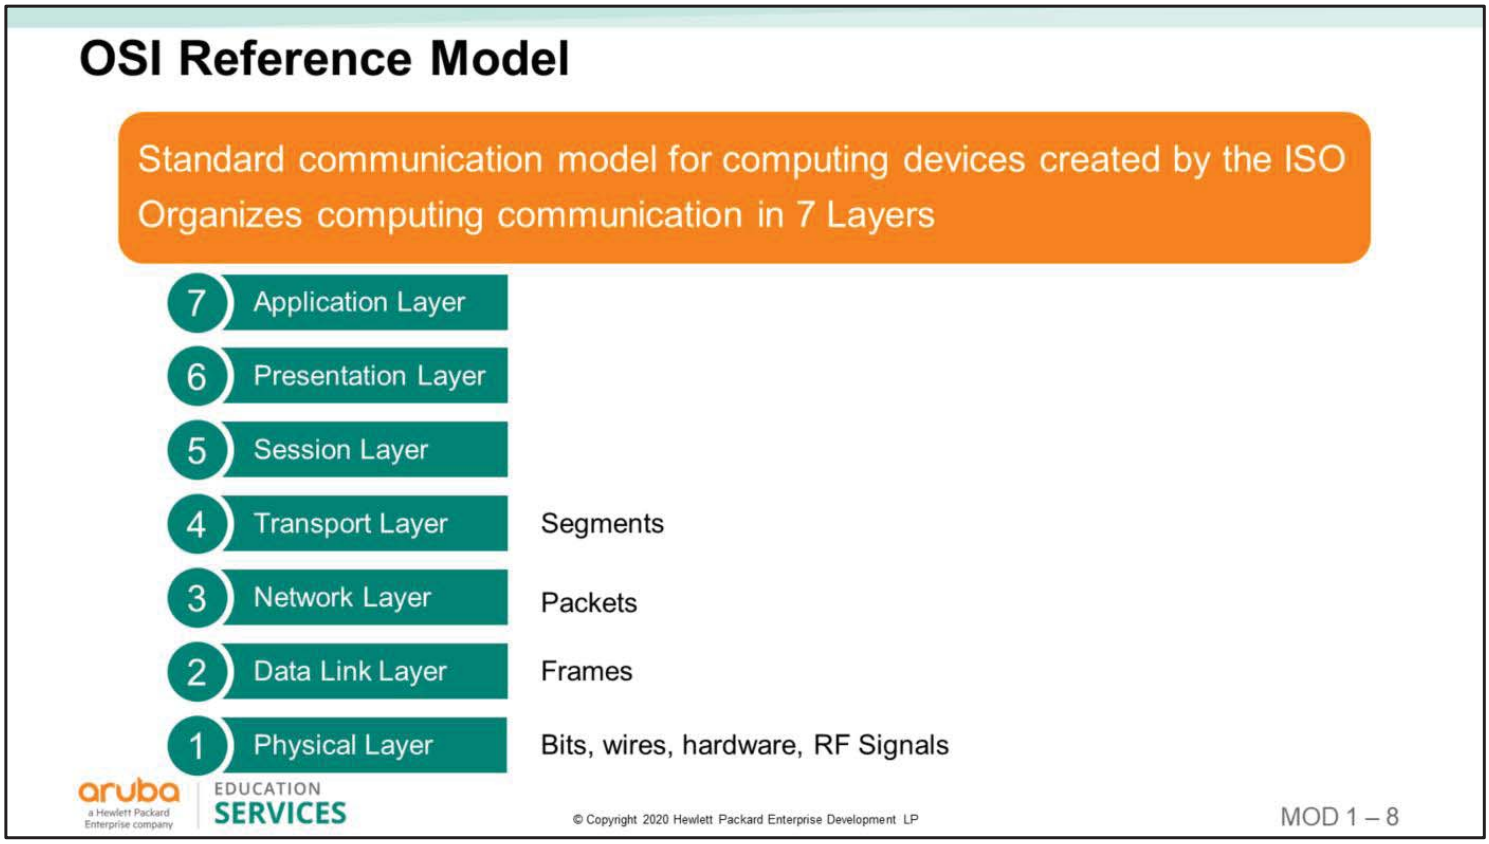

In the mid 1980's during the fast evolution of computing, every vendor wanted to implement their own, proprietary communication protocol. These proprietary protocols created interoperability challenges.

The International Organization for Standardization (ISO) solved the problem by presenting a standard communication model for computing devices – The Open Systems Interconnection (OSI) model.

The OSI model organizes computing communication in 7 layers. Each layer defines a phase of message processing. The OSI layers are shown in the figure, and described below:

Layer 7: Application Layer
Layer 6: Presentation Layer
Layer 5: Session Layer
Layer 4: Transport Layer, which organizes data into segments, as you will soon learn.
Layer 3: Network Layer, which organizes data into packets
Layer 2: Data Link Layer, which organizes data into frames
Layer 1: Physical Layer, which organizes the data into bits, and transmits those bits using physical hardware over wires, fiber optic cable, or RF Signals.

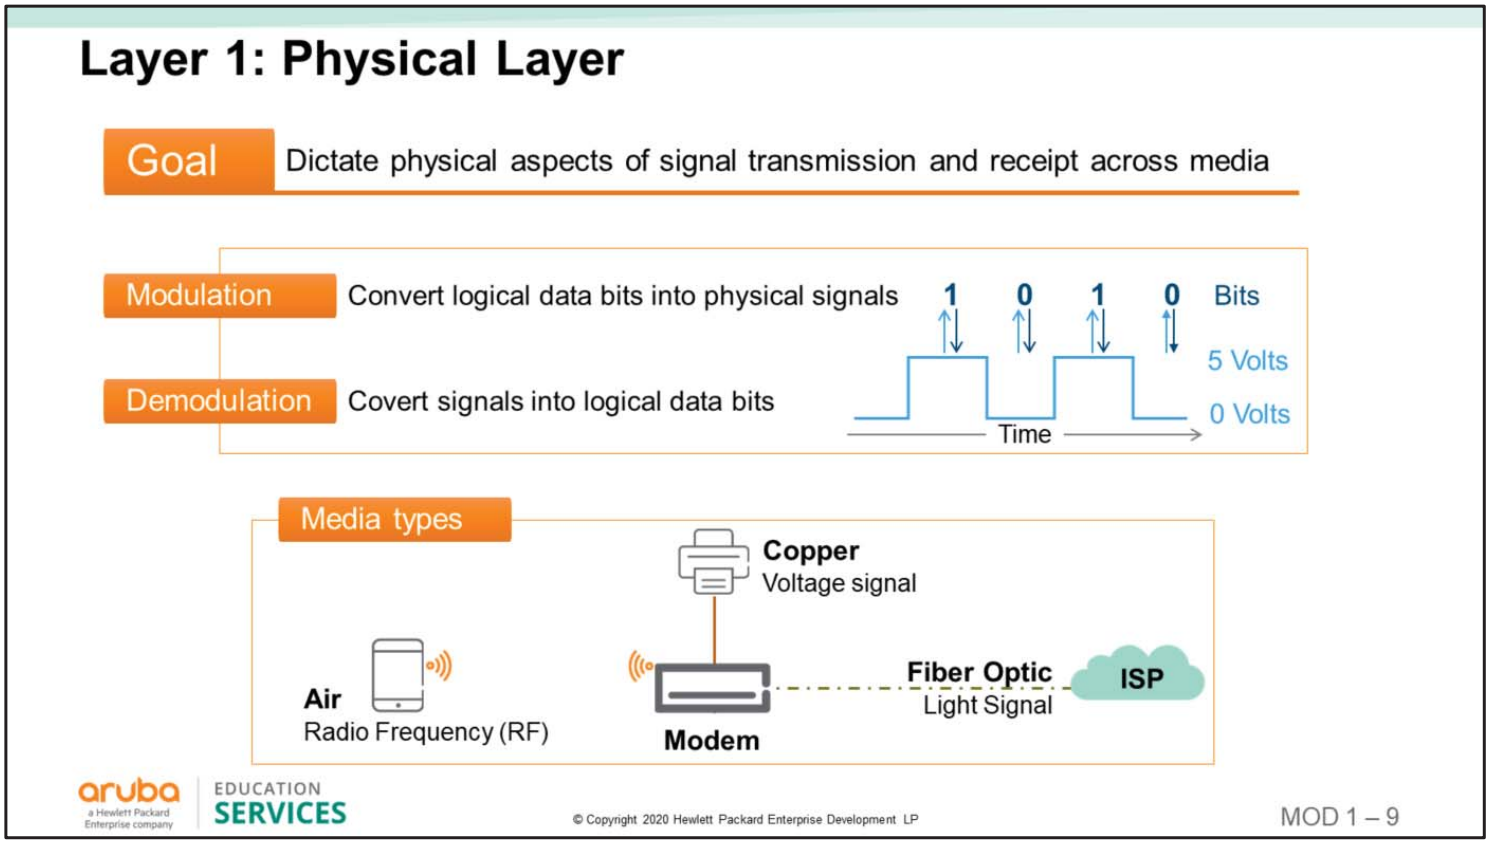

This layer dictates the physical aspects of how signals are transmitted and received across some media.

|Computing devices convert logical data bits into the correct physical principle depending on the media in use, this process is known as modulation.
|The inverse process, converting signals into logical data bits is known as demodulation.

This layer also defines the material characteristics and the components to use to achieve the correct transmission of the messages.

|For example, consider the modem that is used in a home basis, your tablet connects to the modem using Radio Frequency (RF) signals that travel across the air. This modem connects to the Internet Service Provider (ISP) network using Fiber Optic cables. Thus, the router converts data into optical signals. Also, the printer connects directly to the modem via a twisted pair copper patch cable. This means that print jobs are converted into electrical voltage signals before being sent to the printer.

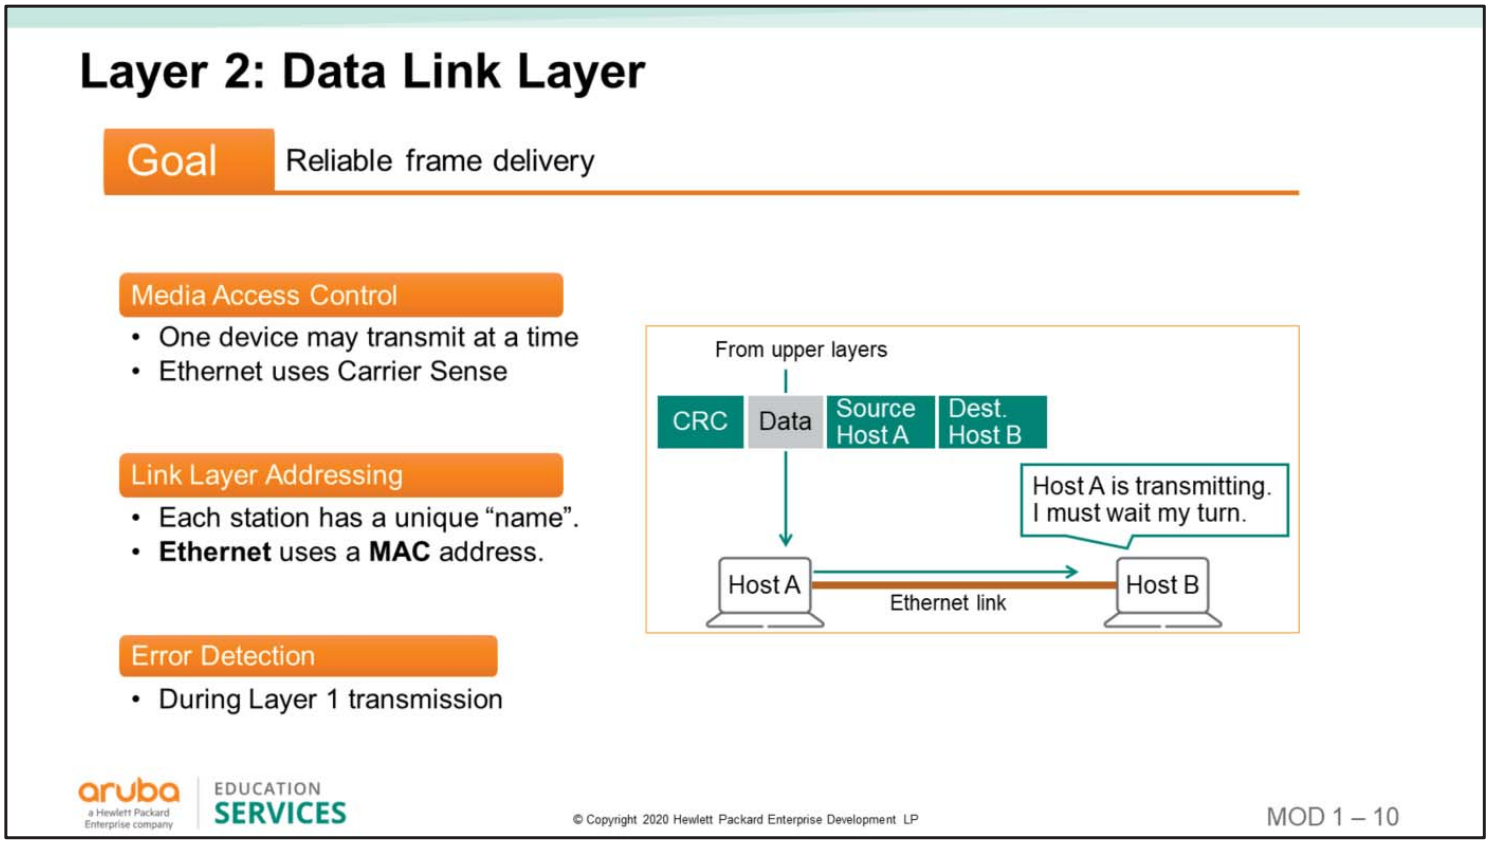

The Data Link provides 3 main functions:

|Media Access Control. In polite human conversation, only one person talks at a time. While Bob talks, Alice politely listens. She detects when Bob stops talking and knows that it is her turn to speak. Similarly, for most media, only one device may transmit at a time, so access to the media must be controlled.

Likewise, many Media Access Control techniques leverage a "Carrier Sense" mechanism – before a station may transmit, it must sense the state of the carrier or media. If a transmission signal is detected, the station must wait to transmit, until the currently transmitting station is complete.

|Link Layer Addressing. In a crowded room, Bob may call out, "Hey, Alice can I buy you an ice cream?" Everyone hears these sounds waves, which travel over a shared media – the air in the room. However, only Alice responds - she was identified as the intended recipient.

Likewise, each station on a LAN has a unique "name". Instead of an alphanumeric name like Alice or Bob, each station is identified by a 6-byte hexadecimal number called a MAC address. All stations on a shared media receive the message, but only the device identified as the intended recipient processes and responds. Like humans, the other stations realize, "This message is not for me" and simply ignore the message. This information is added to the data from the upper layers as so-called "header information", about which you will soon learn.

|Error Detection. For the receiver side, Layer 2 helps to detect errors that could occur during Layer 1 transmission. This avoids unnecessary processing of corrupted or incomplete messages. This is accomplished by adding a "Trailer" to the data.

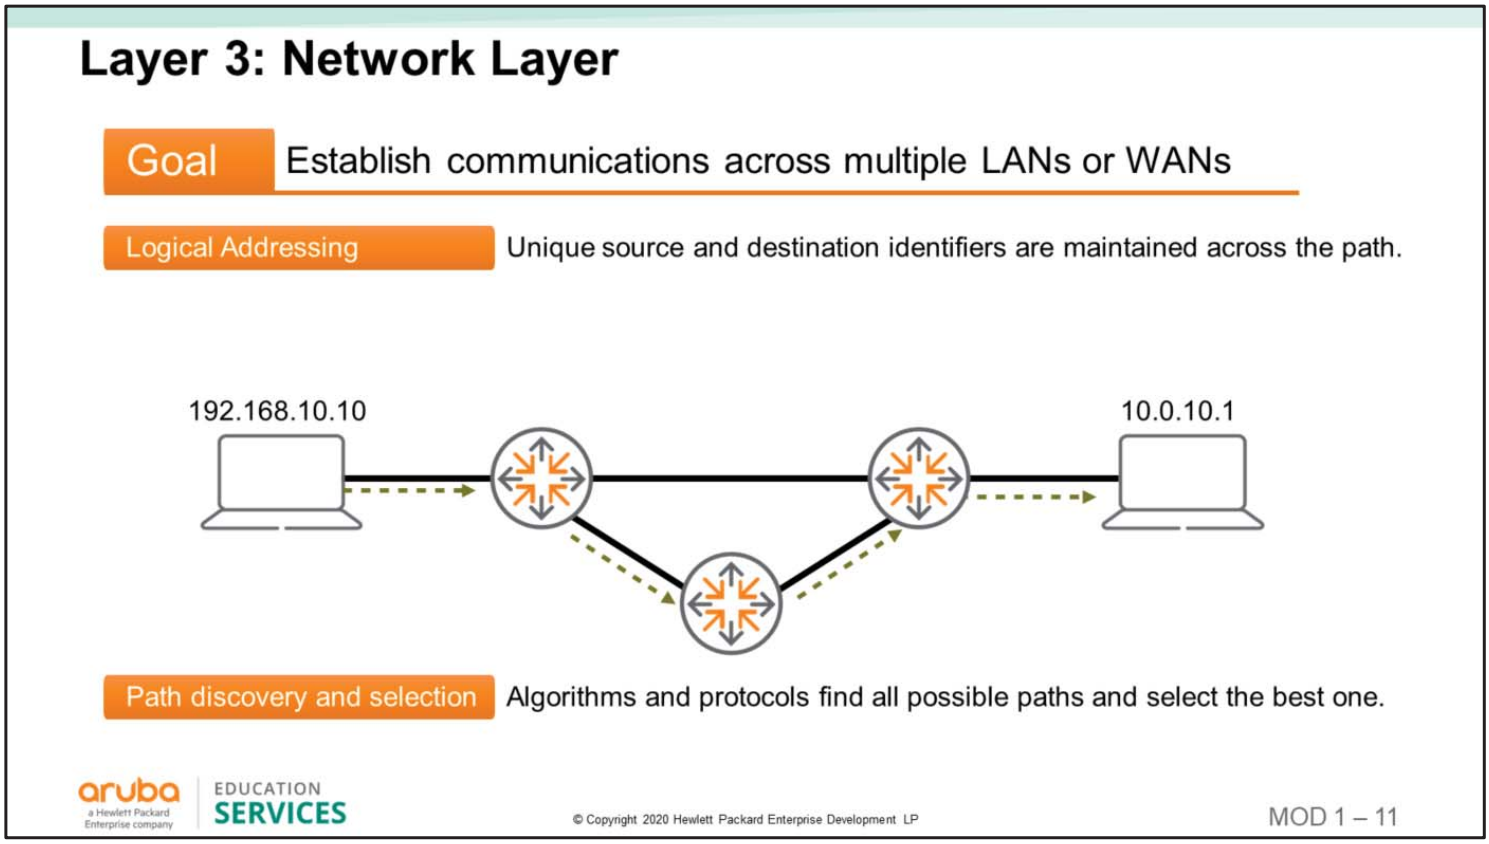

The main goal of the Network Layer is to establish device communications across multiple LANs or WANs, using the best available path. This is achieved using two fundamental techniques.

|Logical addressing. A unique Layer 3 identifier for the source and destination is maintained across the path.
|Path discovery and selection. The Network Layer runs algorithms and protocols to find all possible paths, and then choose the best path. Later in this course, you will learn more about this, and things like the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).

The communication between two computing devices can take a specific path but not necessary the same one will be used in the future. Protocols and algorithms used in this Layer could update the path any time depending of multiple factors. You will learn more about it in this training.

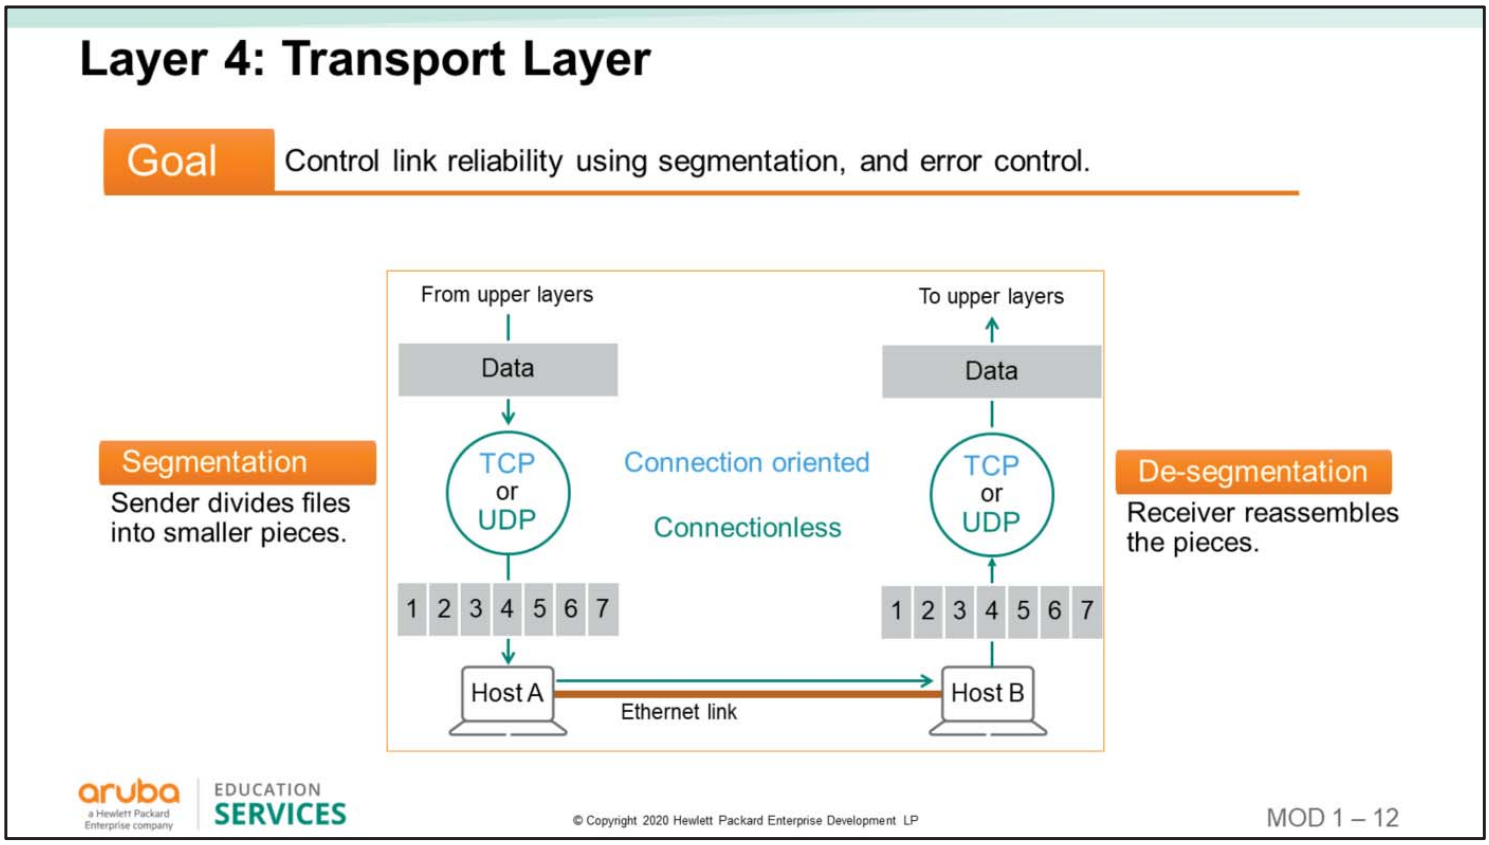

The Transport Layer controls the reliability of a given link through segmentation, de-segmentation and error control. In this layer some protocols, like the Transmission Control Protocol (TCP) are connection-oriented. This means that the transport layer can keep track of the messages and retransmit those that fail. Other protocols, like the User Datagram Protocol (UDP), are stateless or connectionless. This means the transport layer does not keep track of the messages. The advantage of this is that processing these connections is relatively fast and easy to compute.

There are three main responsibilities of the transport layer, as described below.

|Segmentation. The sender's TCP or UDP process accepts files from the application and divides them into smaller pieces (typically 1500 bytes) called segments. Each piece is passed down to the lower layers and transmitted individually – over an Ethernet link in the example shown in the figure.
|De-segmentation. The receiver accepts each segment, puts them back in the correct order if need be, and reassembles the information. This then can be processed by the application.
Error Control. Refers to the verification of the information received to avoid errors that could occur the lower Layers (1-3).

Note. Error detection is a process that happens in different Layers: 2, 4 and sometimes in 7 of the OSI model.

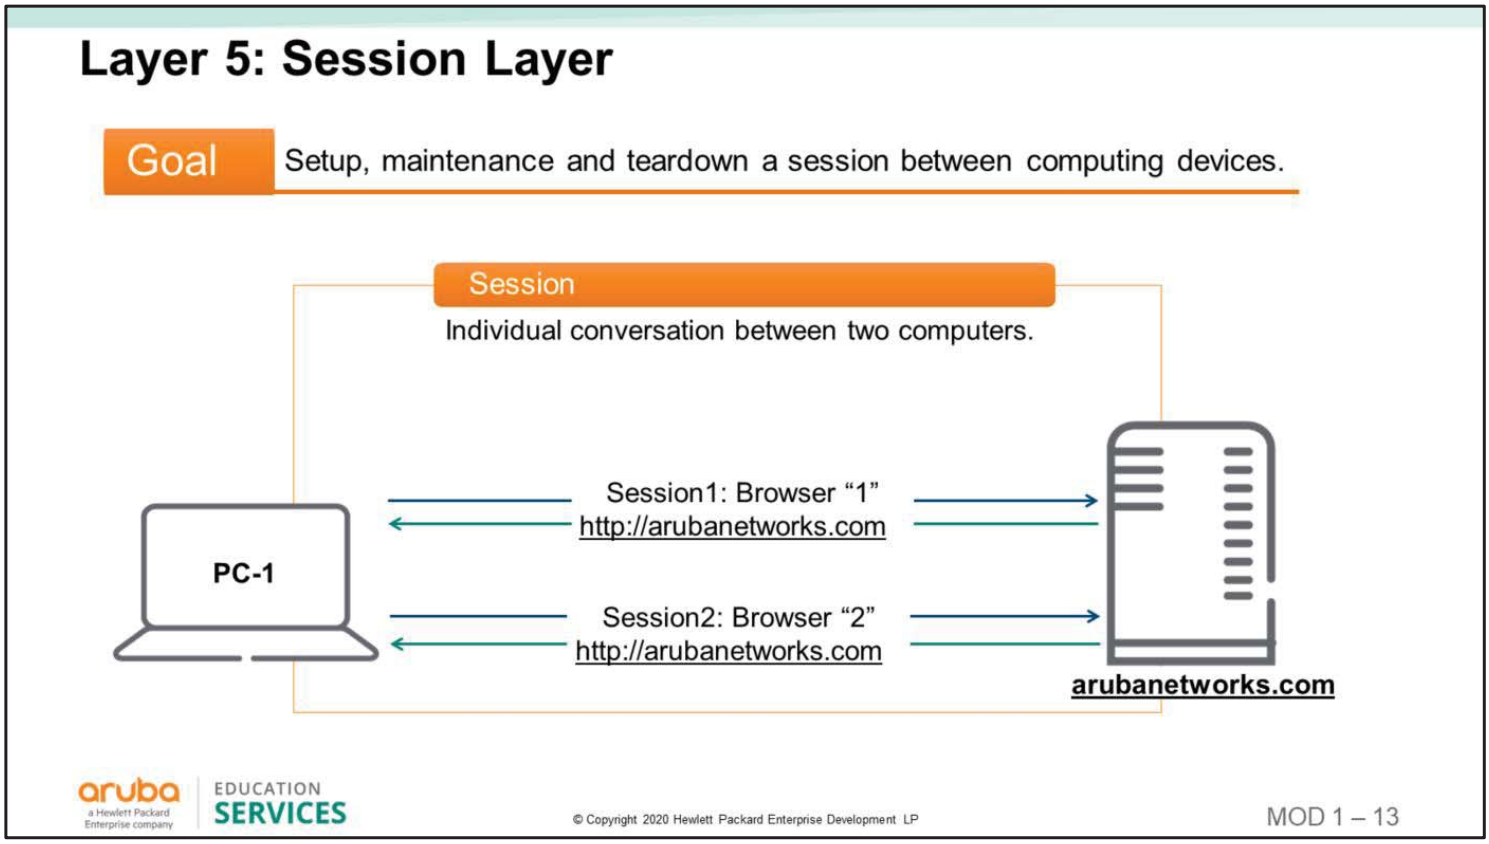

Layer 5 is responsible for setup, maintenance, and tear down of sessions between two computing devices. A session is a particular conversation between two computer devices.

Suppose that some user opens a browser and connects to a web page like http://arubanetworks.com. A session is created. Then the same user opens a different browser to the same destination, since the application is different a new conversion or session is created. Thus, two separate sessions are maintained.

Then the user might establish a connection to a different host for purposes of file transfer, and seven more sessions to remotely configure seven Aruba switches. A typical computer could generate thousands of sessions.

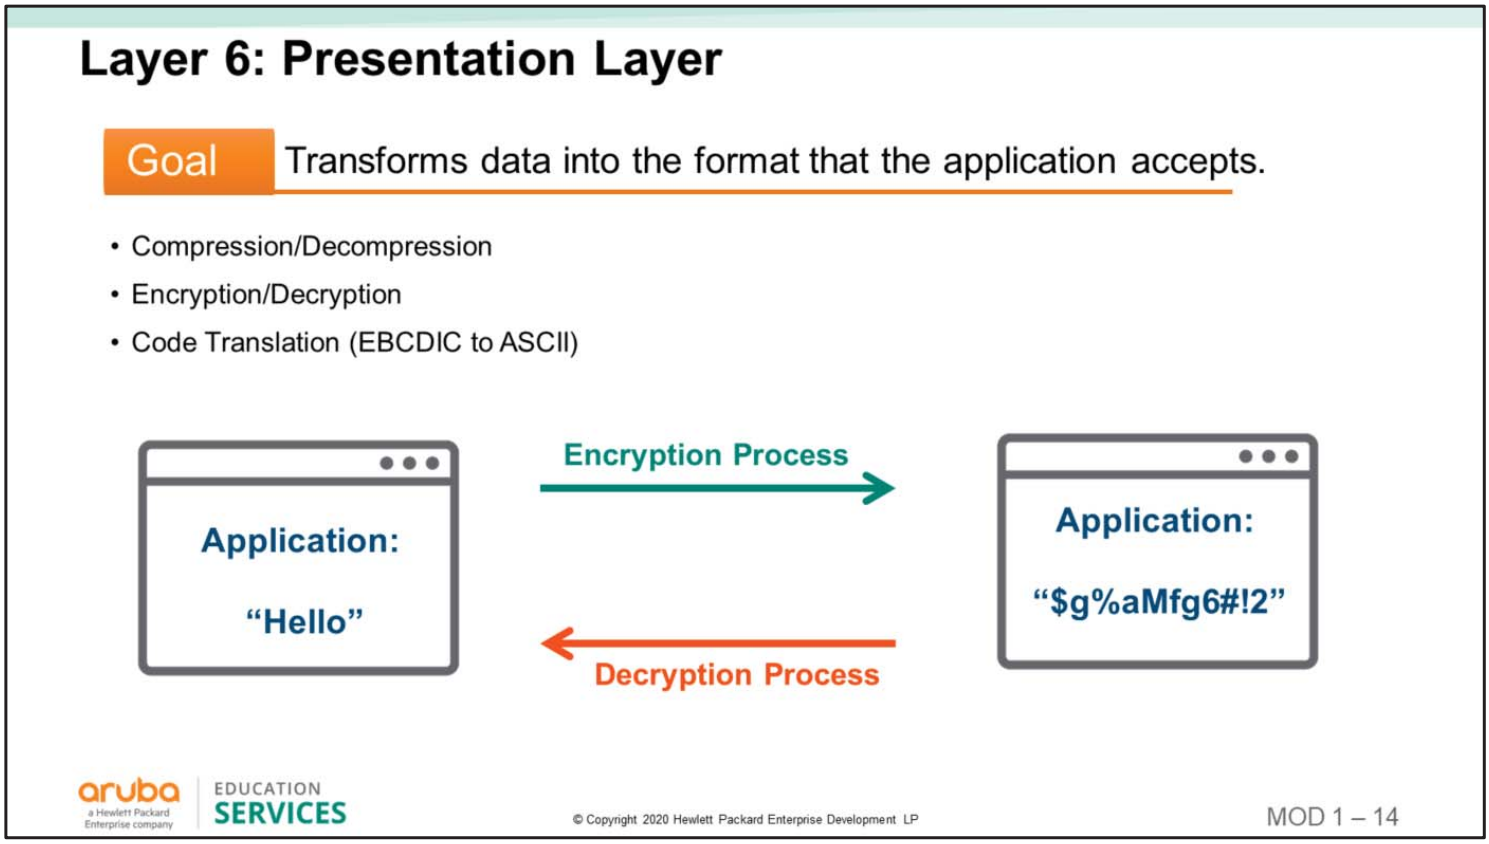

The Presentation Layer transforms data into the formats that the application accepts. Typical processes include:

|Compression/Decompression
|Encryption/Decryption
|Code Translation (EBCDIC to ASCII)

For example, the figure shows how an application passes the clear-text message "Hello" to the Presentation Layer process, which encrypts this message before transmission. This provides confidentiality. If any bad actors or hackers intercept this data, they will not be able to read the message. Of course, upon receipt of an encrypted message, only the intended receiver has the correct digital keys to decrypt the data.

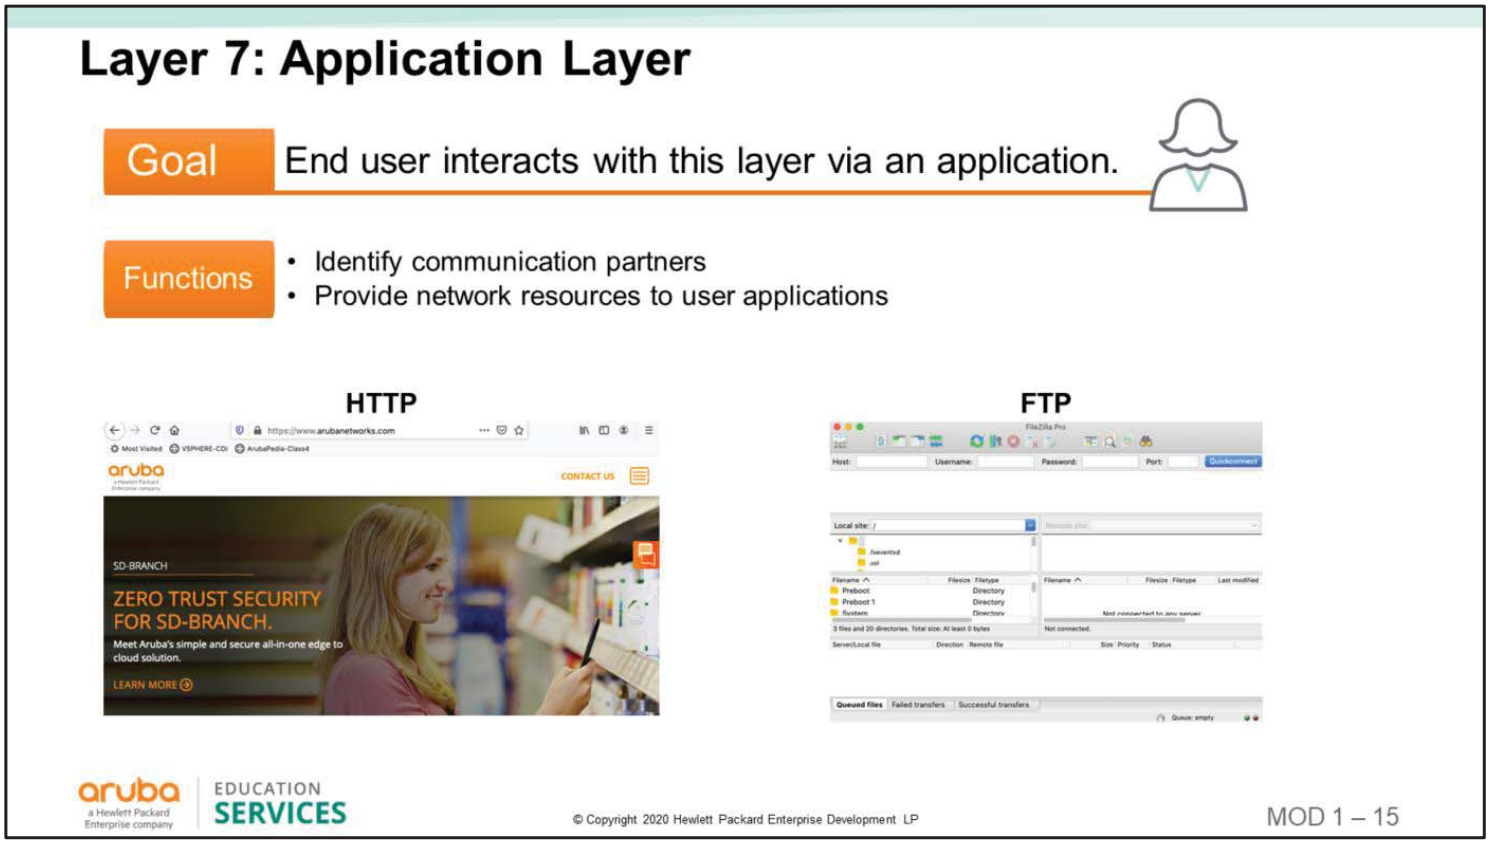

The Application Layer is the closest to the end user, which means that both the OSI Application Layer and the user interact directly with the software application.

|Application Layer functions include:

Identifying Communication Partners. The application layer determines the identity and availability of communication partners for an application with data to transmit.
Provide network resources. This layer provides network services to user applications, such as file transfer, email, video conferencing and many others.

Some examples of application layer include:

Hypertext Transfer Protocol (HTTP), which relies on TCP for transport
File Transfer Protocol (FTP), which also relies on TCP
Domain Name System (DNS), which typically uses UDP, but sometimes also uses TCP
Trivial File Transfer Protocol (TFTP), Which uses UDP as its transport mechanism

Question #1

Computer Networks

Which of the following are concepts or technologies that are **specific** to a LAN?

a. Ethernet

b. Wi-Fi

c. Multi-Protocol Label Switching (MPLS)

d. The Transmission Control Protocol (TCP)

e. Modulation and demodulation

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 16

Answers:
a. Ethernet
b. Wi-Fi

Question #2

The OSI Model

Which of the following are aspects of Layer 4 of the OSI model?

a. TCP
b. IP
c. UDP
d. Segmentation
e. Session Management

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 17

Answers:
a. TCP
c. UDP
d. Segmentation

Encapsulation

Layer Headers and Encapsulation process

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

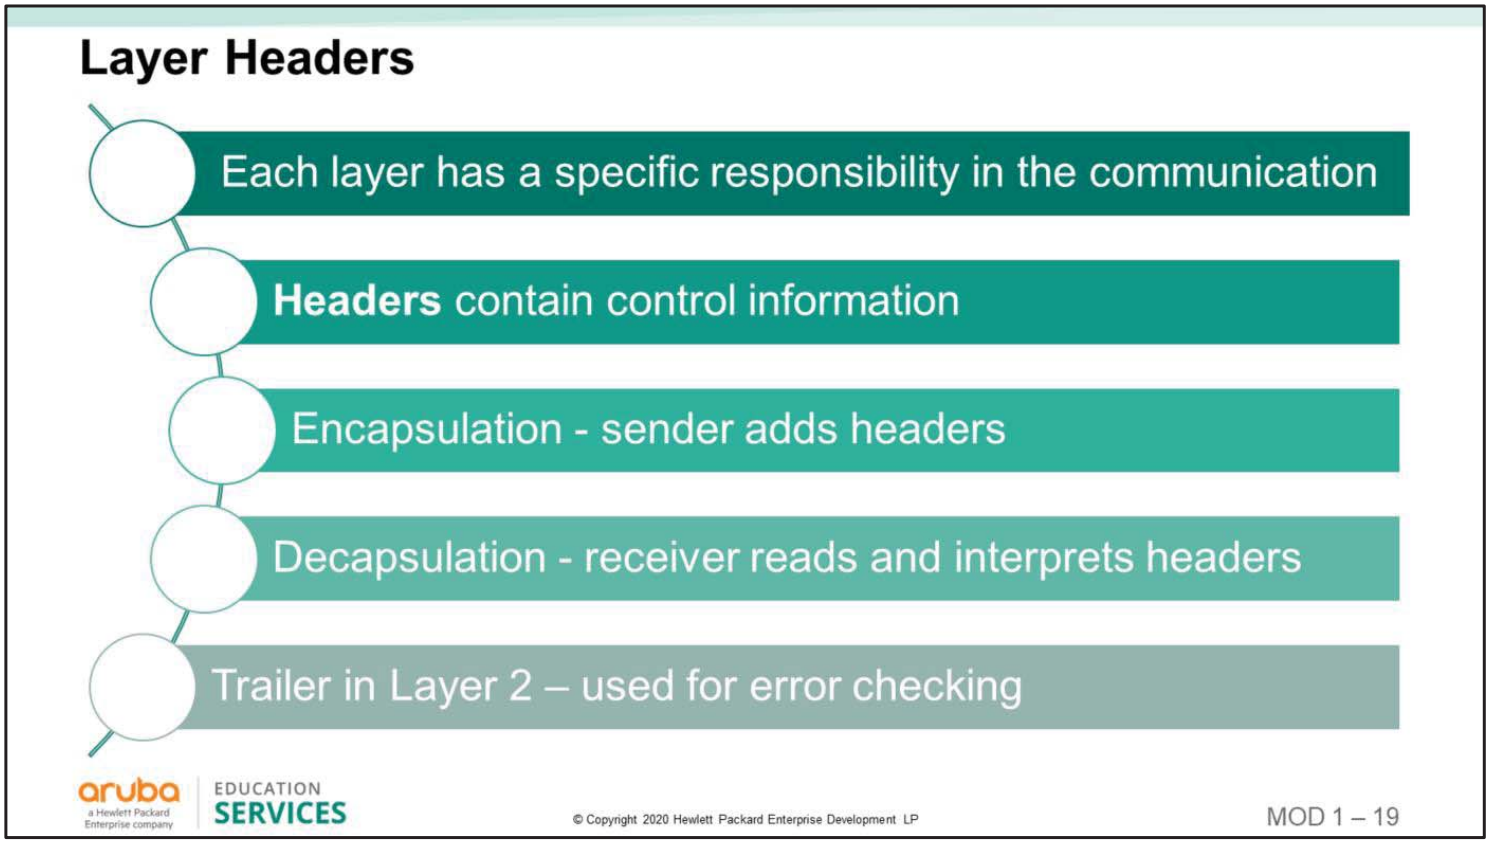

In the OSI model each Layer has a specific responsibility during network communications. In the computing world, devices use headers to establish a communication exchange and to control information on a particular layer, or the layers above it.

Note: A header that is generated on a specific layer by the sender can only be read at the same level on the receiver side.

Encapsulation. Is the process where each OSI Layer adds a header. This process is always done by the sender device.
Decapsulation. Is the process to read and interpret the header information. This process is always done by the receiver device.

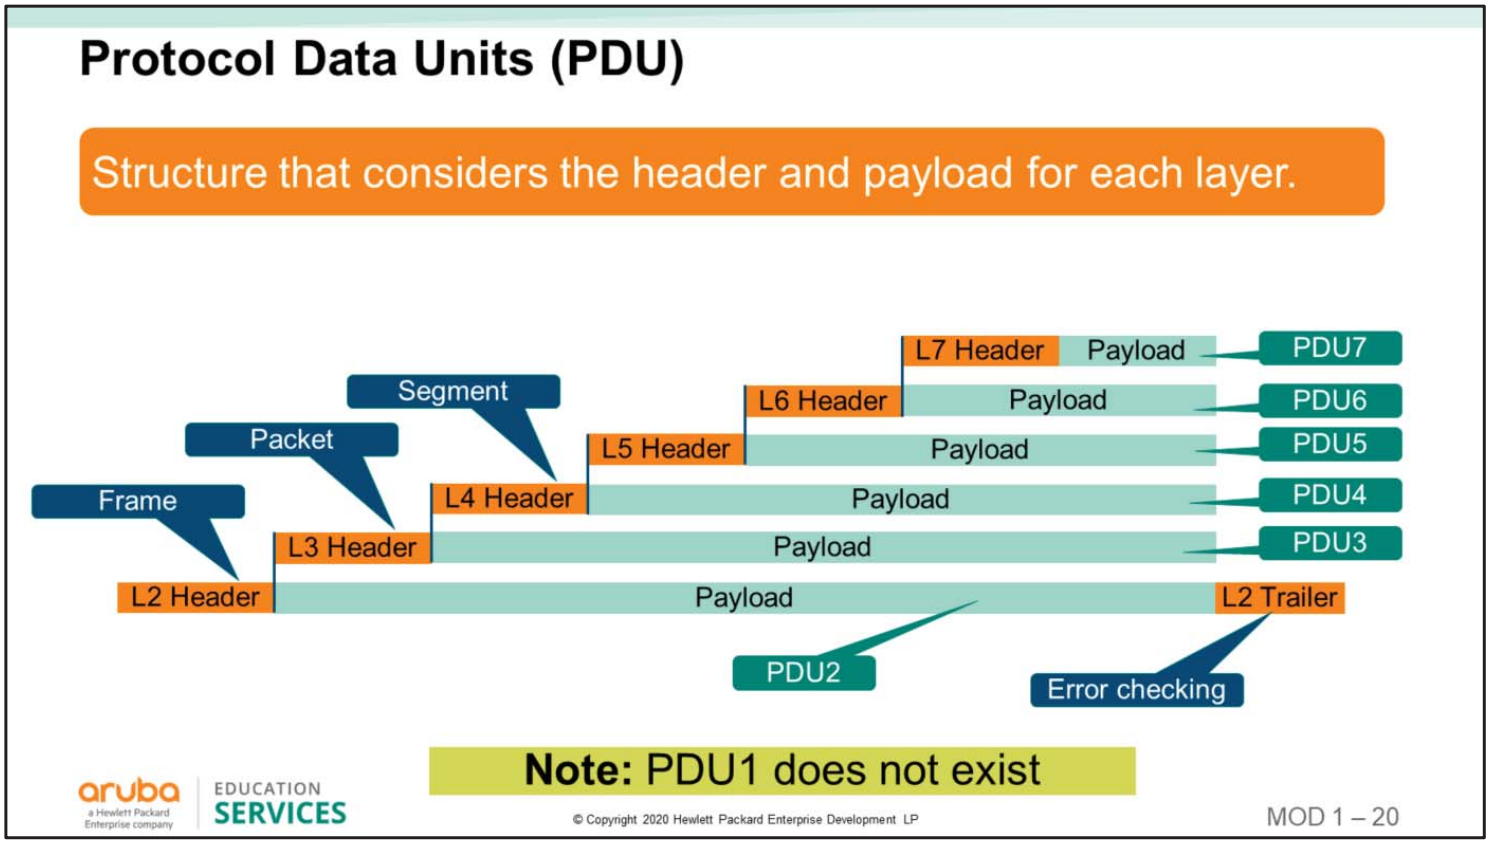

The OSI model introduces the concept of a Protocol Data Unit (PDU). This is simply a structure that considers the header and payload or data for each Layer.
|The following table summarizes the PDUs from Layer 7 to Layer 2.

Note: PDU1 does not exist, remember that Layer 1 refers to signals that cross the media.

There are three key terms related to PDUs that you should know:

|Segment. Refers to the encapsulation that is done in Layer 4. A segment is equivalent to PDU4. So, you might hear networking people speak about a TCP segment or a UDP segment.
|Packet. Refers to the encapsulation that is done in Layer 3. A packet is equivalent to PDU3. You might talk about Layer 3 IP packets.
|Frame. Refers to the encapsulation that is done in Layer 2. A frame is equivalent to PDU2. You might talk about Ethernet frames or Wi-Fi frames.

You might notice that PDU2 – a Layer 2 Frame - not only includes a Header but also a Trailer that is appended after data, labeled “L2T” in the figure.
|The trailer is typically is used to detect errors during the transmission of the message. You recently learned about this during the discussion of Layer 2 of the OSI model. Layer 2 protocols like Ethernet and Wi-Fi add a Trailer, often labeled “Cyclic Redundancy Check” (CRC) or perhaps as a “Frame Check Sequence” (FCS).

Physical Media

Copper, Fiber Optic and Radio Frequency

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

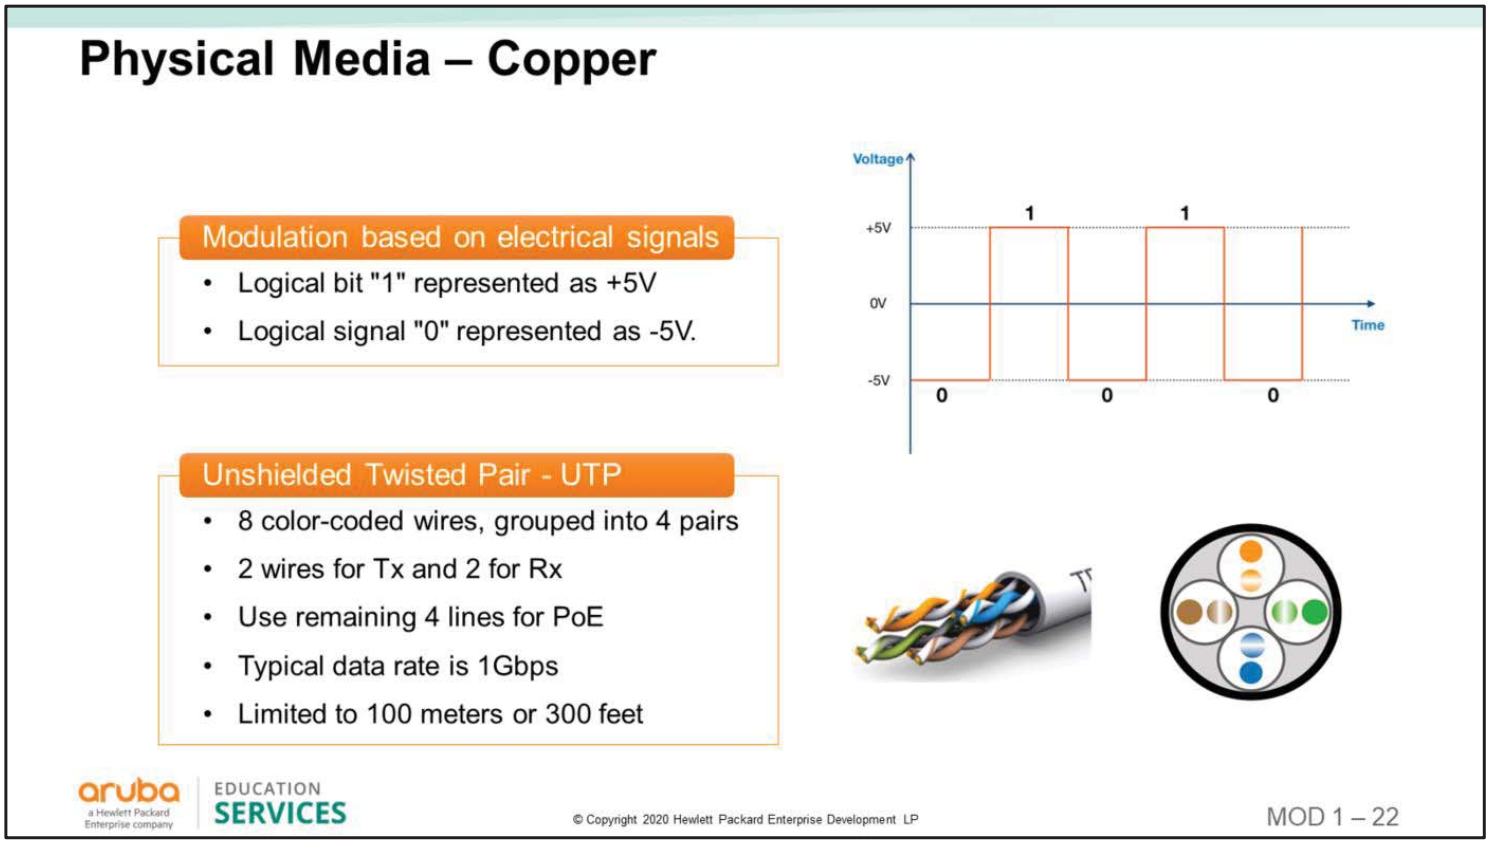

Computing devices might use different media to transmit information, and each media type has different characteristics. Recall that Fiber Optic media modulates light waves, Wi-Fi modulates Radio Waves, and Copper-based media modulates electrical properties like voltage.

Copper cables
You recently learned about how Layer 1 processes modulate or change some aspect of a signal, in order to represent binary data. In the example shown here, +5 volts represents a logical binary 1, while -5 volts represents a logical binary 0.

|The typical cooper cable that is used to transmit digital data in a network is called Twisted Pair. With this media, wire pairs are twisted, to reduce electromagnetic radiation and interference. The most common type of Twisted pair is Unshielded Twisted Pair (UTP). Other variations also exist, such as shielded or Foiled Twisted Pair (FTP) - commonly deployed to provide superior protection from external electromagnetic interference (EMI). This is useful in highly sensitive environments, or those with high levels of EMI.

UTP cabling contains 8 color-coded wires, grouped into 4 pairs. Two wires are used for transmission (Tx) and two are used for reception (Rx). The remaining 4 lines are used to power some devices such as telephones or cameras, using a feature called Power over Ethernet (PoE). The typical connector used with UTP is an RJ-45 connector.

To maintain a data rates of up to 1Gbps, the maximum length of UTP cable cannot be more than 100 meters (300 feet).

Physical Media – Fiber Optic

Modulation based on light signals

- Light travels down the fiber center (core)
- Cladding surrounds the core, traps light using the Total Internal Reflection technique

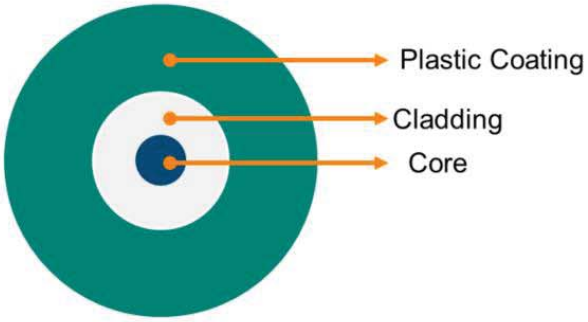

Physical properties

- Distances from 100 meters (300 feet) to 40 Km
- Data rates: 1,10,25,40,50 and 100Gbps
- Types: Multimode (MM) and Single Mode (SM)

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 23

Fiber Optic. Built from glass or plastic fiber, this media uses light signals to transmit and receive data. Light is guided down the center of the fiber, which is called core. The core is surrounded by an optical material called the cladding that traps the light in the core using a technique called total internal refection.

|Fiber optic cables can interconnect devices that are separated by much longer distances than Ethernet UTP's 100 meters, and with higher data rates. Distances and speeds depend on quality, type of fiber, and transceiver type. Common data rates are 1Gbps, 10Gbps, 25Gbps 40Gbps, 50Gbps or even 100Gbps.

Note: A fiber optic transceiver is an optical module installed in the computing device. It is responsible for modulating and demodulating light signals.

Fiber optic is categorized into two main groups - Multimode (MM) and Single mode (SM) fiber. MM is typically less expensive, for relatively shorter distances. SM is often more expensive but can often support very long distances.

There are several standard fiber connectors. However the typical one that is used is LC connector. The image shows an example of this.

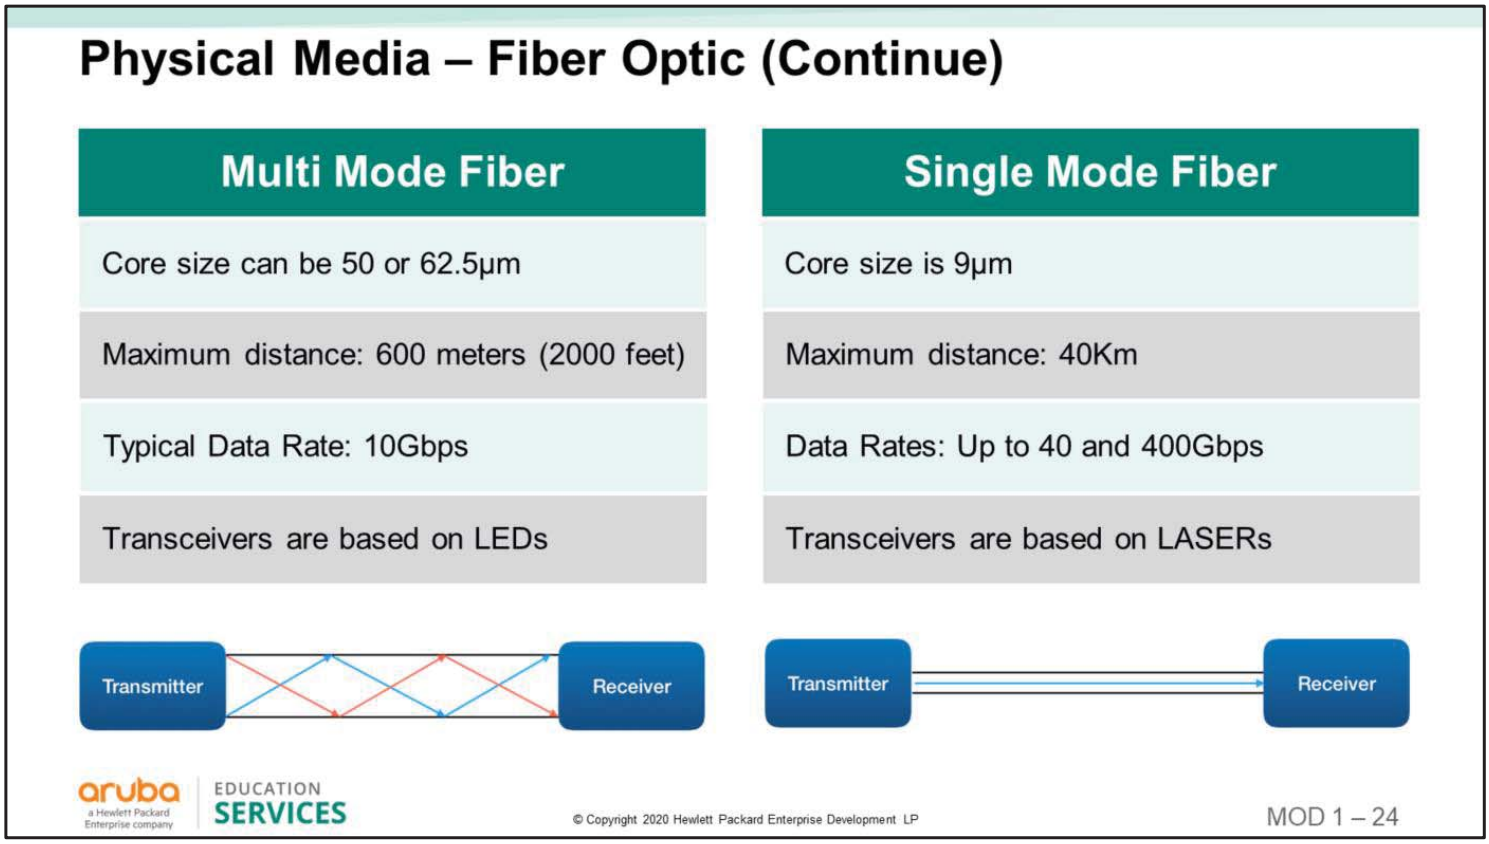

The figure shows specifications for Multimode (MM) and Single mode (SM) fiber.

|Multimode (MM). The core size for this fiber is 50 or 62.5 micrometers (µm). This core size allows greater light-gathering capacity and facilitates the use of less expensive transceivers. Typical distances are up to 600 meters (2000 feet), With typical data rates of up to 10Gbps.

Typically, fiber optic data sheets for Multimode Fiber include terms like 50/125 or 65.5/125, where the first number (50 or 62.5) is the diameter of the core and 125 represents the diameter of the cladding.

Multi-mode fiber with 50µm has a faster light transmission but with shorter distance.

|Single Mode (SM). The core size is only 9µm and carries light directly down the fiber, as a result, light reflection created during light transmission decreases. This lowers attenuation (loss of signal strength) and allows the signal to successfully travel over longer distances. Usually, this fiber is more suitable to be used to interconnect devices using higher data rates such as 40Gbps or even 400Gbps.

As you might imagine 9/125 refers to the fact that the core is 9µm and the cladding is still 125µm.

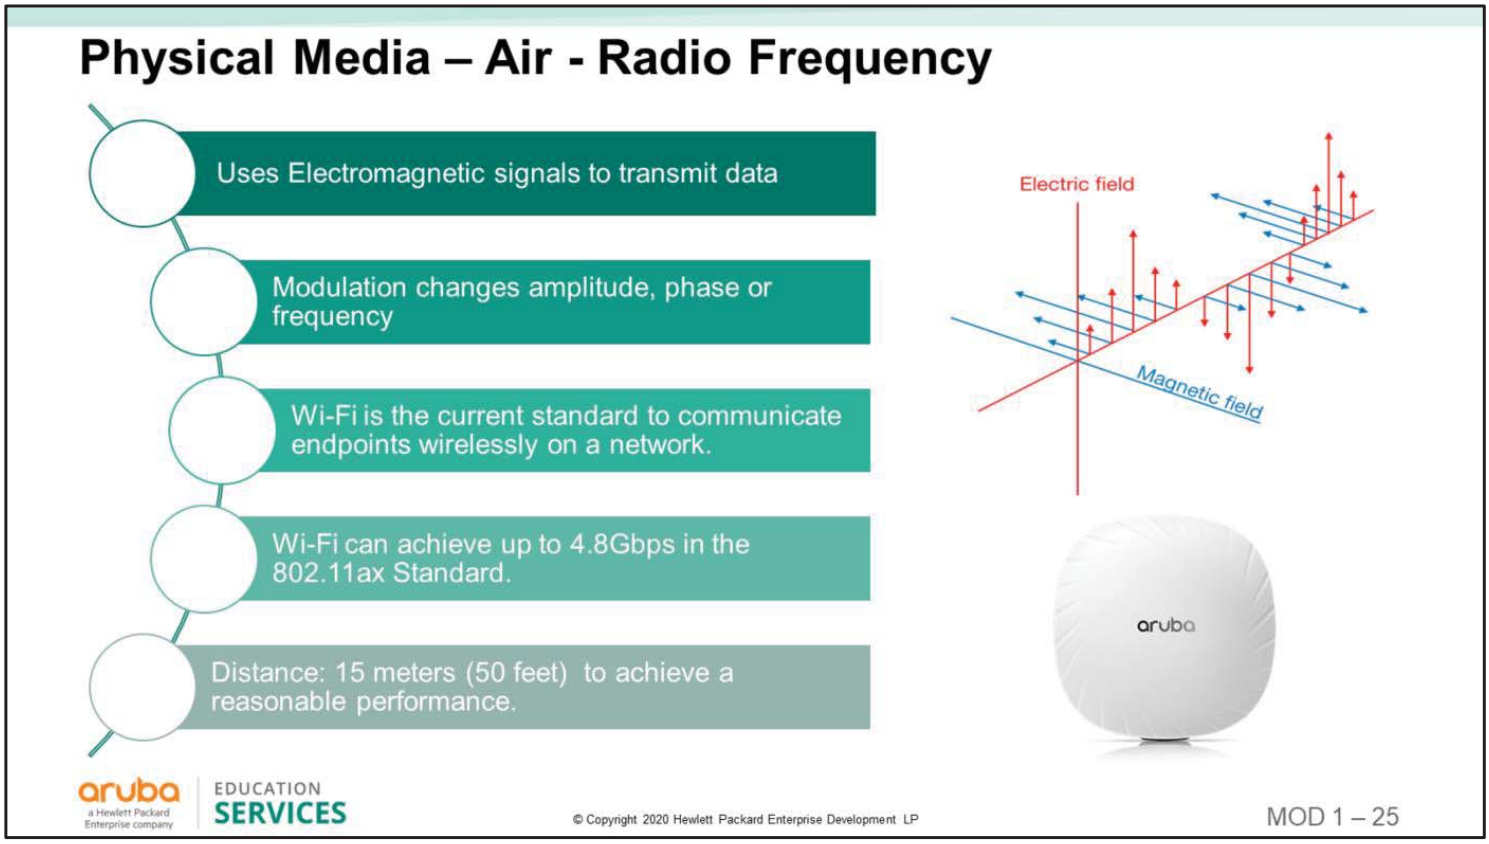

Air (Radio Frequency). To transmit data, this media type modifies some aspects of an electromagnetic signal, such as amplitude, phase, or frequency. These modulated radio waves are then transmitted through the air. This wireless transmission technology liberates users from stationary, wired connections and allows them to roam around.

Radio Frequency (RF) can be used to transmit radio, TV, and satellite signals. In the computer world, Wi-Fi is the technology that permits wireless endpoint communication. This technology works in two radio bands - 2.4GHz and 5GHz. Data rates using Wi-Fi can be up to 4.8 Gbps using the IEEE 802.11ax standard.

Compared to copper and fiber optic media standards-based Wi-Fi transmission supports the shortest distances. Depending on many factors, including device configuration, number and type of obstacles, required performance, and more – maximum distances between devices is typically between 15 and 30 meters (50 to 100 feet).

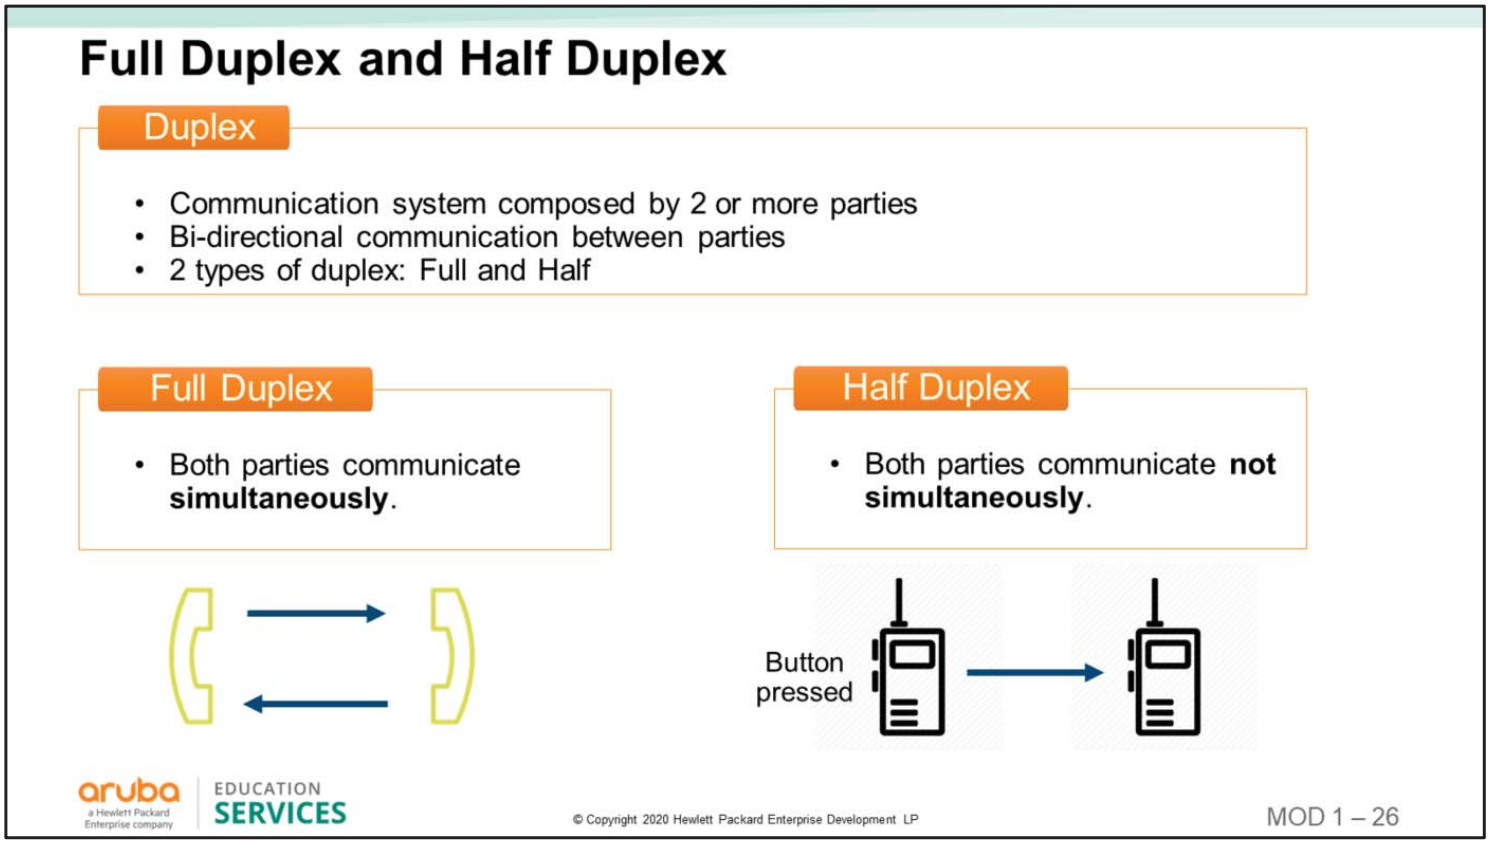

A Duplex communication system is a system composed by two or more connected parties that can communicate with one another in both directions. There are two types of duplex communication systems: Full duplex (FDX) and Half duplex (HDX)

Full Duplex: Both parties can communicate with each other simultaneously. An example of full-duplex is a telephone; parties in both ends can speak and can be heard by the other party simultaneously.

Half Duplex: Both parties can communicate with each other but not simultaneously; the communication is one direction at the time. An example of half-duplex is a walkie-talkie; on this communication each person must press a ”push-to-talk” button when they want to talk, when the button is pressed then the user cannot hear the remote person. To listen the button has to be released.

Question #3

Physical media

Under which circumstances is it most appropriate to use Single Mode fiber optic mode?

a. When you need to connect two devices in a single equipment rack
b. When you must connect two buildings that are 10km apart
c. When you are concerned about electromagnetic interference
d. When you want to use UDP instead of TCP

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 27

Answers

b
c

Types of Traffic

Unicast, Multicast and Broadcast

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

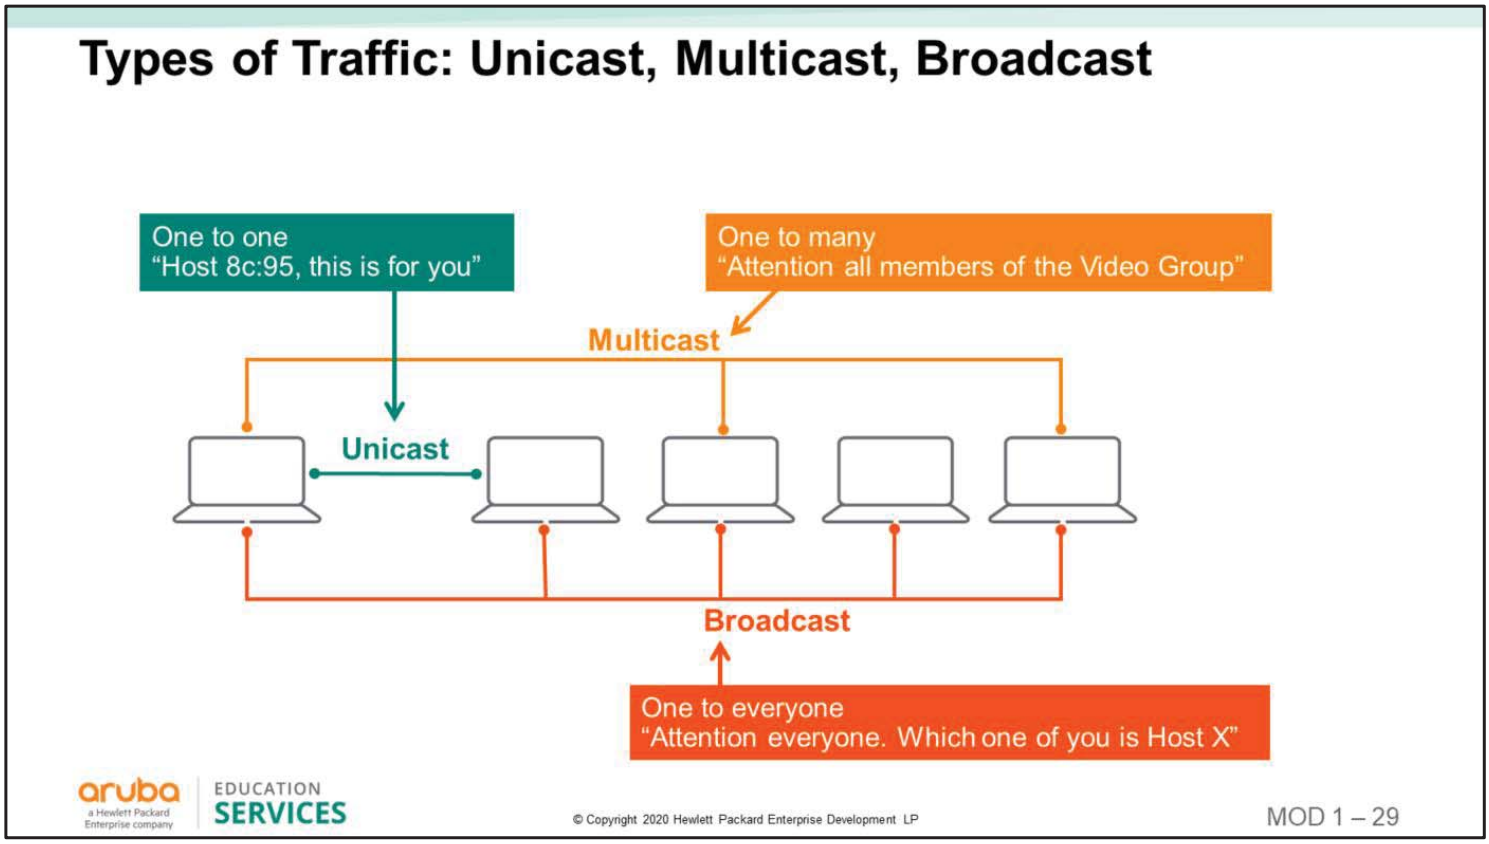

Computing communications can be classified in three types:

Unicast

This traffic refers to one-to-one communication. One transmitter and one receiver. Imagine that there are several learners in a classroom. Bob is the instructor. He calls out, "Alice, I have a message for you from the front desk". This message came from one source – Bob and is destined to a single destination – Alice.
Similarly, when a PC needs to transfer a file to a server, the two devices use unicast communications.

Multicast

This traffic refers to one-to-many communication - one transmitter and multiple receivers. In our classroom analogy, suppose that lunch has been brought in for the learners. Bob may call out, "All vegetarians can find their meals on the green table". The message came from one source – Bob. The message is destined for several people in the room – all of those with a vegetarian diet.
A common networking example is video streaming, where a video source (transmitter) sends a video stream to multiple devices that are capable and interested in receiving that information – perhaps on their PC, tablet, or smartphone.

Note. Multicast traffic can include many devices but not all the devices in the network.

Broadcast

This type of traffic refers to the communication one-to-all. In our classroom analogy, Bob calls out, “Attention everyone. It is break time, and there are free doughnuts in the lobby”. The special word “everyone” means that all people in the classroom are intended to receive this message.
Similarly, there are special network addresses that all stations will receive and process. At layer 2, this is the MAC address FF:FF:FF:FF:FF:FF. At Layer 3, this is the IP address 255.255.255.255.
This helps a particular computing device to discover others in a specific network.

Numerical Systems

Binary, Decimal and Hexadecimal

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

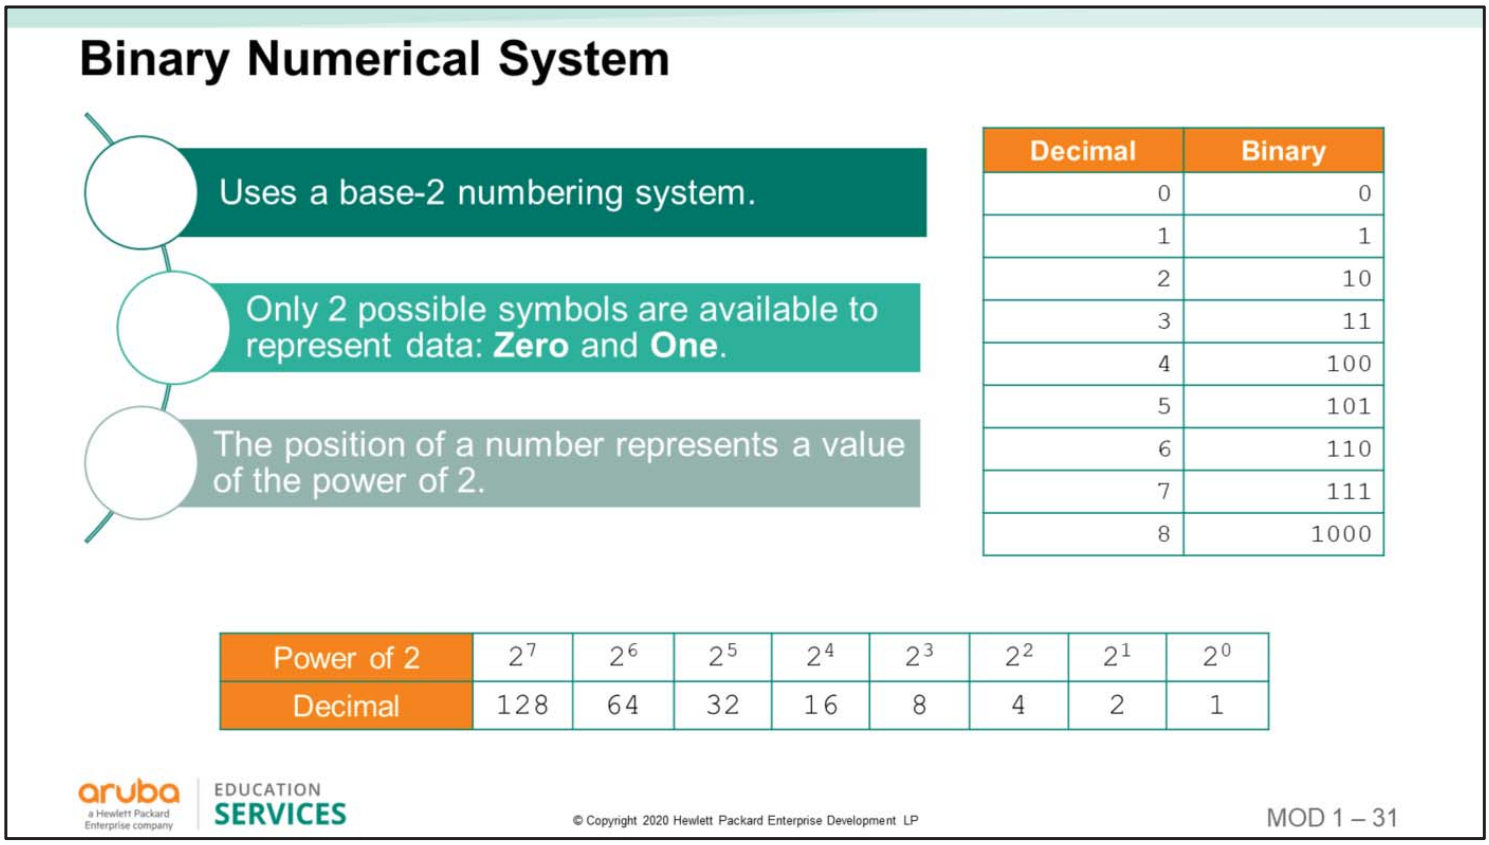

Familiarization with the binary numerical systems is key to understand how computing systems process and communicate information.

The binary system uses a base-2 numbering system, this means that there are only 2 possible symbols available to represent data: zero and one.
In the binary system the position of a number represents a value of the power of 2. The table below shows the first 8 positions and the decimal number associated to each.

Note. To avoid any confusion this text will use an index after a number to indicate the base numerical system, for example 10010_2 will represent the number 100 (one hundred) on the decimal system.

Counting in binary

The table shows the first eight decimal numbers and their representation in binary. The first two values only require one digit (2^0), and so zero and one are the same for both numerical systems. Please notice that for decimal number 2 in binary you must add a new number to the left indicating 2^1 and the number to the right 2^0. Let's make a quick comparison of binary and decimal numbering systems.

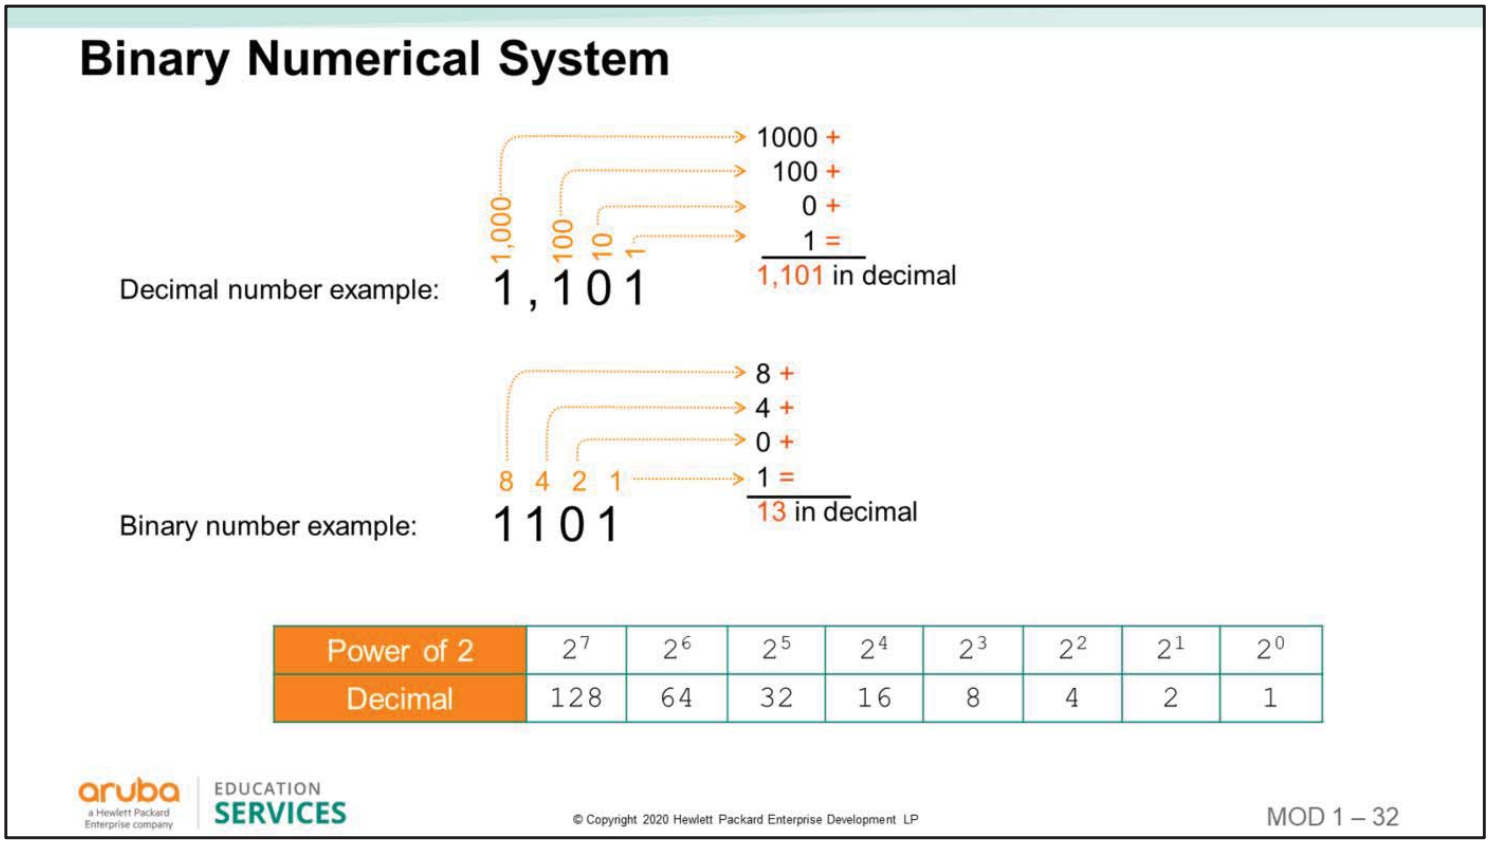

Consider the decimal number 1,101, as shown in the figure. You know that the right-most digit (3) is in the “1’s” column, then 0 is in the 10’s column, the 1 is in the 100’s column, and the 1 is in the 1,000 column. Therefore, the first 1 (left to right) does not merely represent a quantity of 9, it represents a quantity of one thousand – 1 X 1000 = 1,000. Similarly 1 x 100 = 100. 0 x 10 = 0, and 1 x 1=1.
You know this already, right? And 1000 + 100 + 0 + 1 = 1,101

Binary works exactly the same way. Yes, the numbering system changed from base 10 to Base 2, but the fundamental rules never change. The only difference is that instead of 1000, 100, 10, and 1 – the columns are 8, 4, 2, and 1.

For example, consider the binary number 1101. The left-most number is 1, and it is in the 8’s column. 1 x 8 = 8. The next number is 4, and 1 x 4 = 4. The next number is 0, and 0 x 2 = 0. The right-most number is 1, and 1 x 1 = 1. Now add them up, just like before. 8 + 4 +1 =13

Let’s take this a bit further, using a more methodical step-by-step process.

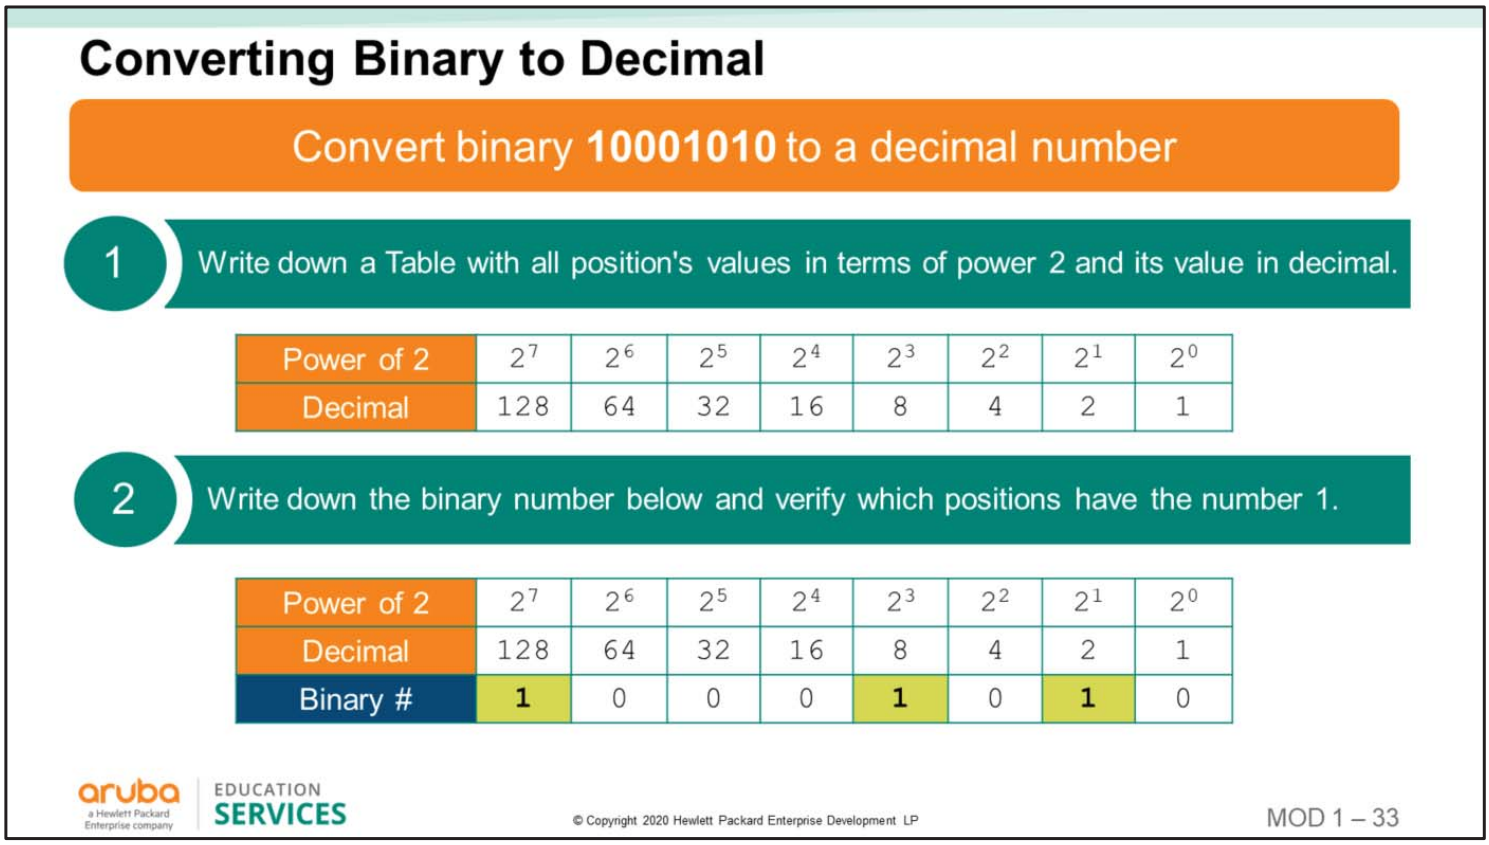

There are several methods to convert a binary number into decimal, however the comparison method is the easiest to learn. This method simply uses the position value of each number (remember that in binary system each position is a power of 2) and simply sums all the values where the binary number is set to 1.

As an example, let's convert the binary number 100010102.

Step 1. Write down a table with all the positions values in terms of power of 2 and its value in decimal. This example has 8 binary characters, so you need to write down the first 8 values.
Step 2. Write down the binary number below and verify which positions have the number 1.

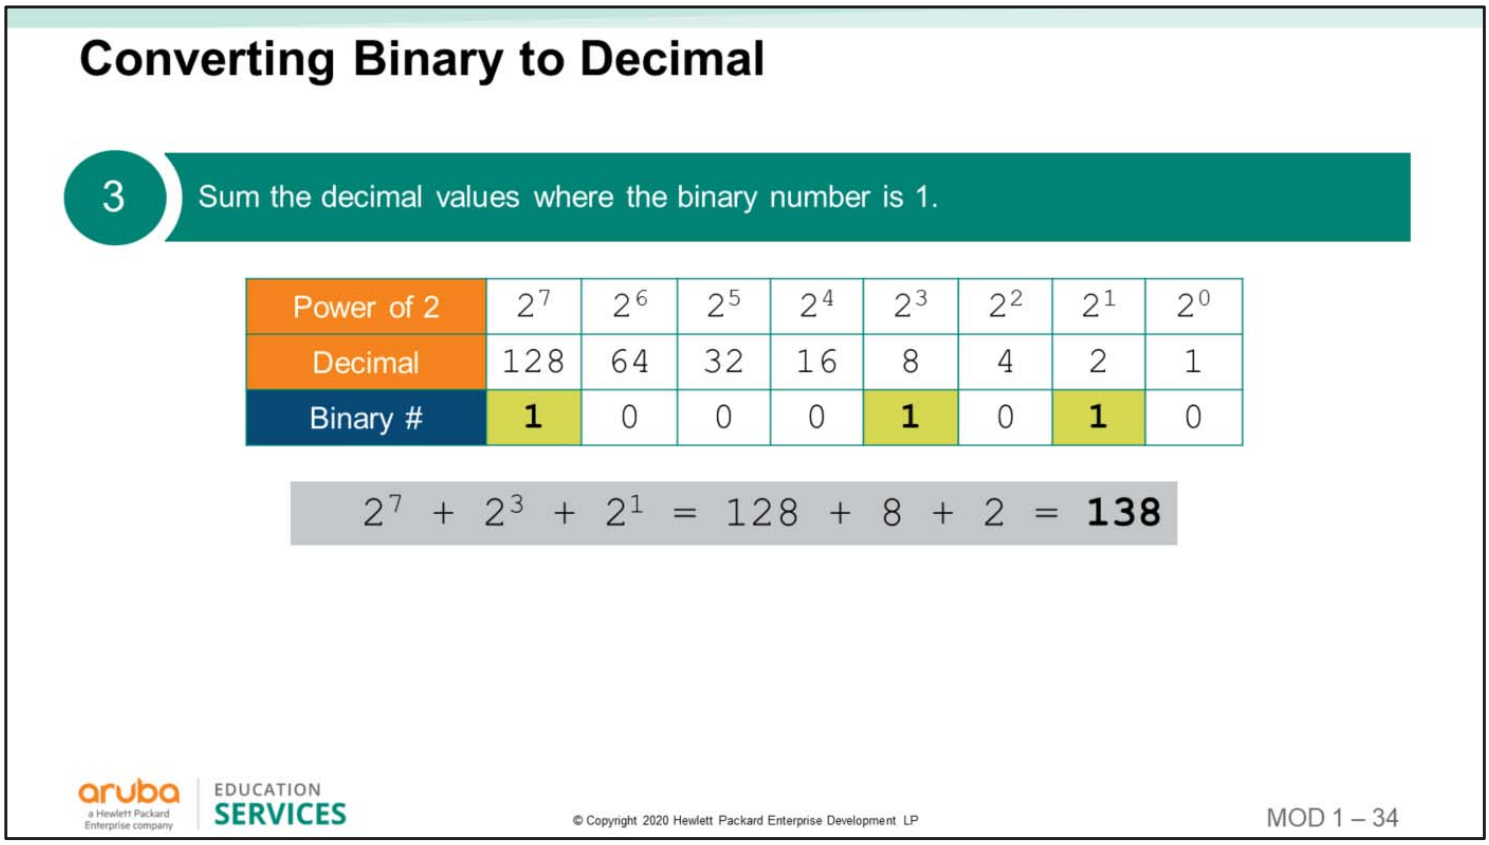

Step 3. Sum the decimal values where the binary number is 1.

27 + 23 + 21 = 128 + 8 + 2 = 138

As a result, you can conclude that 10001010 in binary is equivalent to 138 in decimal.

In order to practice please try to determine the decimal numbers associated to the following binary numbers:

101011102
010001012
111111102
100000012

Answers:

17410
6910
25410
12910

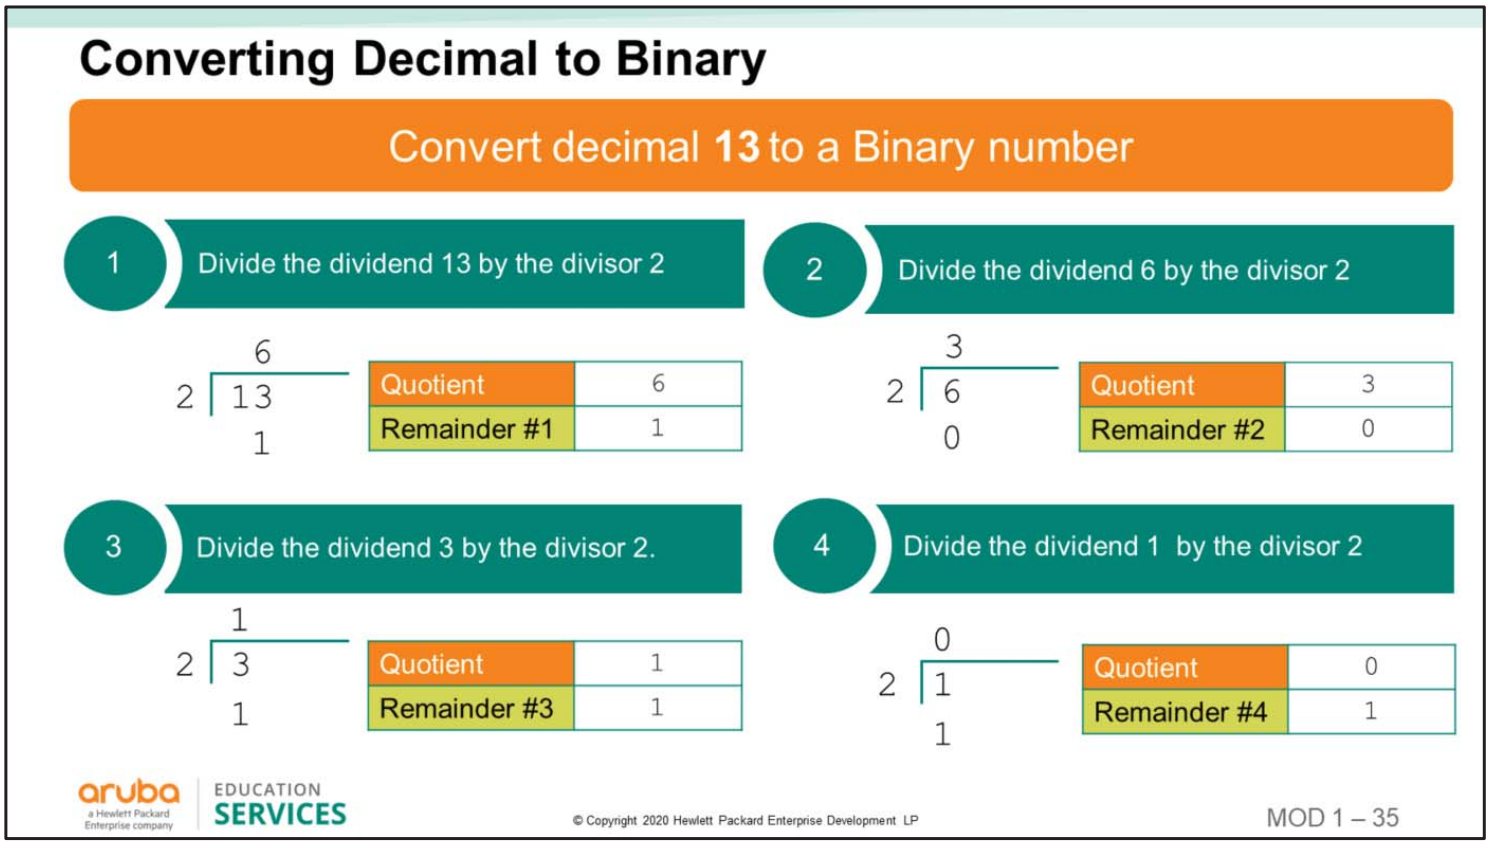

The conversion from decimal to binary is not based on a sum but in a repeated divide-by-2 process. Start by dividing the decimal number by 2. Notice of the quotient and the remainder. Continue dividing the quotient by 2 until you get a quotient of zero, then just write out the remainders in the reverse order.

Consider the following example, where the number 1310 is converted into its binary representation.

|Step 1. Divide the dividend 13 by the divisor 2. This operation will have a result of 6 as the quotient and 1 as the remainder.
|Step 2. Take the quotient of the previous operation (6) and divided by 2. This operation will have a result of 3 as the quotient and 0 as the remainder.
Step 3. Take the quotient of the previous operation (3) and divided by 2. This operation will have a result of 1 as the quotient and 1 as the remainder.
Step 4. Take the quotient of the previous operation (1) and divided by 2. This operation will have a result of 0 as the quotient and 1 as the remainder.

Note. On step 4 you know that the quotient of the division is, 0.5 and the Remainder 0, however the process only focus on the integer part of the quotient in this case is 0 and on the last Remainder before the operation is done, in this case the last Remainder is 1.

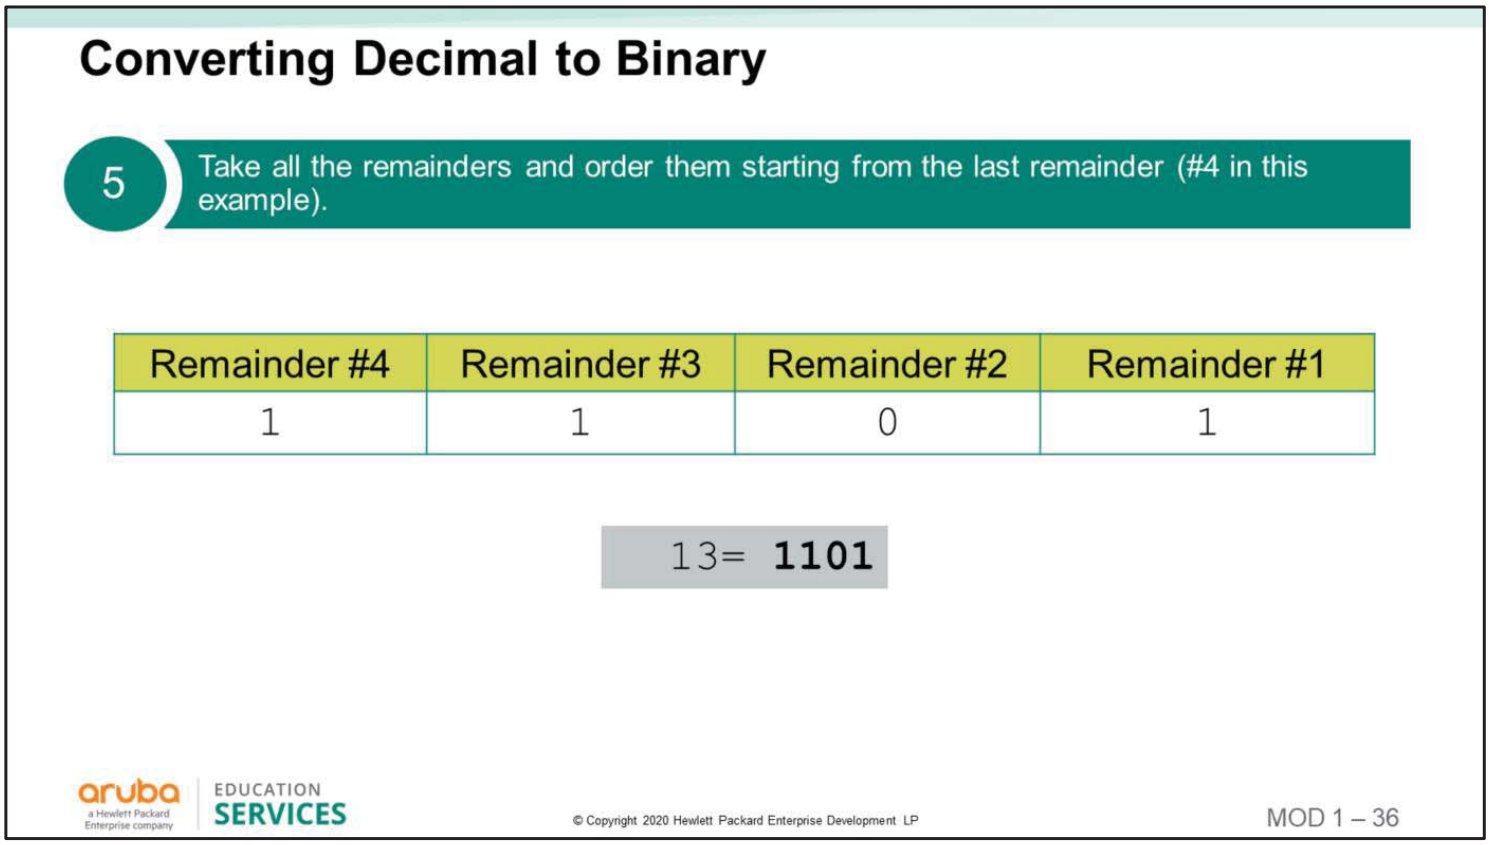

Step 5. Take all the Remainders and order them starting from the last Remainder (Remainder #4)

This implies that the binary number for 1310 is 1101 in binary.

Let's briefly review another approach.

Converting Decimal to Binary – Alternate method

- You wish to convert decimal "13" into binary.
- Is 13 greater than or equal to 128? Then place a 0 in the 128's column.
- Is 13 greater than or equal to 64, 32, and 16? Then place 0's in those columns.

Power of 2	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
Decimal	128	64	32	16	8	4	2	1

0 0 0 0 1 1 0 1 = 13

- Is 13 greater than or equal to 8? Yes, then place a 1 in the 8's column.
- Add the following column (8+4), is 13 greater than or equal to 12? Yes, then place a 1 in the 4's column
- Add the following column (12+2), is 13 greater than or equal to 14? No, then place a 0 in the 2's column
- Add the following column (12+1), is 13 greater than or equal to 13? Yes, then place a 1 in the 1's column

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 37

As before, you wish to convert decimal 13 to binary.
Look at the column values for the binary number system and compare. Is 13 greater than or less than 128? Less than, therefore you must place a 0 in the 128's column.
13 is also less than 64, 32, and 16, so those columns must all have a 0, as shown in the figure.

Now, 13 is obviously greater than 8, so you must place a 1 in the 8's column. We're still not to 13 yet, so keep going.
Add the next lowest column: 8+4 = 12. You have not reached 13, so keep going.
12+2 = 14 and this value is greater than 13 then place a 0 on the 2's column
12+1 = 13, which is equal to the number you are finding, then place a 1 on the 1's column.

So 13 in decimal = 00001101 in binary.
Let's try one more example.

Converting Decimal to Binary – Compare and Add

- You wish to convert decimal "187" into binary.
- Is 187 greater than 128? Then place a 1 in the 128 column.
- Now add 128+64= 192. That's too high! Put a 0 in the 64 column.

Power of 2	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
Decimal	128	64	32	16	8	4	2	1

1 0 1 1 1 0 1 1 = 187

- Add 128+32 = 160. That's lower than 187, so put a 1 in that column.
- Add 160+16 = 176. Still lower, so put a 1 in the 16's column.
- 176+8 = 184. This is lower than 187, so put a 1 in the 8's column.
- 184+4 = 188. Put a 0 in that column.
- 184+2+1 = 187 = 10111011 in binary.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 38

Now you wish to convert decimal 187 to binary.
Look at the column values for the binary number system and compare. Is 187 greater than or less than 128? Greater than, so that column gets a 1.
Now add 128+64=192. That's too high, so the 64 column gets a 0. Go to the 32's column.

128+32 = 160. Lower than 187, so put a 1 in the 32's column.
160+16 = 176. Lower than 187, so put a 1 in the 16's column.
176+8 = 184. Lower than 187, so put a 1 in the 8's column.
184+4=188. Too high, so put a 0 in the 4's column.
184+2 = 186...a 1 in the 2's column.
And 186+1 = 187...a 1 in the 1's column.

So, 187 in decimal = 10111011 in binary.
With just a small bit of practice, you will soon have the columns memorized - 128, 64, 32, 16, 8, 4, 2, 1.
With just a bit of experience, you will begin to learn certain patterns, and this will become more intuitive for you.
Let's look at some common time-saving patterns that might give you an intuitive edge and speed your conversion efforts.

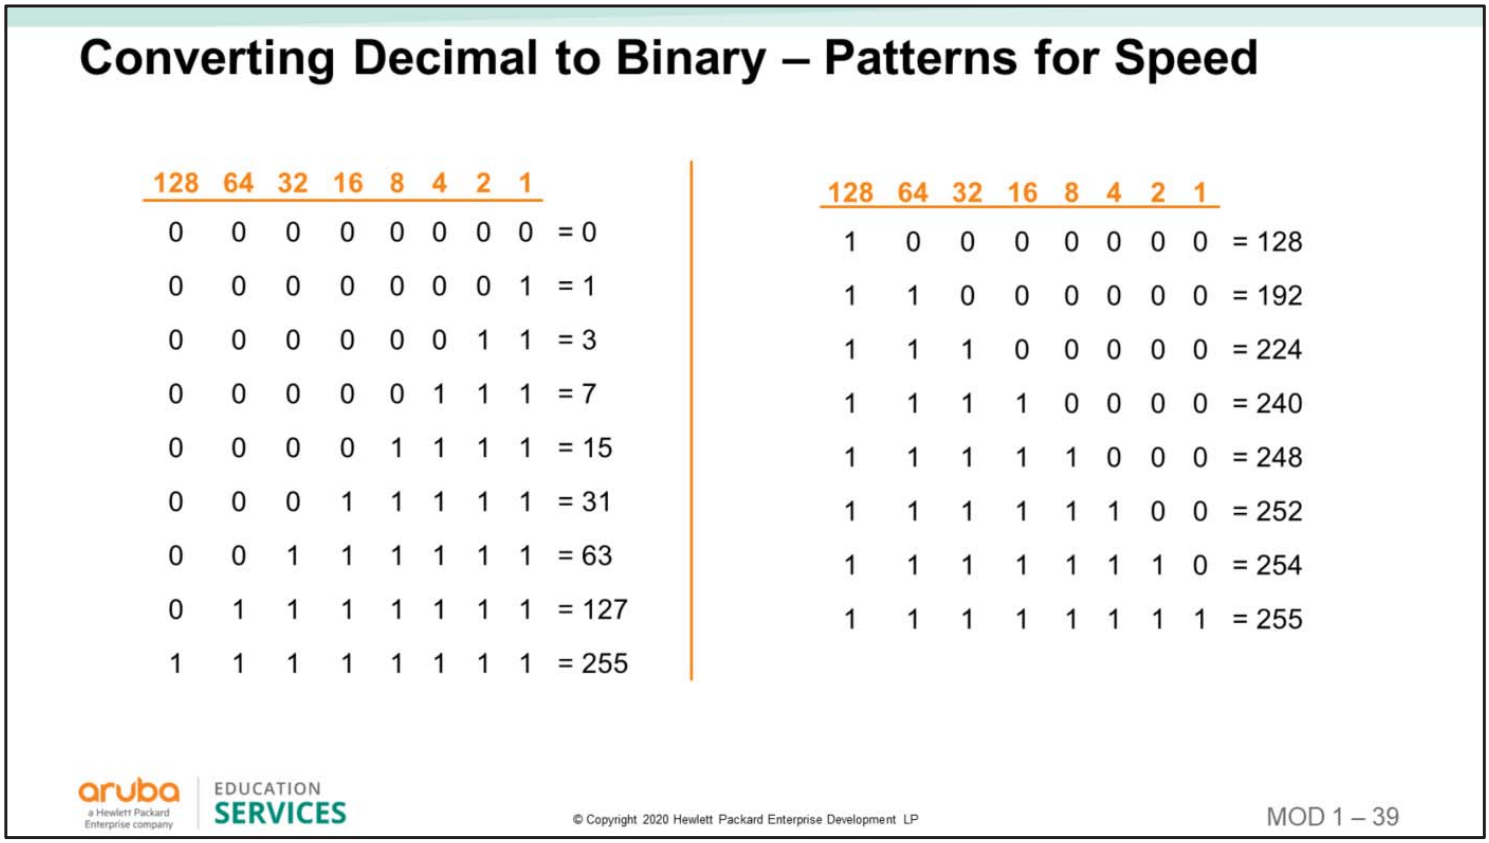

Look at this left-most set of binary to decimal conversions. Notice that it goes from “all zeros” at the top example, to “all 1’s” in the bottom example.
Now look at the third example from the top: 00000011 = 3. Suppose you then add a 1 to the 128 column. You shouldn’t need a lengthy process to convert this to decimal.
You know that the one in the 128’s column represents 128. Then add 3 to that and you get 131. With some practice, this begins to seem intuitive – you can do it in your head.

Consider the next example: 00000111 = 7. Then suppose that say, the 32’s column had a 1 instead of a 0. 32+7=39.

Do you also notice that 3 is one number less than the next column to the right? The 4’s column. And that 7 is one less than the next column to the right – the 8’s column. Do you see the pattern? Once you understand this, you don’t even have to memorize the chart shown in the figure. If you see the binary sequence 00111111, you don’t need to remember that this equals 63. You know that the sequence of binary 1’s end right before the 64’s column, and so this = 63.

Now look at the right-hand column, which shows another set of sequential patterns. Look at the second example from the top: 11000000 = 192. Now suppose you saw another identical example, except that the 4’s column was also set to 1. 192+4 = 196.

With just 30 to 60 minutes of practice, this conversion process will continue to become ever-more intuitive. You will discover other useful patterns on your own, and quickly be able to convert from binary to decimal. Revisit and stay sharp with this skill. This ability to convert numbers is useful for passing Aruba exams. Perhaps more importantly, it will help you as you advance in your training, for things like complex subnet masking, advanced IP address

assignment, advanced route filtering, and other concepts that you will learn about when you take more advanced courses.

Now you know about decimal and binary, let’s learn about another important numbering system for computers and networks – hexadecimal.

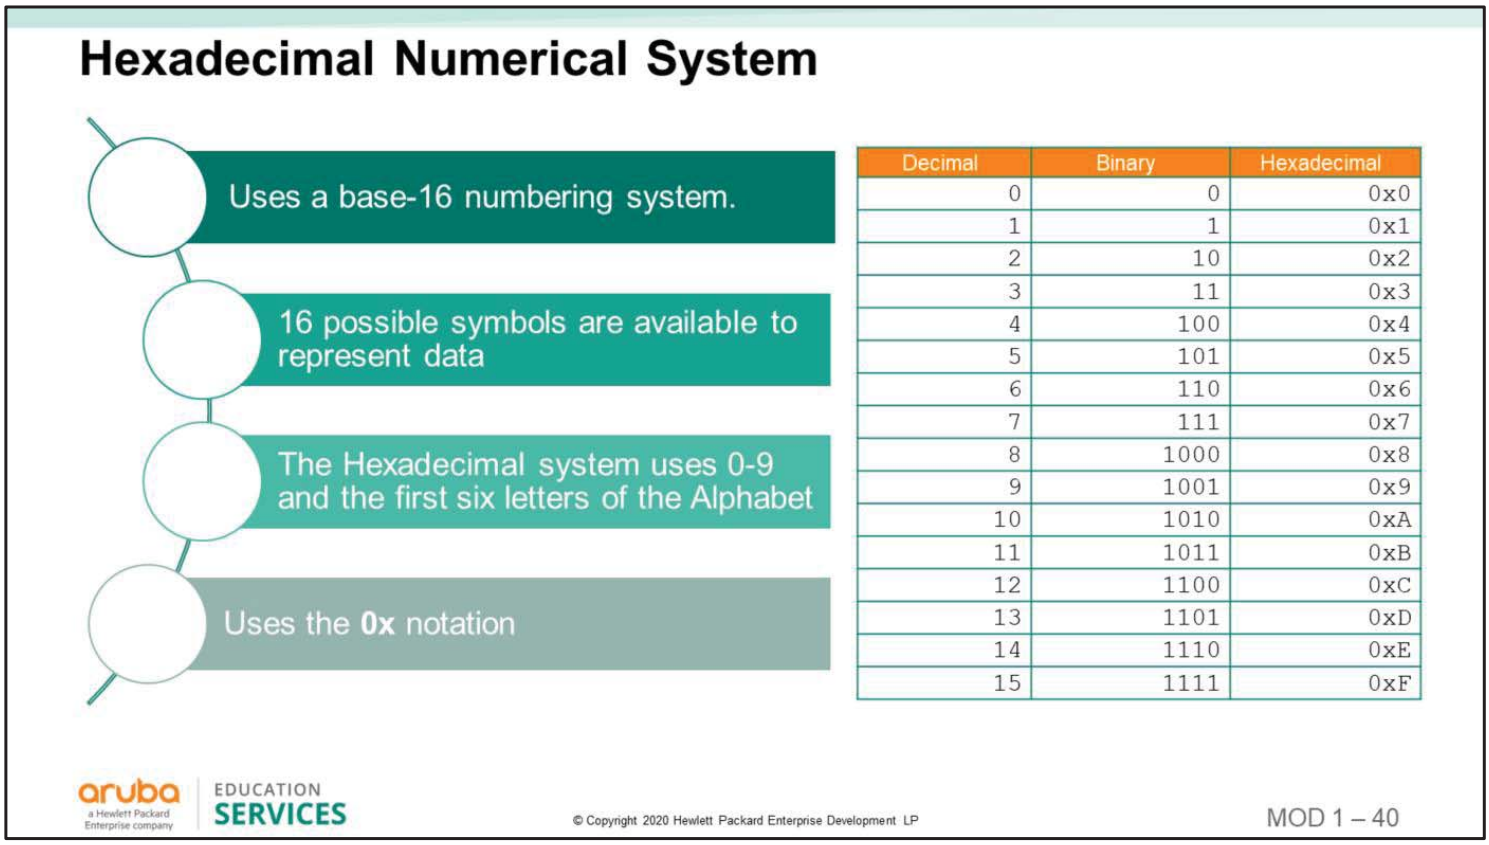

The Hexadecimal system uses a base-16 numbering system, this means that there are 16 possible symbols available to represent data. In this case the hexadecimal system uses numbers (0-9) and the first six letters of the alphabet (A-F).

Note: It is common to represent hexadecimal numbers with a preceding 0x. This notation helps to differentiate hexadecimal numbers from decimal numbers. The hexadecimal number 0x29 is a very different value from the decimal number 29.

In the hexadecimal system the position of a number represents a value of the power of 16.

Converting Binary to Hexadecimal

Four binary digits represent a single hexadecimal number

Binary	Hexadecimal	Binary	Hexadecimal
0000	0x0	1000	0x8
0001	0x1	1001	0x9
0010	0x2	1010	0xA
0011	0x3	1011	0xB
0100	0x4	1100	0xC
0101	0x5	1101	0xD
0110	0x6	1110	0xE
0111	0x7	1111	0xF

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 1 – 41

The conversion between Binary to Hexadecimal is simple. Just know that four binary digits represent a single hexadecimal number. Based on this property, the conversion process is just a substitution, as shown in the figure. Let's look at this in another way.

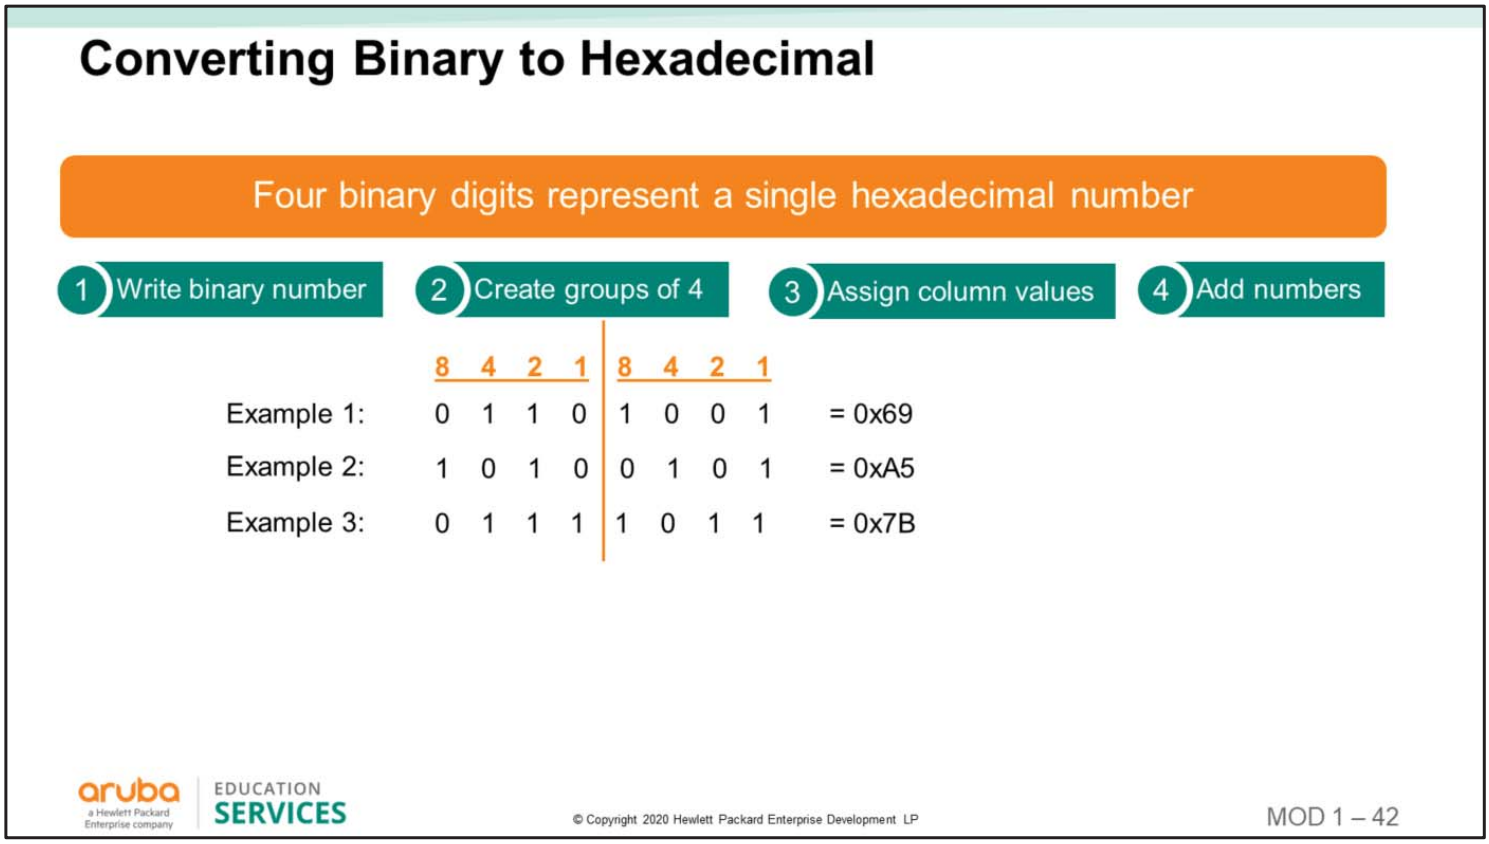

When you are going to convert a binary number to hexadecimal. Group 4 binary numbers (nibble), and assign column values of 8,4,2,1 to each group, as shown in the figure.

Now look at example 1. In the most significant nibble, there are binary 1's in the 4's column and the 2's column. 4+2 = 6. In the least significant nibble, there are binary 1's in the 8's column and the 1's column, and 8+1 = 9. Therefore, 0110 1001 = 0x69 in hexadecimal.

Now consider the second example. The high nibble has 1's in the 8's and 4's column. 8+4 = 12 in decimal, and 12 = 0xC in hex. The low nibble has 1's in the 4's and 1's column, and 4+1 = 5. Therefore, 1010 0101 in binary = 0xC5 in hex.

In the third example, the high nibble is 4+2+1 = 7. The low nibble is 8+2+1 = 11 in decimal, which equals 0xB in hex. Thus, 0111 1011 = 0x7B in hex.

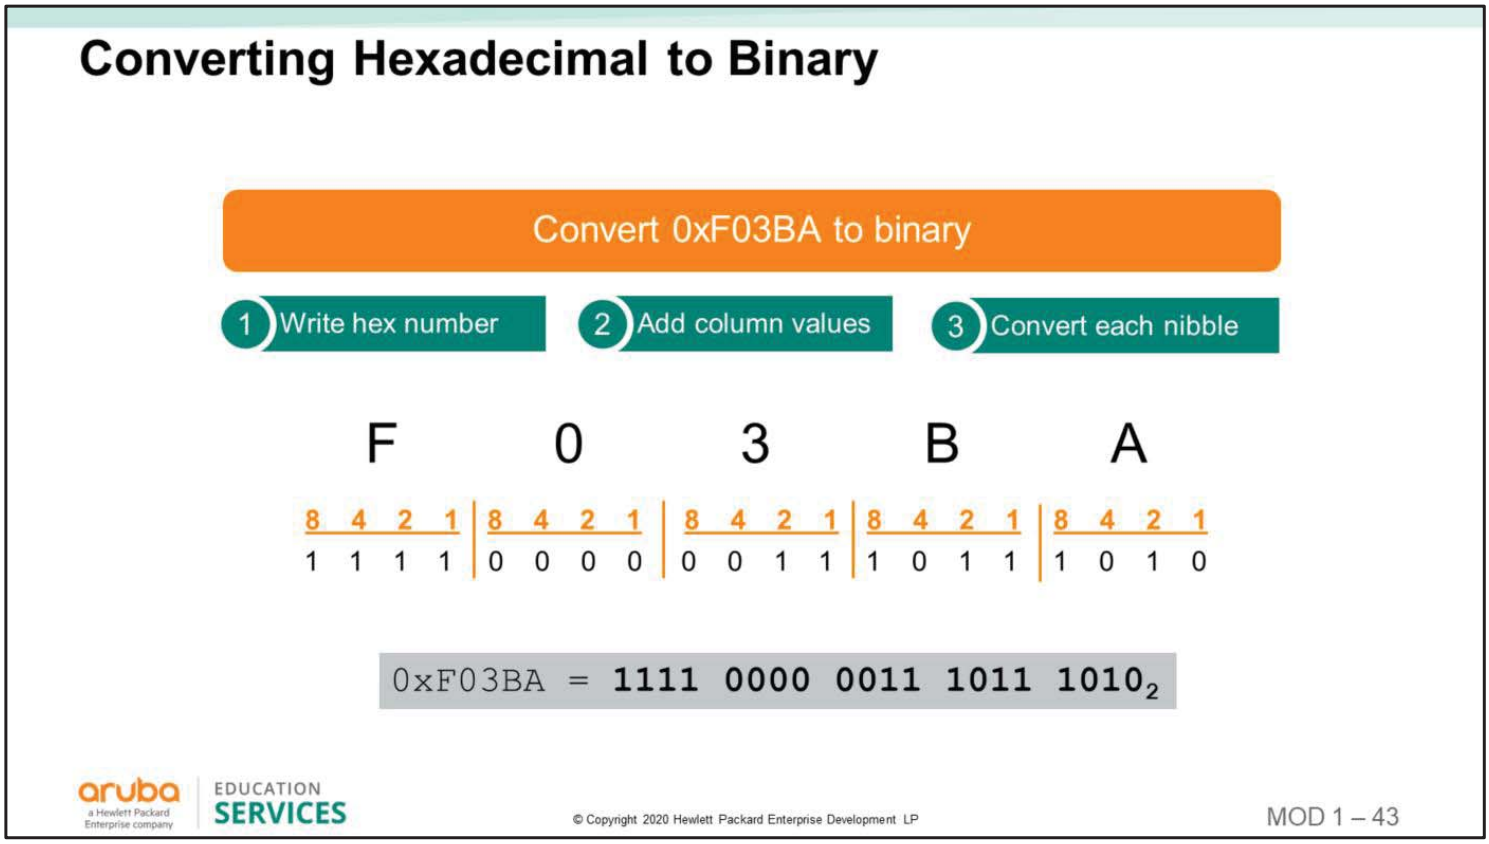

Now suppose you need to convert from hexadecimal to binary.
Simply write the hexadecimal number down, leaving ample space for 4 bits underneath each hex number.
Then add the 8,4,2,1 column values if you like.
Now simple convert each hex value to its binary equivalent. Reference the chart until you have the values memorized.
You know that 8+4+2=1 = 15 in decimal, which equals 0xF in hex.
Of course, 0 in hex = 0000 in binary, and so on.

Thus, 0xF03BA in hex = 1111 0000 0011 1011 1010 in binary.

Question #4

Binary to decimal conversion

What does 11010110 equal in decimal?

a. 212

b. 198

c. 214

d. 218

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

Answer: c

Lab Activity

Lab 1: Numerical Conversion

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Additional Content

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Decimal to Hexadecimal Conversion

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

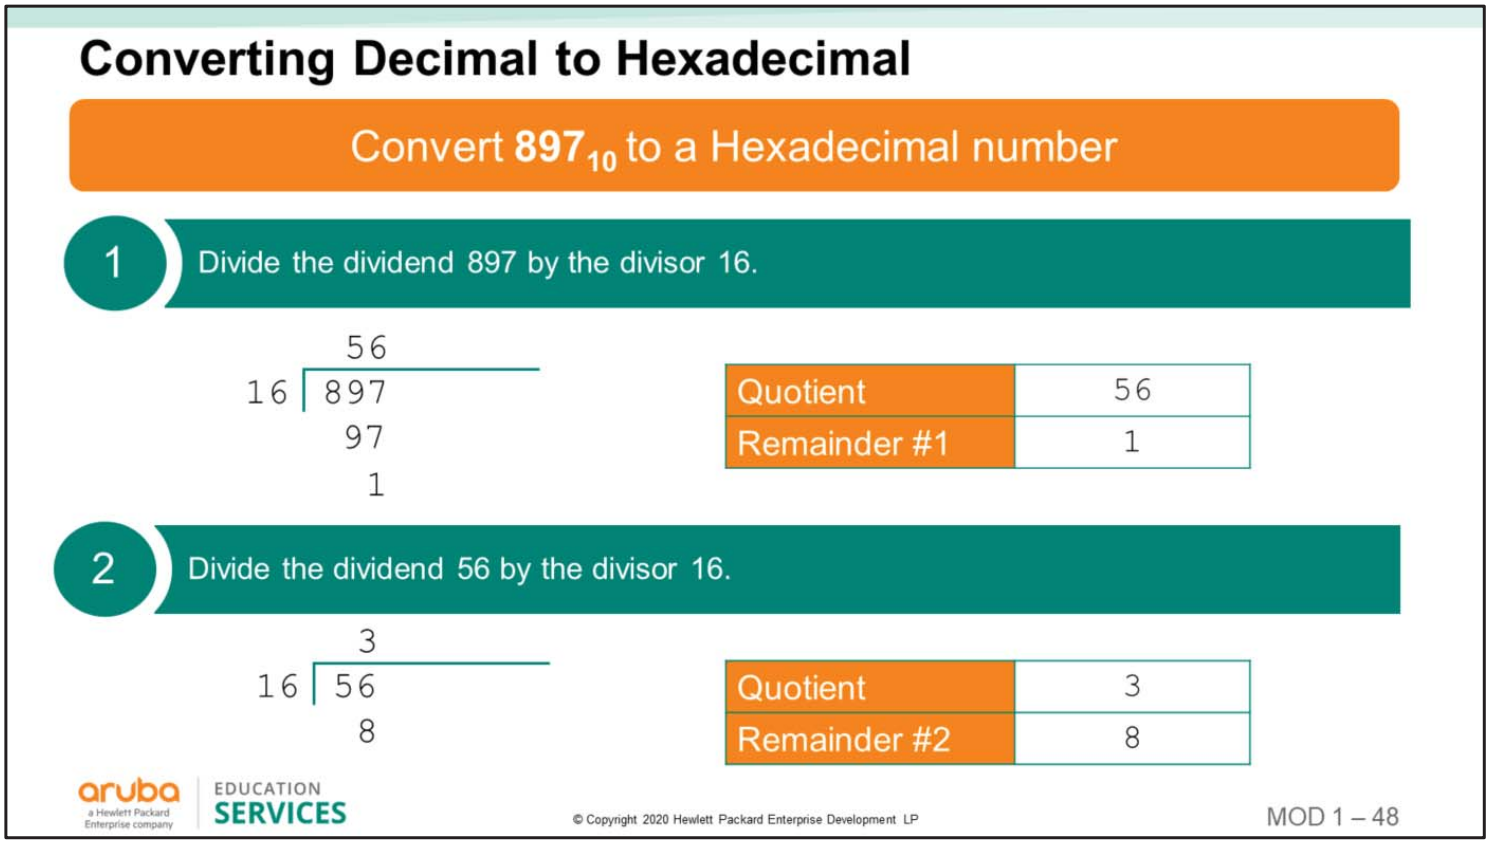

The process relies on a repeated division-by-16 process. Start dividing the decimal number by 16. Keep track of the quotient and the remainder. Continue dividing the quotient by 16 until you get a quotient of zero, then just write out the remainders in the reverse order.

Consider the following example, where the number 89710 is converted to its Hexadecimal representation.

Step 1. Divide the dividend 897 by the divisor 16.
Step 2. Divide the quotient obtained in the previous step (56) by 16.

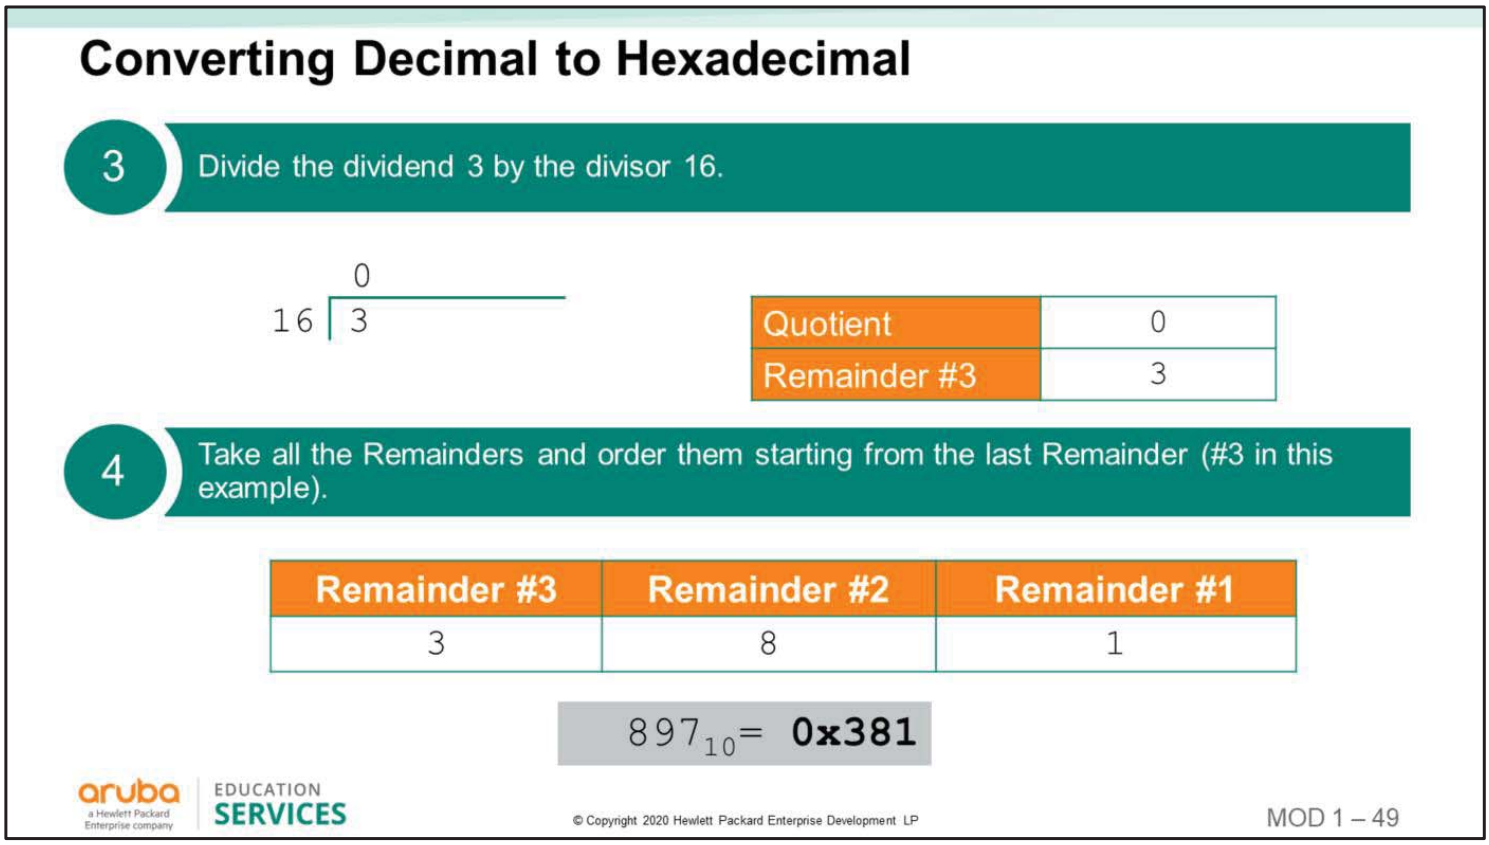

Step 3. Divide the quotient obtained in the previous step (3) by 16.
Step 4. Take all the Remainders and order them starting from the last Remainder (Remainder #3)

Note. On step 3 you know that the quotient of the division is actually a fraction number (0.1875). However, the process only focuses on the integer part of the quotient and the last remainder before the operation is actually done.

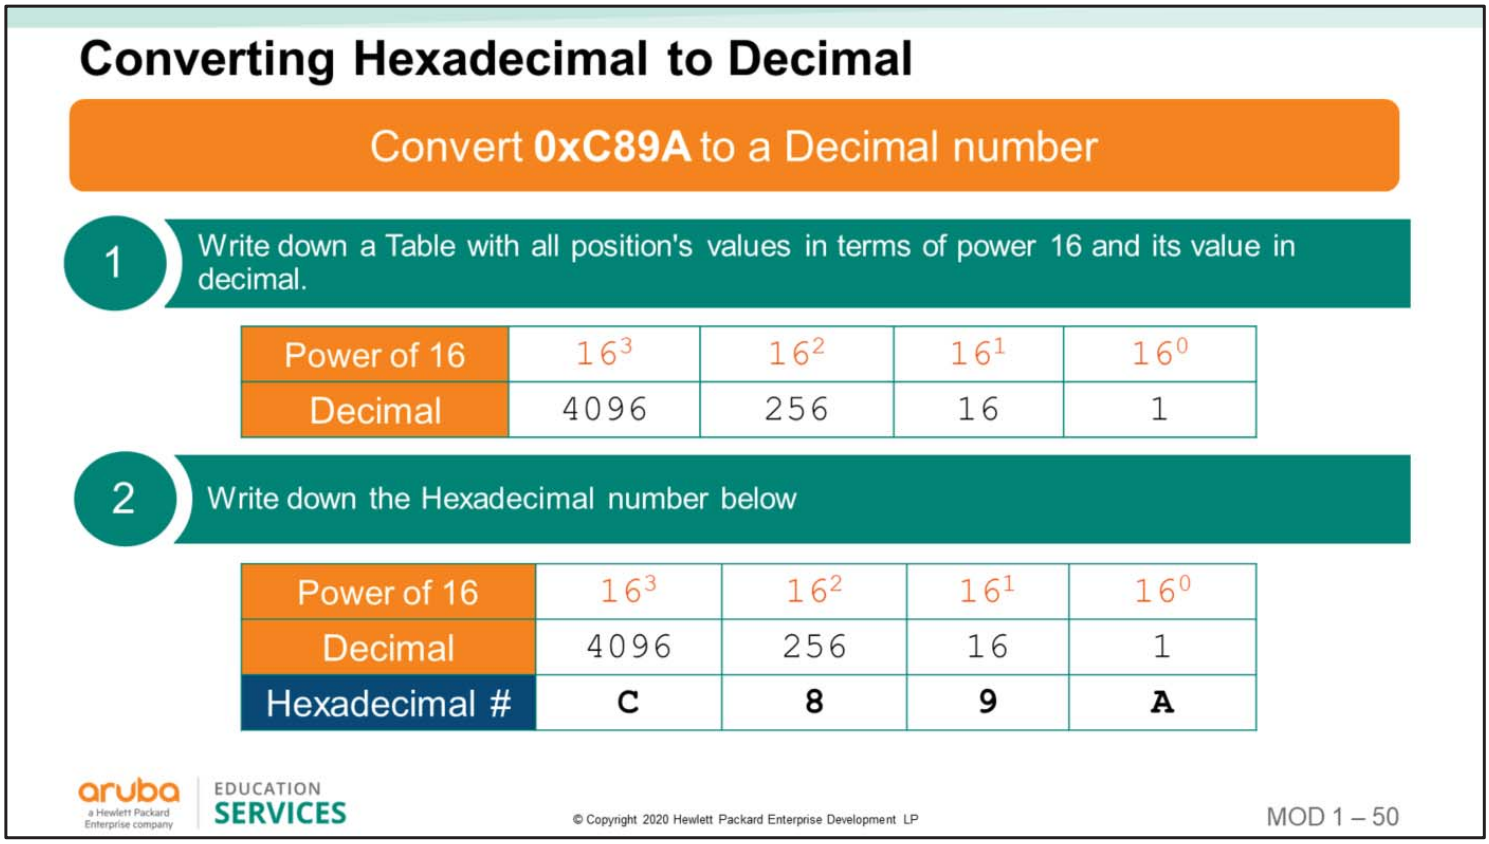

There are several methods to convert a hexadecimal number into decimal, however the comparison method is the easiest to learn. This method simply uses the position value of each number (remember that in hexadecimal system each position is a power of 16) and then sum the values.

As an example, let's convert the hexadecimal number 0xC89A to decimal.

Step 1. Create a table that includes the values of each position in power of 16 and the associate decimal value.
Step 2. Write down the hexadecimal number below that, as shown in the figure.

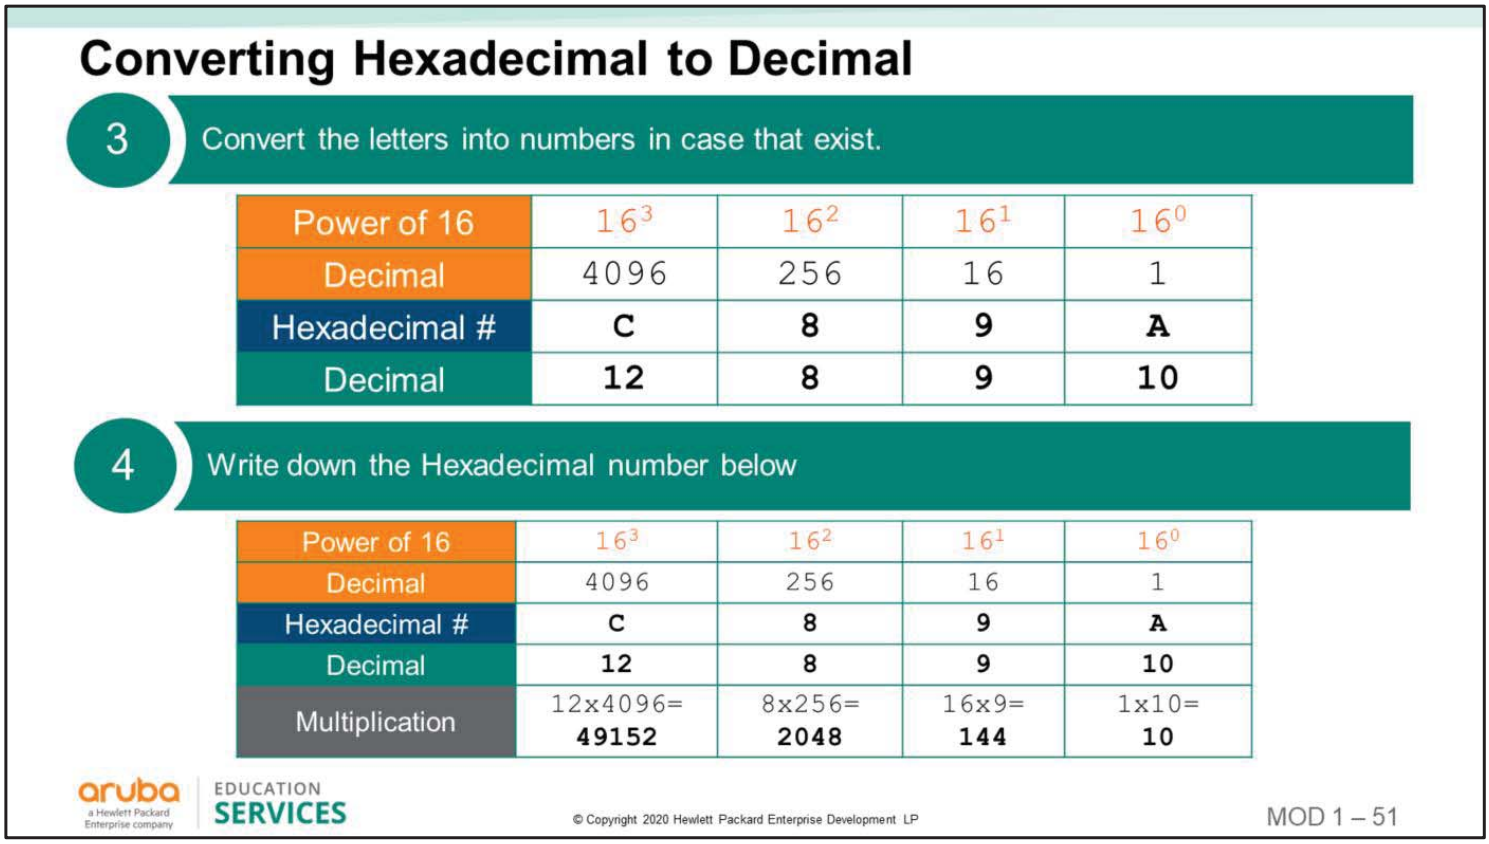

Step 3. Convert any existing letters into numbers. Please refer to the table presented previously. You will soon have this memorized, with only a bit of practice.
Step 4. Multiply the number obtained in the previous step times the value of the position.

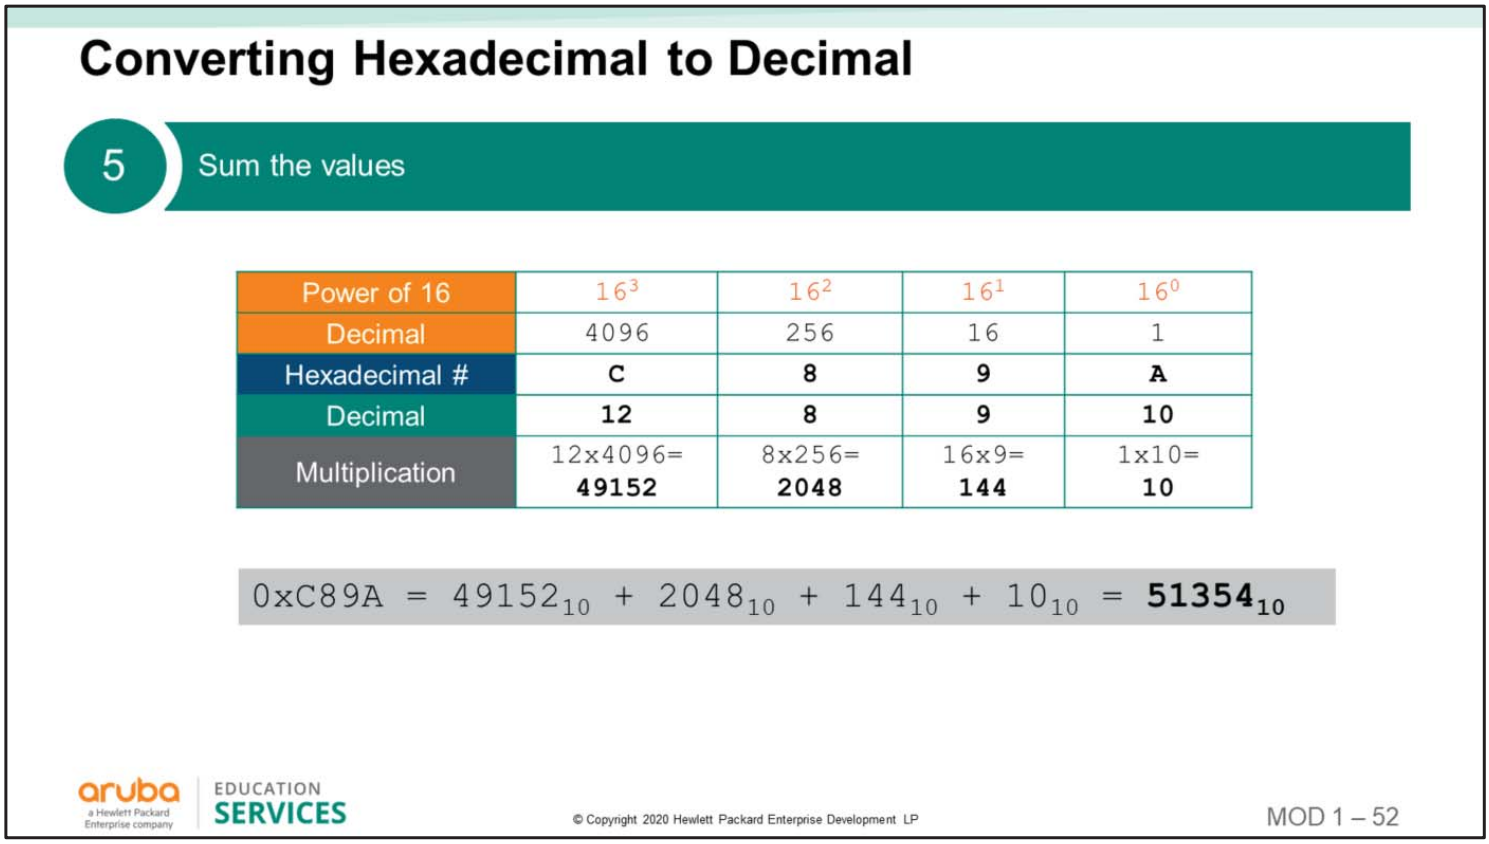

Step 5. Sum the values

0xC89A = 4915210 + 204810 + 14410 + 1010 = 5135410

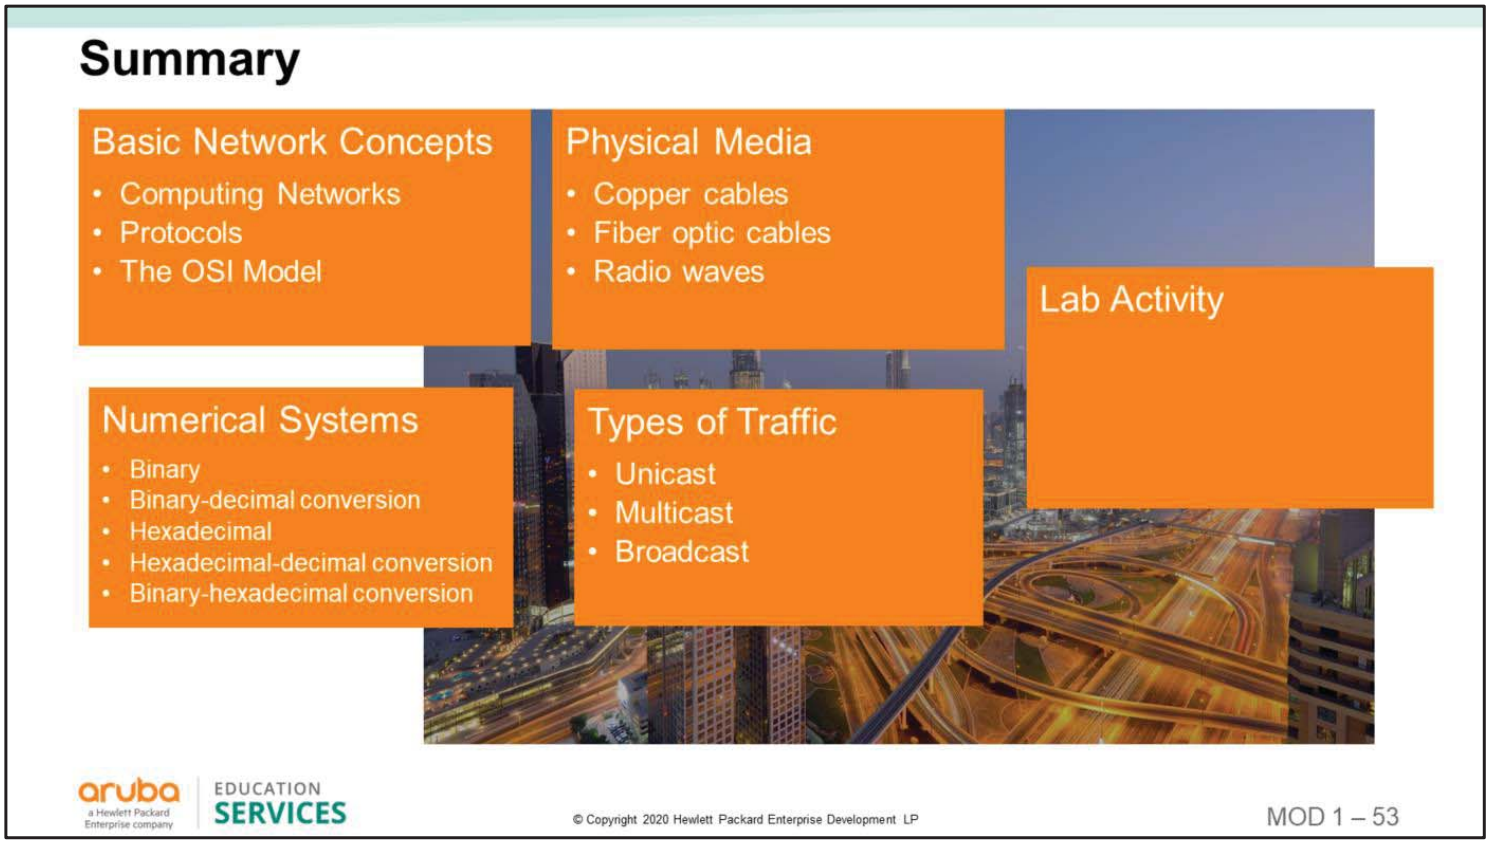

Congratulations! Now you have the essential knowledge and skills to grasp advanced concepts and build your career.

You learned about computer networks and related concepts – LAN, WAN, protocols, and the OSI model.

You learned which protocols map to each layer of the OSI model, and how each protocol encapsulates the data with header information to create segments, packets, and frames. You learned about physical transmission medium – copper wires, fiber optic cables, and radio waves.

You also explored the binary and hexadecimal numbering systems, and how to convert between them.

You explored the TCP/IP protocol stack, with a special focus on Ethernet and Wi-Fi frames, before learning about unicast, multicast, and broadcast traffic types.

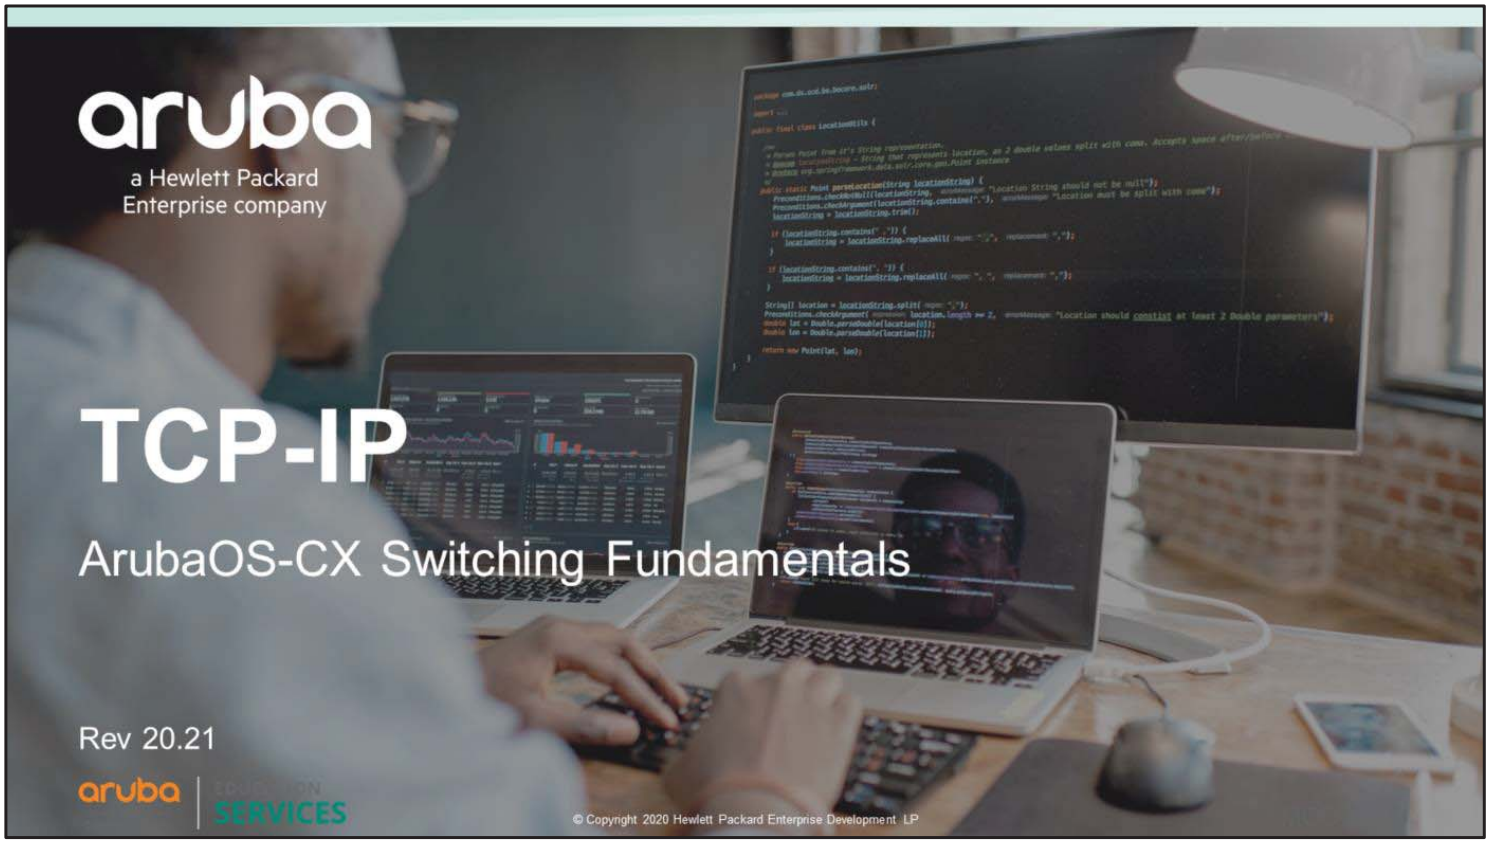
aruba
a Hewlett Packard Enterprise company
TCP-IP
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP

Objectives

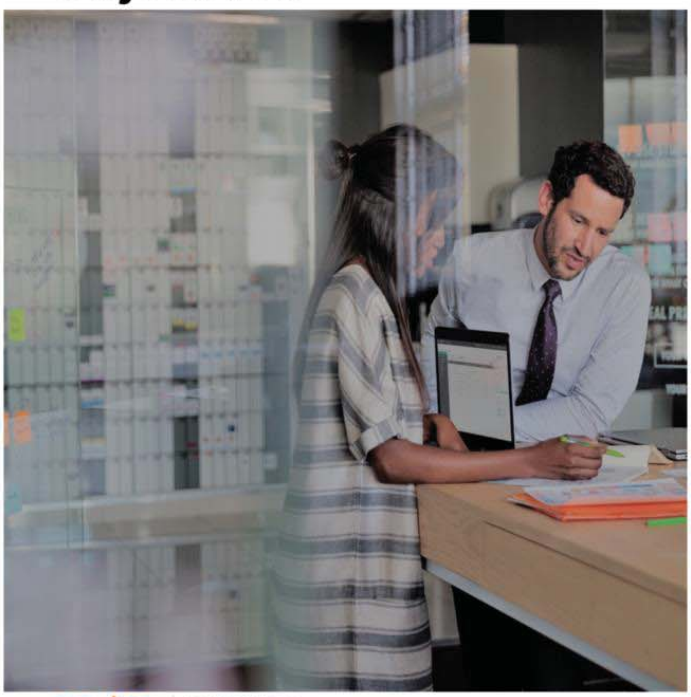

–Learning Objectives

- –Describe network devices and common network services
- –Explain network hierarchical models

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 2

After completing this module, you will be able to describe the typical devices that are used to create a network – switches, routers, multi-layer switches, access points, firewalls, and servers. You will also be able to explain common networking services that are used over these networks, including DHCP, DNS, HTTP, Telnet and SSH, and FTP.

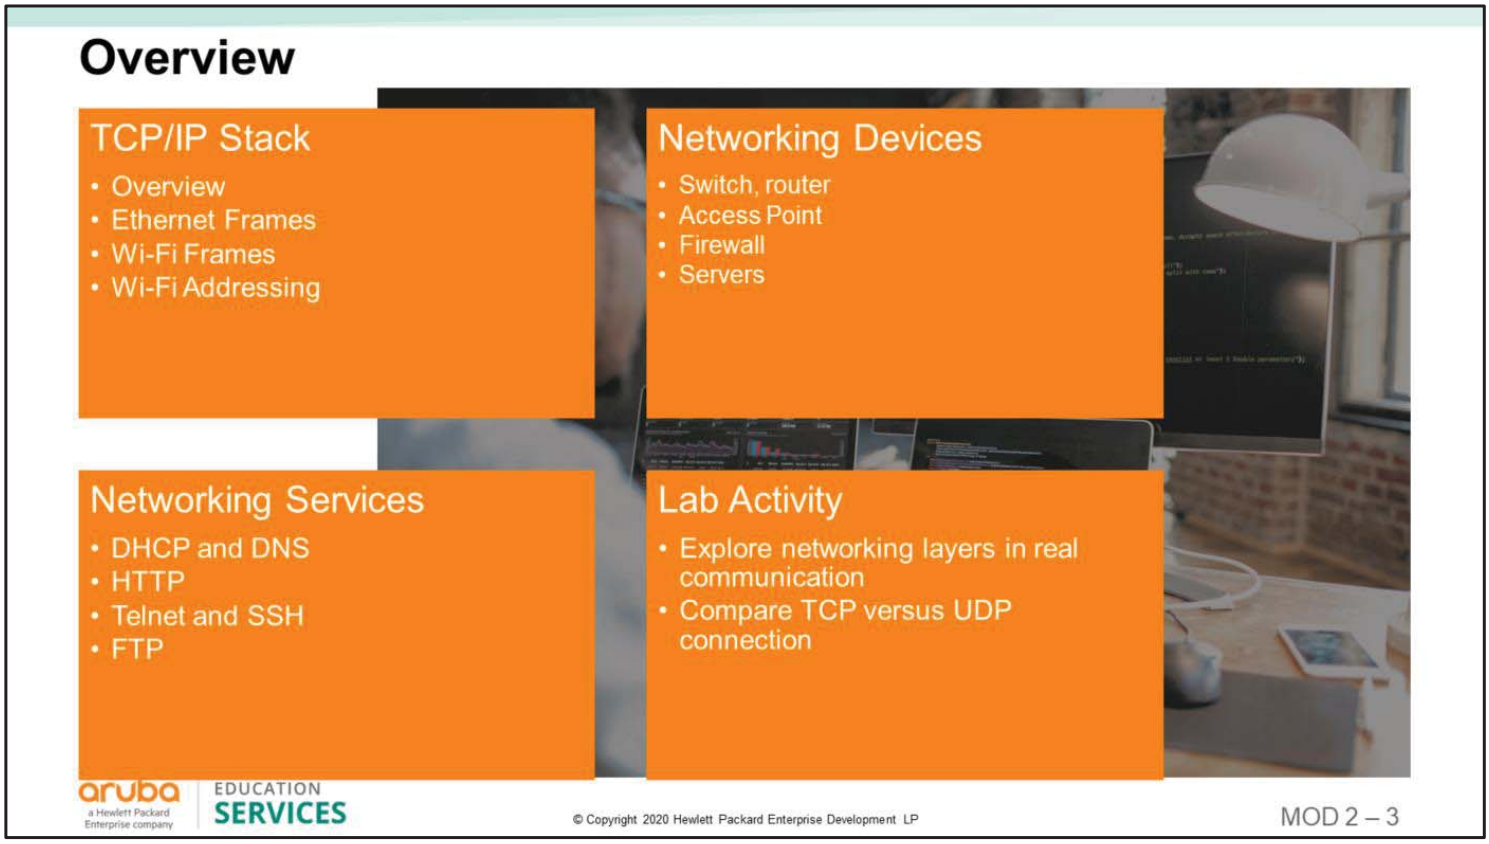

You begin this module by exploring the networking devices used to build a typical network. Understanding the basic of these components is the first step in being able to design, implement, and troubleshoot networks. You will also learn about important network services that are vital for proper network function and end-user interactions. Finally in the lab activity you will be able to explore the Networking layers in a real communication and compare a TCP connection versus UDP connection.

TCP/IP Stack

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

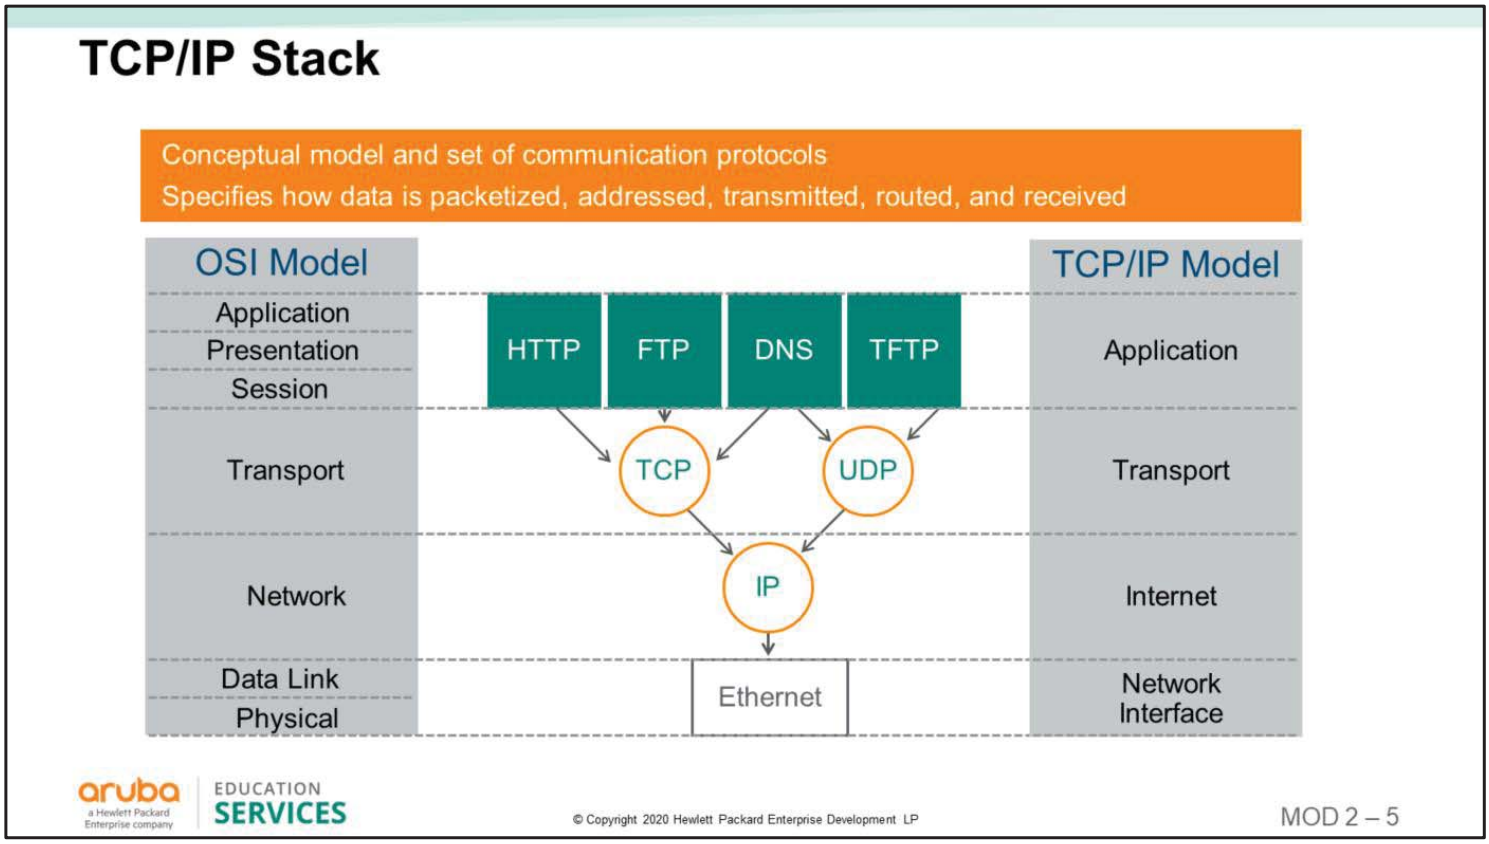

The Internet Protocol suite is a conceptual model and set of communication protocols used in Internet and computing networks. It is commonly known as the TCP/IP stack because the foundational protocols in the suite are Transmission Control Protocol (TCP) and the Internet Protocol (IP). This model provides end-to-end data communication. It specifies how data should be packetized, addressed, transmitted, routed, and received. This functionality is organized into four abstraction layers.

|The TCP/IP model is often compared to the OSI model, as shown in the figure. You see that the OSI model has seven layers, which map to the four layers of the TCP/IP model. As you can see the functionality of the OSI model's Application, Presentation, and Session layers all map to a single TCP/IP Application layer.

|This means that TCP/IP-based applications are responsible for interfacing with users and creating the data (application layer), putting it in the proper format (Presentation Layer), and managing sessions (Session Layer).

The transport layers of each model directly map. Recall that this layer is responsible for flow and error control. This layer establishes an end-to-end connection of two devices whose logical connection traverses a series of networks.
|The two choices used in a TCP/IP network are UDP and TCP.

The OSI Network Layer maps to the TCP/IP Internet layer. Both perform the same function - routing data across logical network paths by defining a packet format and an addressing format.
|IP version 4 (IPv4) and IP version 6 (IPv6) are the protocols that works on this layer.

The OSI models Data Link and Physical functions are all defined in a single layer in the TCP/IP model – the Network interface layer. As you learned in the section about the OSI model, these layers contain protocols relating to the physical communication medium (copper wires, fiber optics, and radio waves). |Recall that physical media access control also occurs here. Ethernet and Wi-Fi are very common protocols that work at this layer.

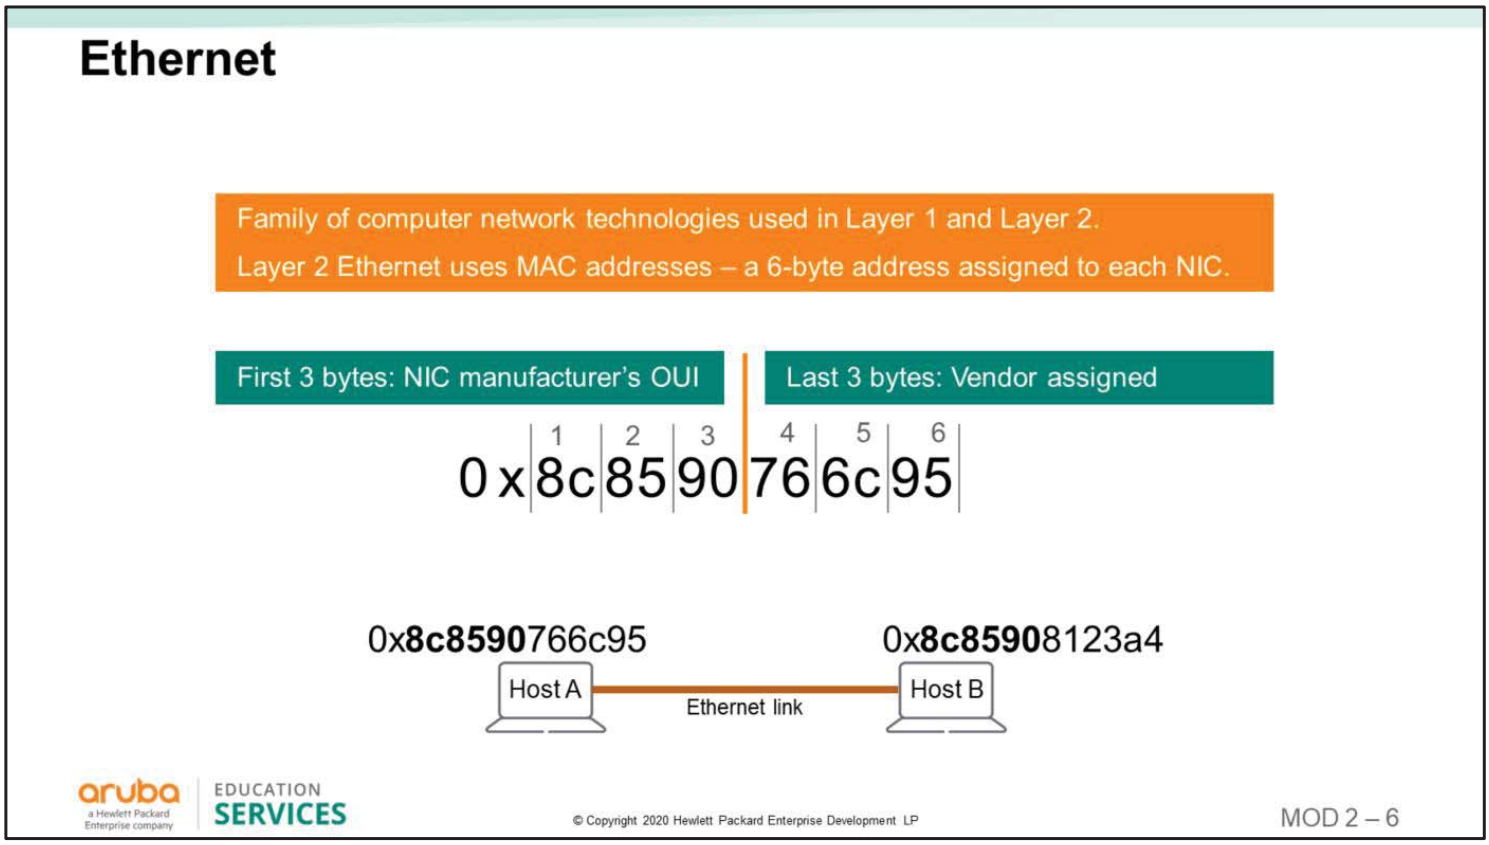

Ethernet is a family of computer networking technologies that are used at Layer 1 and Layer 2 of the OSI model. At Layer 2, this standard provides for the assignment of 6-byte (48-bit) Media Access Control (MAC) addresses to each device. Each manufacturer assigns a globally unique hexadecimal MAC address to each Network Interface Controller (NIC). This address is stored in permanent Read Only Memory (ROM) on the NIC. In other words, every Ethernet NIC in the world has a unique MAC address. How is this possible?

|The most significant first three bytes of each MAC address are referred to as the manufacturer's Organizationally Unique Identifier (OUI). A standards body assigns each manufacturer in the world a unique 3-byte number. All NICs from that vendor therefore have the same number in the first 3 bytes.

|The last three bytes of the MAC address comprise the NIC serial number. Recall that with hexadecimal a single byte is comprised of two digits. Thus, each MAC address is represented with 12 hexadecimal characters, each Byte or octet is typically separated by either a colon or dash.

For example for the MAC address 8c:85:90:76:6c:95, the OUI is 8c:85 :90, while the NIC serial number is 76:6c:95. As you progress through this course, you will often see fictitious, 2-byte MAC addresses such as 8c::95. This just makes it easier for you to focus on fundamental concepts, as opposed to always trying to read out large, 6-byte hexadecimal numbers.

To verify the vendor of a NIC by its MAC address, access a web site like https://macvendors.com , which allow you to perform OUI lookups. If you were to lookup the address in this example, you would determine that the manufacture is Apple.

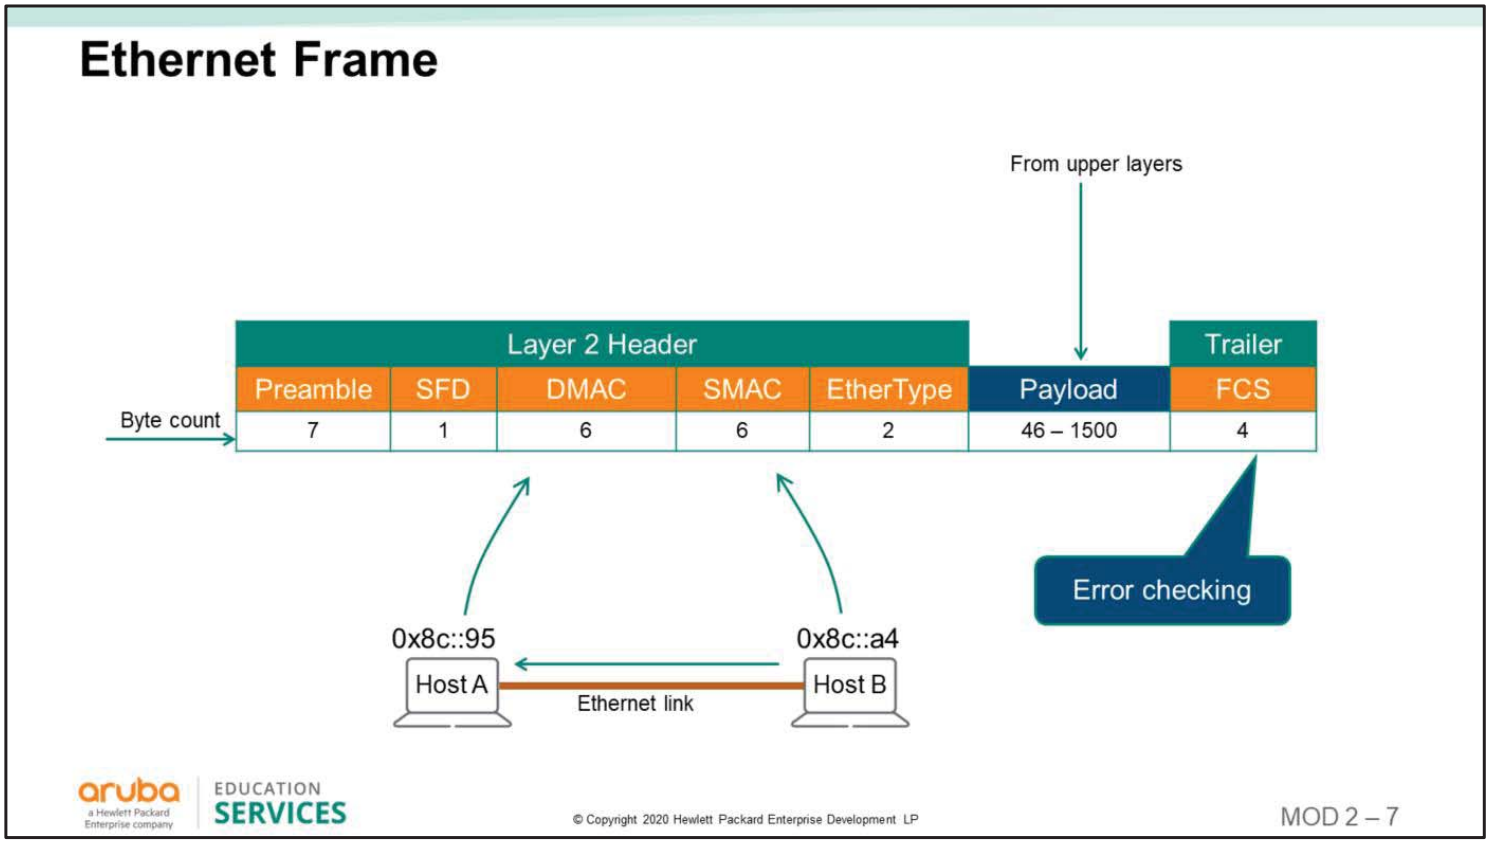

Suppose that Host B receives data from the upper layers of the OSI model. Think of Ethernet as an employee of the upper layers. They have given Ethernet this payload, and asked Ethernet to get this over to host A. To do this, the Ethernet process must wrap this payload with a header and trailer, as shown in the figure. Each field in the header and trailer is described below:

Preamble. Informs the receiving system that a frame is starting and enables synchronization.
SFD (Start Frame Delimiter). Indicates that the Destination MAC address field begins with the next byte. There might be several hosts on this Ethernet segment. The preamble and SFD are seen by every host as a kind of signal, “Hey, here comes a frame, and it might be for you!” How does the receiving station know whether the frame is for it, or some other host? By reading the next field – the destination MAC address.
Destination MAC(DMAC). Identifies the NIC of the receiving computing system. In this example Host A reads the MAC address and knows, “Hey, this is my MAC address – I should accept this frame”. All other stations read the MAC address, see that the message is not for them, and drop the frame. This is like someone yelling in a room, “Hey Marta, can I buy you an ice cream?” Only Marta will respond.
Source MAC(SMAC). Identifies the NIC of the sending computing system. This is how Host A knows where the message came from, and therefore to whom it should respond.
EtherType. Defines the Layer 3 protocol inside the frame, for example IPv4 or IPv6. Host A may be running an IPv4 protocol stack and an IPv6 protocol stack. The type field tells host a which stack to pass the payload to for processing.
Payload. Contains the data that to be transmitted. FCS (Frame Check Sequence). Allows detection of corrupted data using the Cyclic Redundancy check (CRC) or checksum method. As Host B creates this Frame, it performs a mathematical calculation on most of the

header and payload bits. It places the result of this calculation in the FCS field. When Host A receives the frame, it performs the same calculation. If it comes up with the same result that is stored in the FCS field, the frame is not corrupted. But if you have bad wiring, or strong electromagnetic interference (EMI) the data may become corrupted. In this case, Host A's calculation will not match the FCS, and so the frame is discarded – no need to accept bad data.

It is important to understand that the Ethernet payload not only carries application data – it also includes the header information from the upper layers. Remember, TCP adds a header, which IP considers part of the payload. Then IP adds a header, which Ethernet considers part of the payload. Let us look now at this IP header information.

Note: As promised, the figure only shows the first and last bytes of each 6-byte MAC address, to ease the learning experience.

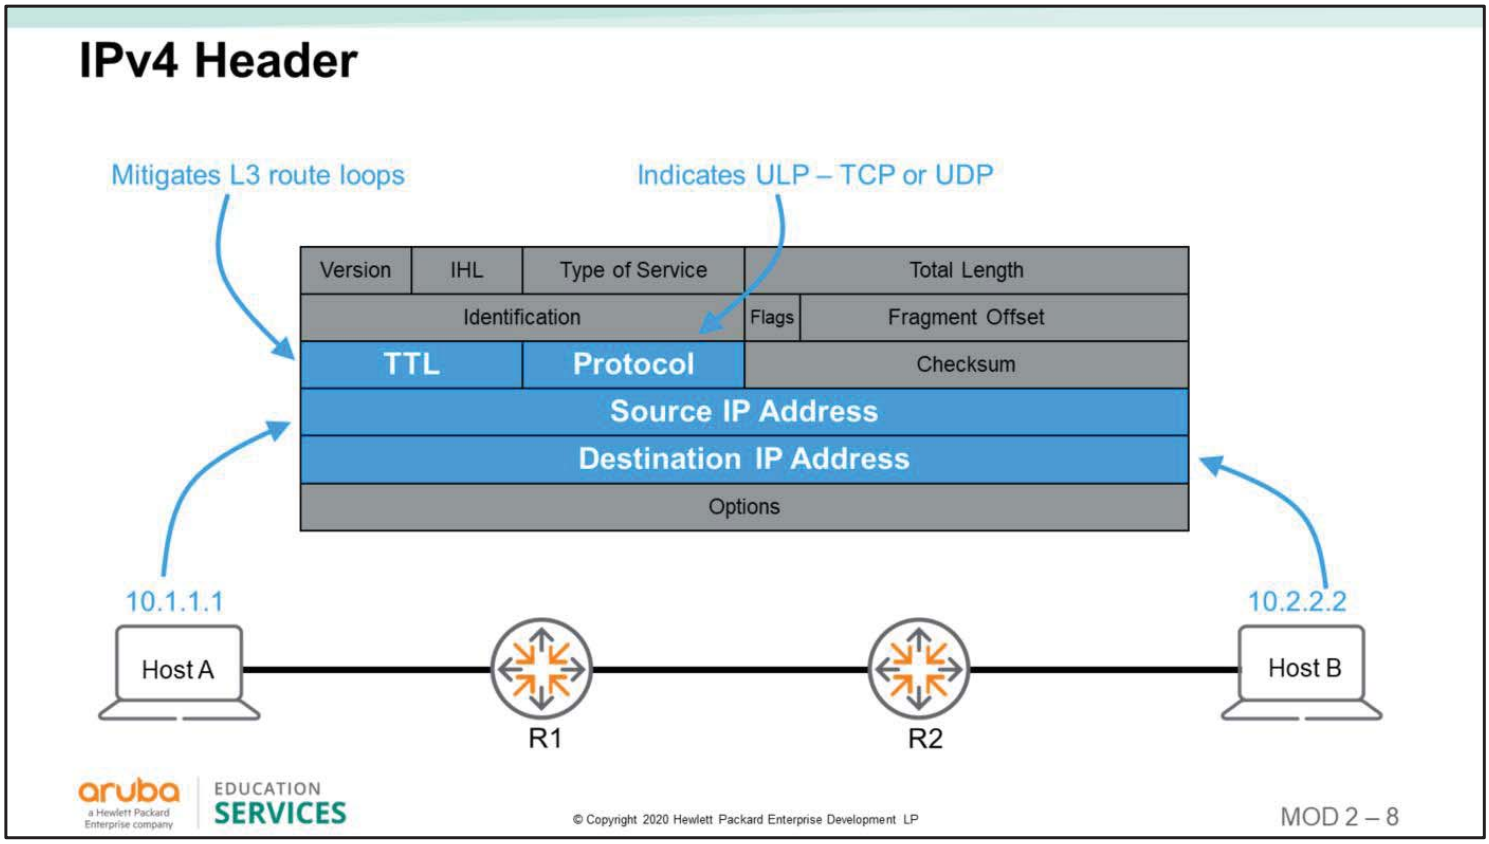

The figure shows the IPv4 header information added to a packet at Layer 3, with a special emphasis on the most important fields in this header, as described below:

|Time To Live (TTL): This field mitigates Layer 3 route loops. If routers are misconfigured, IP packets could be bounced around between two or more routers forever. To prevent this, the original places a number in this field from 0 to 255. For example, suppose that TTL – 15. Each router that routes this packet decrements this field by 1. So, after this packet “hops” through the first router, TTL = 14. After the next router, TTL = 13, and so on. The router that decrements the TTL field to zero discards the packet.
|Protocol: This field indicates whether this IP packet carries TCP traffic or UDP traffic. Recall that application data is passed down to TCP or UDP, which adds a header and passes this on down to IP. From IP perspective, TCP, UDP, and other protocols are at higher layers in the model, and so can be collectively referred to as Upper Layer Protocols (ULP).
|Source IP Address: This originator of this IP packet adds its IP address to the source address field. In the figure, host 10.1.1.1 sends a packet to host 10.2.2.2. Thus, the Source IP Address fields = 10.1.1.1, and the Destination IP address field = 10.2.2.2.
Destination IP Address: The ultimate destination of the packet, as described above.

|The other fields are shown here in grey, because they are not relevant to your learning in this and several of the modules that follow. You will learn about these other fields as they become relevant.

Remember, this is the header that Layer 3 routers use to do their job. They receive an inbound packet, analyze the destination IP address field, and forward packets along the best path toward that destination. In the figure, router R1 receives this packet from host 10.1.1.1 and forwards it to R2. R2 receives it and forward it to destination host 10.2.2.2

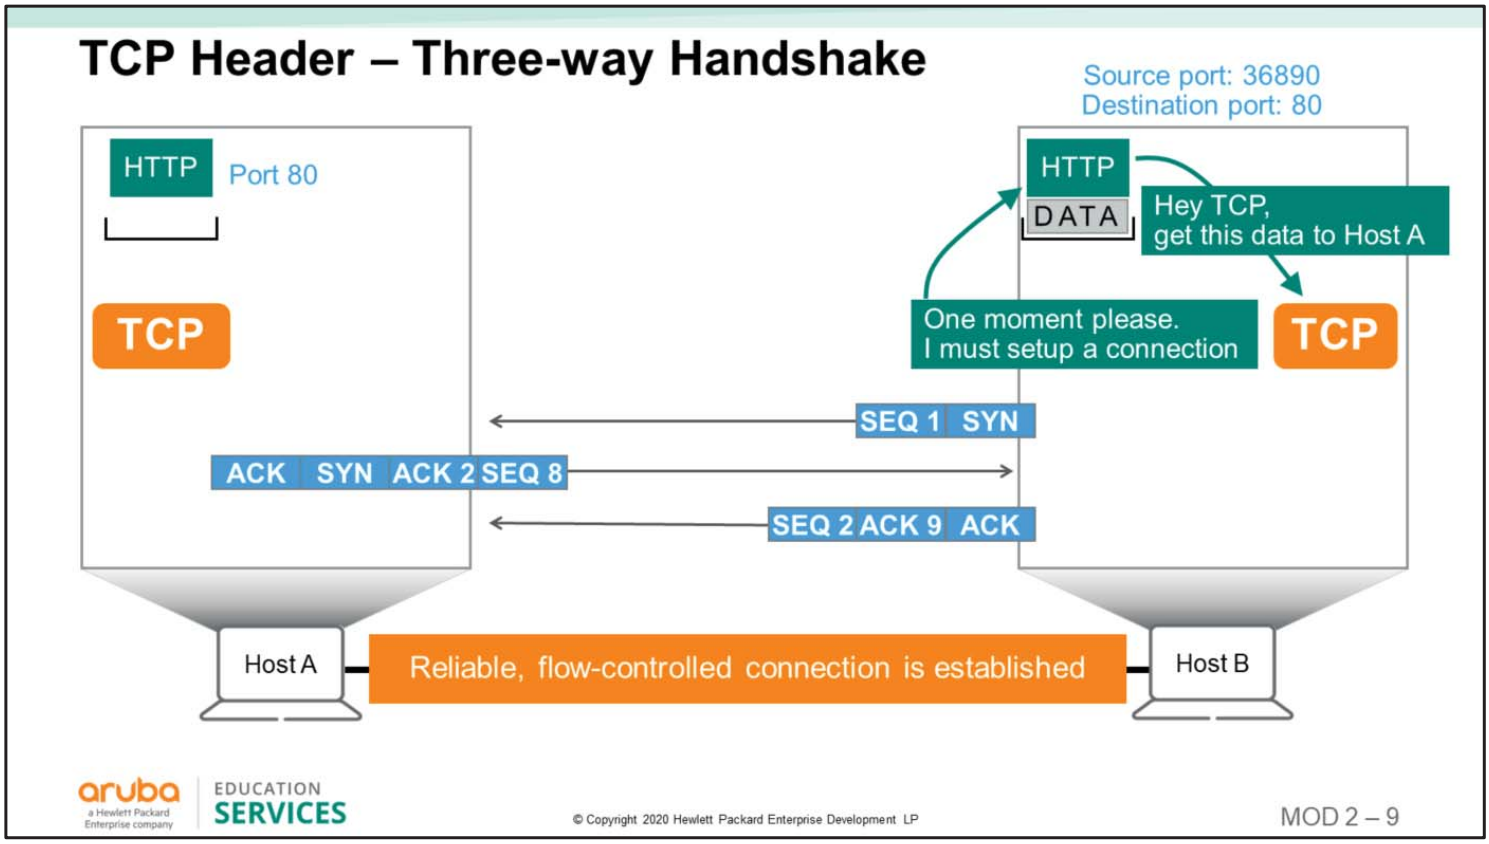

Recall that TCP acts as a kind of “employee” for applications, which rely on TCP to get data to some destination host. Before TCP processes any application data, it must first contact the target device and establish a reliable, flow-controlled connection.

In the figure, the Host B’s HTTP application has data – sitting in a kind of “out box”, waiting to be transmitted to Host A.
| It is as if HTTP says, “Hey TCP, I need you to get this data to Host A”.

|TCP says, “OK boss, but first I must establish a connection with Host A – to ensure that your data is sent in a controlled, reliable manner”.

There are four primary fields in the TCP header that are used to establish this connection – The sequence (SEQ) number, the acknowledgement (ACK) number, the Synchronize (SYN) flag, and the Acknowledgement (ACK) flag.

|Host B starts by creating a TCP packet. This packet carries no data – it only signals Host A, “We should establish a connection”. In this example, Host B sets SEQ = 1, and raises its SYN flag (sets it to a binary 1).

Host A receives this and knows, “This packet came from Host B with the SYN flag raised – Host B wants to connect”. Host A must inform Host B that it received SEQ = 1, and so expects SEQ = 2 next.
| To do this, Host A sends a packet with ACK = 2, and both the SYN and ACK flags are raised. Understand that Host A is independent of Host B. Host B tagged its packet with SEQ = 1. Host A chooses to tag its packet with SEQ = 8. Host A is saying, “I’m using my packet number 8 to acknowledge receipt of your packet number 1”.

|Host B receives and responds to this packet, “You expected SEQ=2 next, so here it is...I am using my packet number 2 to acknowledge receipt of your packet number 8. I expect packet number 9 next. As you can see, both the SYN and ACK flags are raised in this packet.

|At this point, a reliable, flow-controlled connection is established. Each host has successfully received and acknowledged a sequence of packets from the other. Now TCP is ready to start transferring that data for the application.

HTTP uses a specific port to transmit the information in this case the destination port number is 80 the source port in this case is randomly selected on Host B, in this case Host B select the port 36890 as the source port.

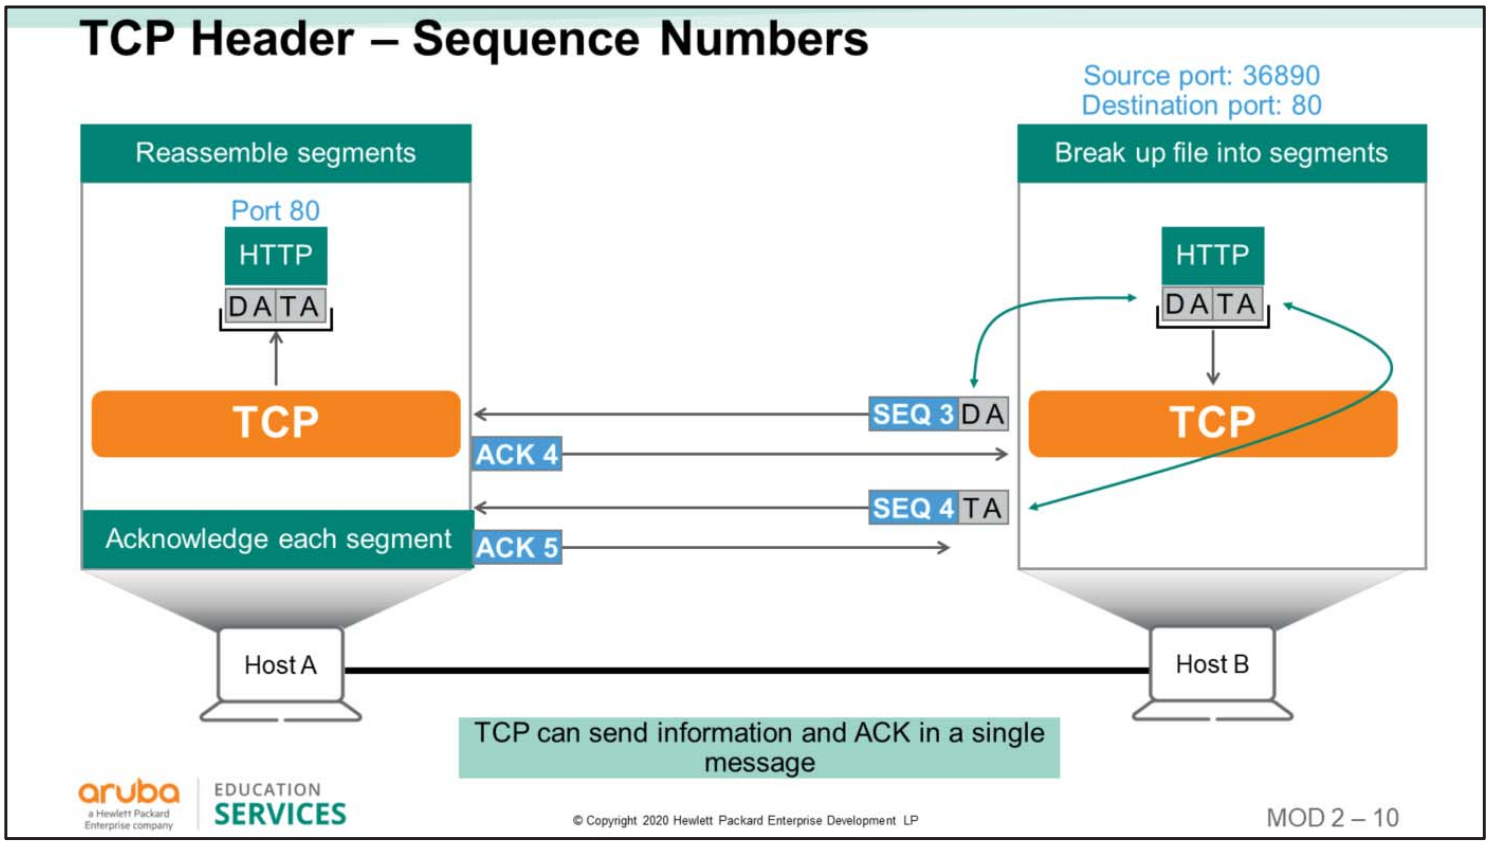

TCP may be asked to transfer very large files – too big to transfer all at once. Therefore, Host B's TCP process breaks up the large file into smaller pieces called segments and assigns a unique sequence (SEQ)number to each segment. In this example you want to transmit the word "DATA" Host B then creates 2 pieces one for the letters "DA" and other for the letters "TA"

| In the figure, TCP assigns Sequence = 3 to the first piece (DA) and transmits it to Host B. Host A acknowledges this by sending a packet with Acknowledgement (ACK) number = 4. Essentially, Host A is saying, "I received segment 3, and so I expect segment 4 next".

| Host B then takes the next piece (TA), adds Sequence = 4, and sends it to Host A. Host A acknowledges this by sending ACK = 5.

| As Host A receives this data, it reassembles the pieces back into a complete file and passes it up to the application.

Note. TCP can send a block of information and Acknowledge receipt of data blocks using a single ACK message.

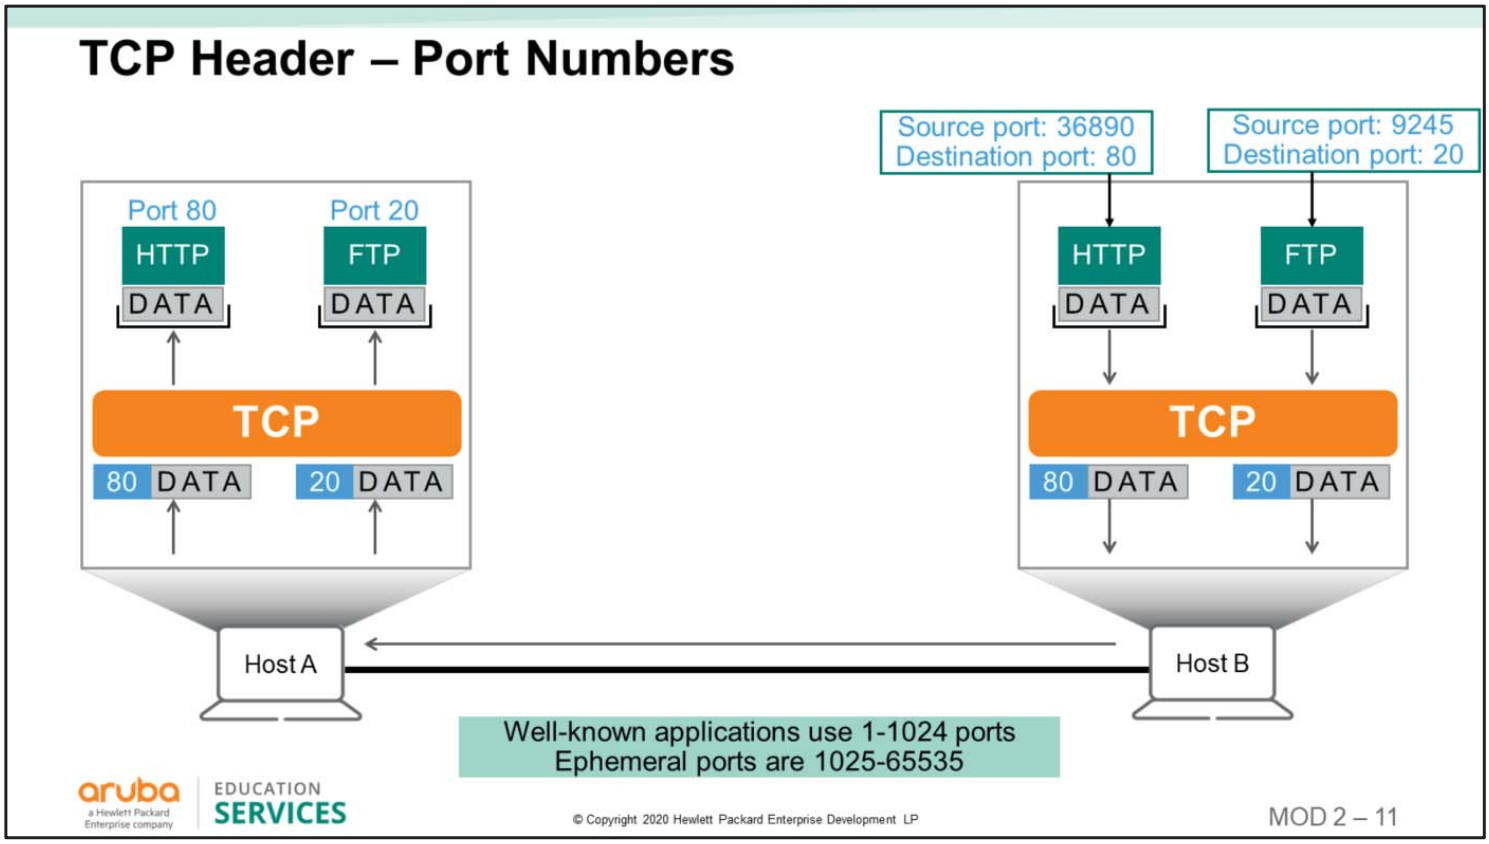

Recall that TCP acts as a kind of “employee” for applications. TCP receives requests from HTTP, FTP, and many other applications. It is as if the application says, “Hey TCP – establish a connection with the target host, and make sure my data gets there reliably”. TCP uses the destination port number to keep track of each application it serves.

To accommodate this, each application is assigned a standard port number. FTP uses port 20 for data transfer, HTTP uses port 80, and so on. When Host B’s TCP process accepts data from HTTP, it sets Destination Port = 80.

| Thus, when Host A receives this data, it knows to pass this data up to its HTTP application, as opposed to DNS or FTP.

In TCP you can use up to 65535 ports however the well-known applications use the first 1024 ports. From 1025 to 65525 ports are known as ephemeral Port numbers.

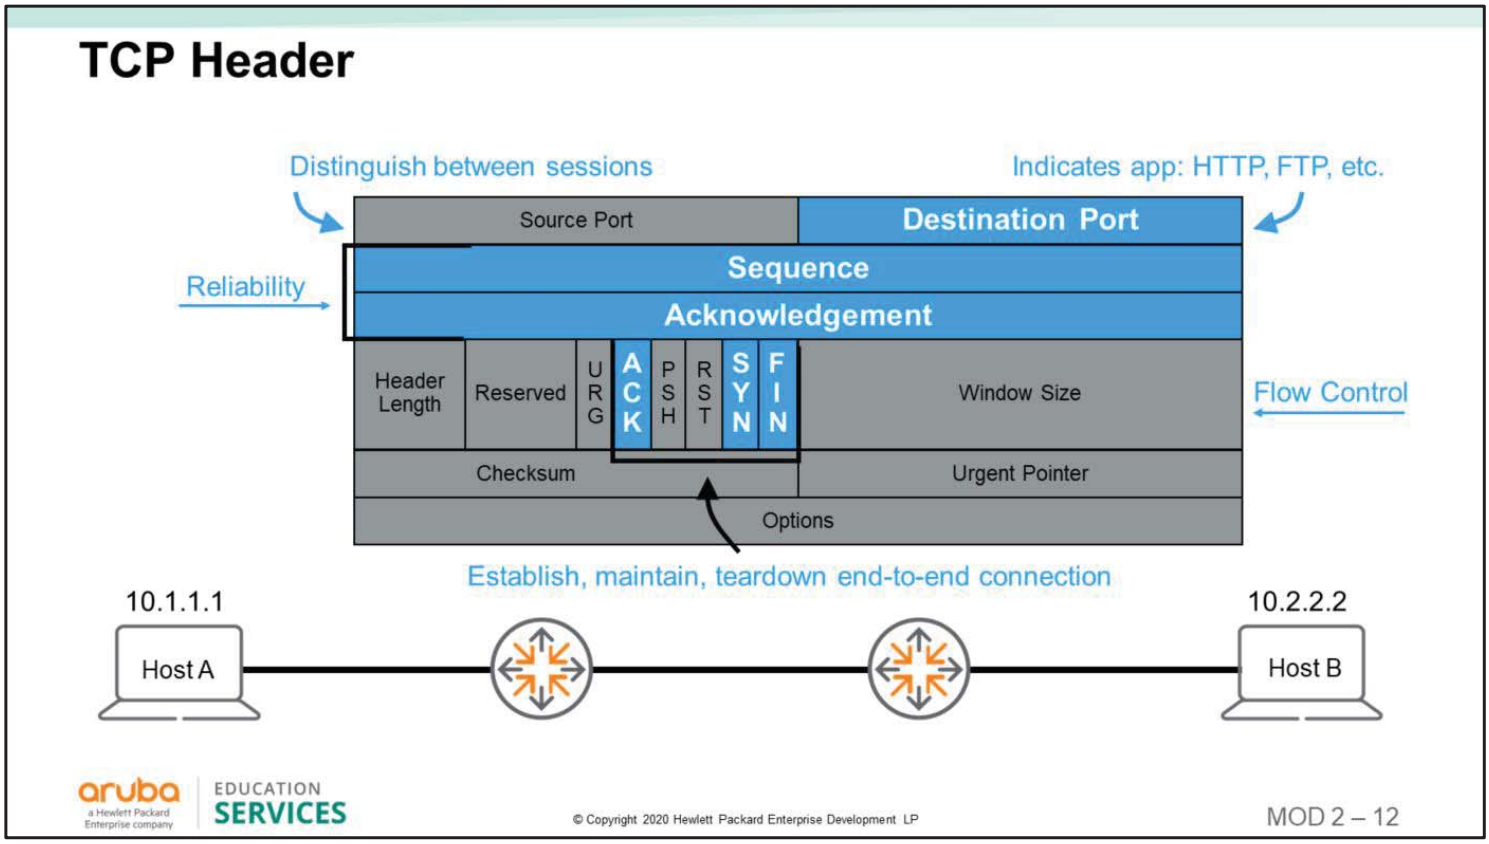

The figure shows the entire TCP header information added to a Layer 4 segment, with a special emphasis on the most important fields in this header, as described below:

Destination Port
|The first fields added are the source and destination port numbers. You just learned hosts use the destination port to distinguish the various applications that it serves. But what if you form two or more HTTP connections? You point your browser to your favorite search engine. Then open another tab in the browser and point it to www.arubanetworks.com. Bot sessions have Destination Port = 80. How does TCP know which data is for the search engine, and which data is for arubanetworks.com?

|TCP sets the source port to a random number. The session to the search engine has Source Port = 58936, Destination Port = 80. The session to Aruba has Source Port = 57576, Destination Port = 80. Thus, TCP can distinguish multiple simultaneous sessions for the same application.

Sequence and Acknowledgement Numbers
|You also learned about the sequence and acknowledgement numbers. They are used along with the ACK and SYN flags to ensure reliability and flow control. Host B sends SEQ=5, Host A responds with ACK=6, and so on. But what if Host A responded with ACK = 5? Host B says, “Hmm. I guess host A did not receive segment 5, so I will resend it”. In this way, TCP ensures reliability, resending any packets that are lost in transmission.

Flags
|You learned how the SYN and ACK flags are used to formally synchronize TCP packet (sequence) numbers and acknowledge receipt of packets. There is also a Final (FIN) flag,

which tells the receiver, “Hey Host A, this is the last piece of the file.”

|There is the urgent (URG) and Push (PSH) flags to indicate that certain data is urgent and should be quickly pushed up to the receiving application.
|The URG flag is used in conjunction with the Urgent Pointer field, which indicates where the urgent data is located. There is also a Reset (RST) flag, used to reset the connection.

Window size
|There’s a checksum for reliability, and a Window size field, used for flow control. If Host be is sending too much data at a time, Host A may not be able to handle the load. Host A may lower the window size, telling Host B, “Slow down, please do not send so much data at a time”.

The remaining fields are not either rarely used or not relevant to this initial discussion. You will learn about these fields as you continue your training, as appropriate.

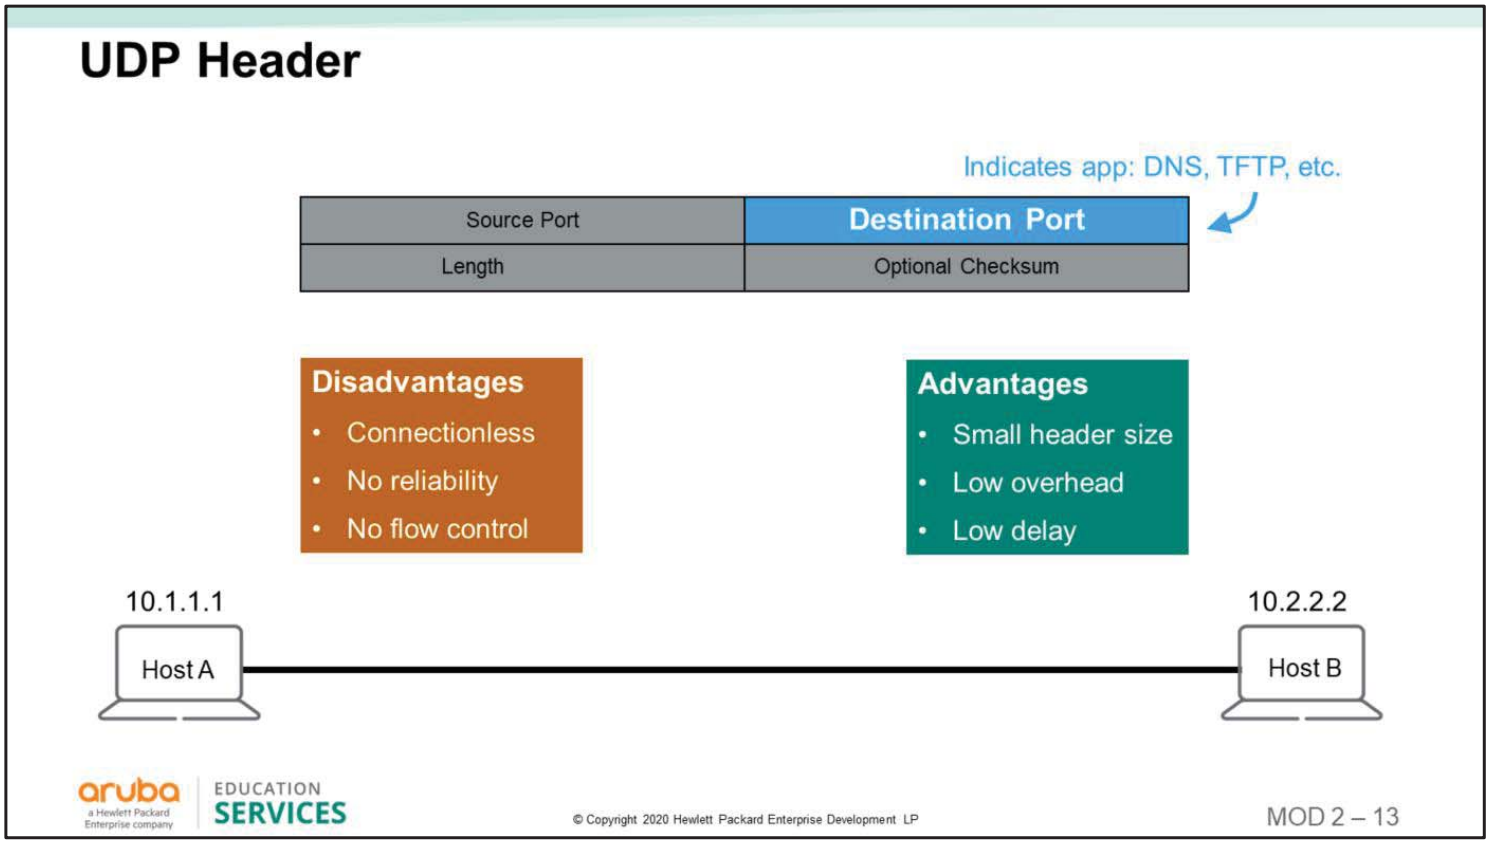

The figure shows the entire UDP header information added to a Layer 4 segment. Look how simple this is. There is only source and destination ports, used in exactly the same way as for TCP. Then there is a Length field, and an optional checksum. For applications, TCP is a very dedicated employee – it established a connection with the other side, retransmits any lost packets, and makes sure the receiver does not become overwhelmed with data.

In contrast, UDP is like a lazy employee – no connection, no reliability, and no flow control. Did the packets actually get to the destination? UDP says, “I don’t care”. Is the network or destination host overwhelmed with the data that I transmit? “I don’t care”.

| UDP might be lazy, but it is inexpensive. The small header size means that there is low overhead. The lack of services means that UDP can transfer data with lower delay. It is not reliable, but it is quick. This is perfect for applications like Voice over IP (VoIP) or video streaming. These applications do not need all of TCP’s reliability and they do not want to pay for it.

Programmers who build applications can simply build the needed reliability into the application itself.

In UDP you can use up to 65535 ports however the well-known applications use the first 1024 ports. From 1025 to 65525 ports are known as ephemeral Port numbers, same concept as TCP.

Question #1

The OSI Model

Which of the options below accurately describe MAC addresses?

a. They are used at Layer 2 of the OSI model
b. They are used at Layer 3 of the OSI model
c. OUI is contained in the last 3 bytes of the address.
d. They are 48-bits long
e. They are used for Ethernet and Wi-Fi technologies, among others.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

Correct answers:
a.
d.
e.

Networking Devices

Switches, Routers, Multilayer Switches, APs, Firewalls, Servers

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

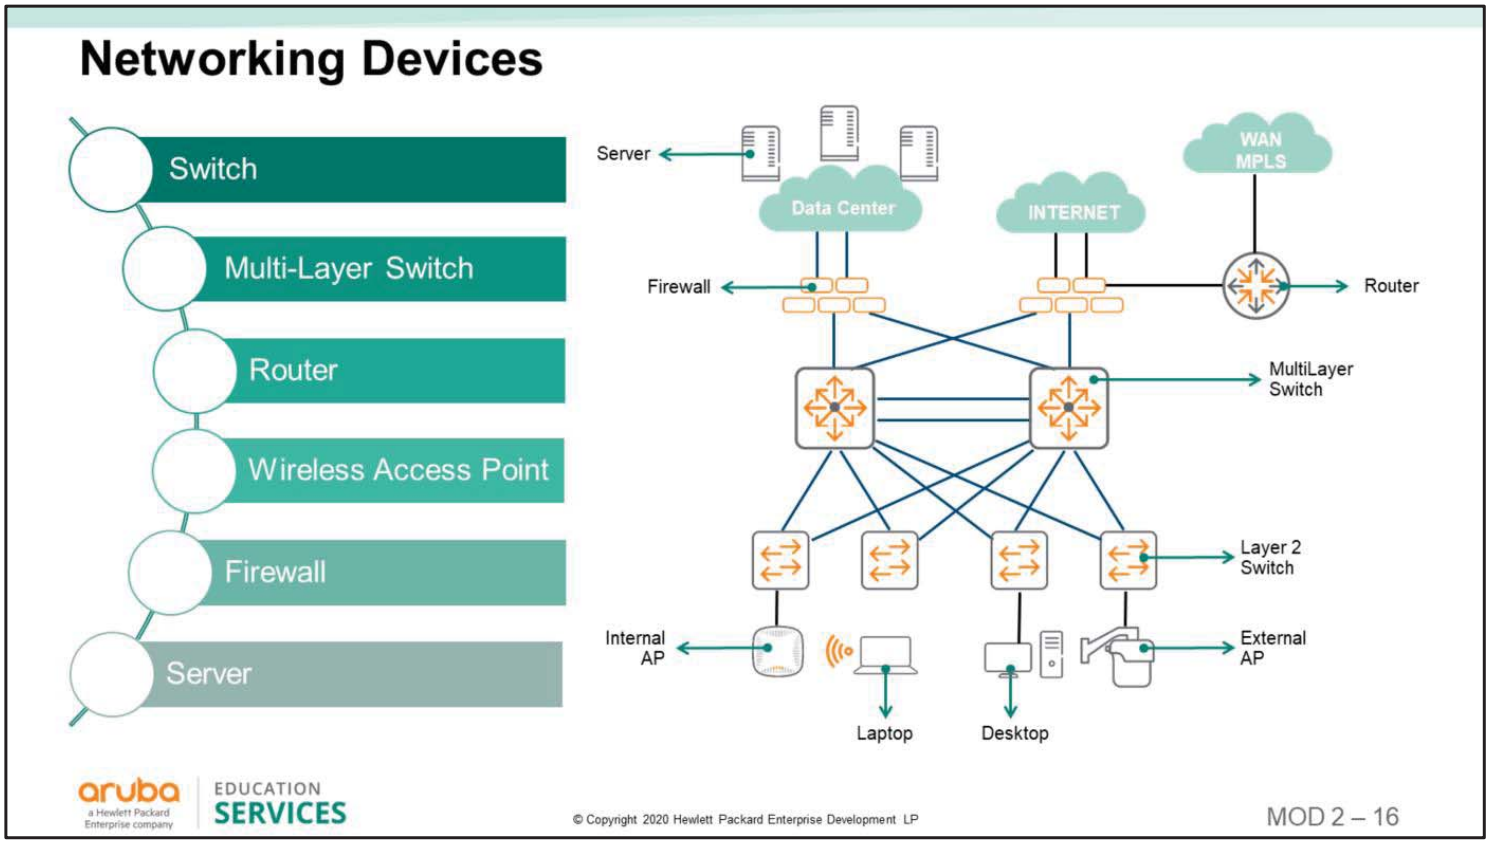

A corporate network that serves hundreds or thousands of computing devices uses different network elements to achieve proper digital communication. This section introduces you to the network devices that are typically deployed in a LAN Campus network. The figure shows a typical deployment of these devices and how they could interconnect.

Layer 2 Switch
Layer 3 Router
Multi-Layer Switch
Wireless Access Point
Firewall
Server

Networking Devices – Switch

Layer 2 device forwards Ethernet frames based on destination **MAC** addresses

Multiple ports connect endpoints

- PCs, printers, cameras, and more
- Commonly 8, 24, 48, or more physical ports
- Transparent to endpoints

Special L2 protocols

Performance, reliability, security

- STP
- LLDP
- 802.1q

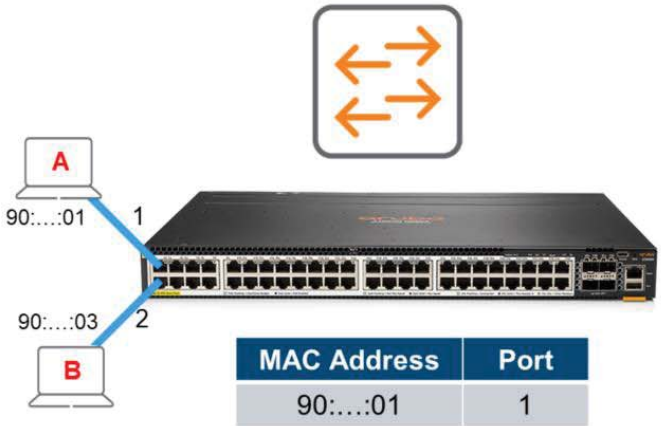

MAC Address	Port
90:...:01	1
90:...:03	2

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 17

Switches have multiple ports that are used to connect computing devices into one or more Local Area Networks (LAN). This can include PCs, printers, video cameras, Voice-over-IP (VoIP) phones, and more. You may encounter switches that have 8, 24, 48, or more ports. Understand that switches are “transparent” to endpoints. Connected devices are not aware of the existence of the switch – they perceive themselves as being directly connected to each other.

A switch is a Layer 2 network device that forwards Ethernet frames, based on destination MAC addresses. Recall that MAC addresses are part of the Data Link Layer, and so switches are considered Layer 2 network devices. The figure shows the Aruba icon used for a switch, and below that a physical switch example.

|You also see a simplified MAC address table that is maintained in switch memory. The switch has learned that a device with MAC address 90::01 is connected to port 1 of the switch, and a device with MAC address 90::03 is connected to port 2.

|Switches also have special languages or protocols that they use to increase network performance, reliability, and security. Examples of these protocols include the following:
Spanning Tree Protocol (STP)
Link Layer Discovery Protocol (LLDP)
802.1q

You will learn more about these protocols later in this course.

Networking Devices - Router

Layer 3 device that forwards packets based on destination IP address

Functionality and features

- Connect separate networks into an inter-network
- Offer WAN connectivity to networks

L3 routing protocols

Learn all possible paths, then choose a best path

- RIP
- OSPF
- BGP

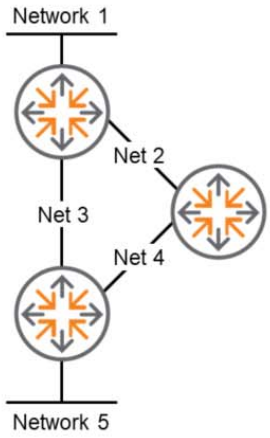

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 18

A router is a Layer 3 network device that forwards packets based on the destination IP address. Since IP addresses are part of the Network Layer, routers are considered Layer 3 devices.

|You know that switches connect computing devices into one or more networks. It is a router's job to connect those separate networks together, to create an inter-network. The Internet is the largest, most common example of an inter-network. The figure shows four routers, used to interconnect five networks (links between the routers are also consider networks).

While switches might have a few dozen or a couple hundred ports, routers have a relatively small port count – often between 2 and 6, which are used to connect to WAN networks.

When you take a long trip in your car, there are often several paths you could take to arrive at your destination. You consider all possible paths, and then choose the best path. You judge the best path based on some criteria that is important to you, such as whether the route is the fastest, or the shortest, or whether the route is particularly scenic. Then you drive along that path.

|Similarly, routers run protocols between them to learn all possible paths between all available networks. They then choose a best path for each destination, and forward packets along those best paths.

Examples of these routing protocols are:
Routing Information Protocol (RIP)
Open Shortest Path First (OSPF)

Border Gateway Protocol (BGP)

You will learn more about these protocols later in this course.

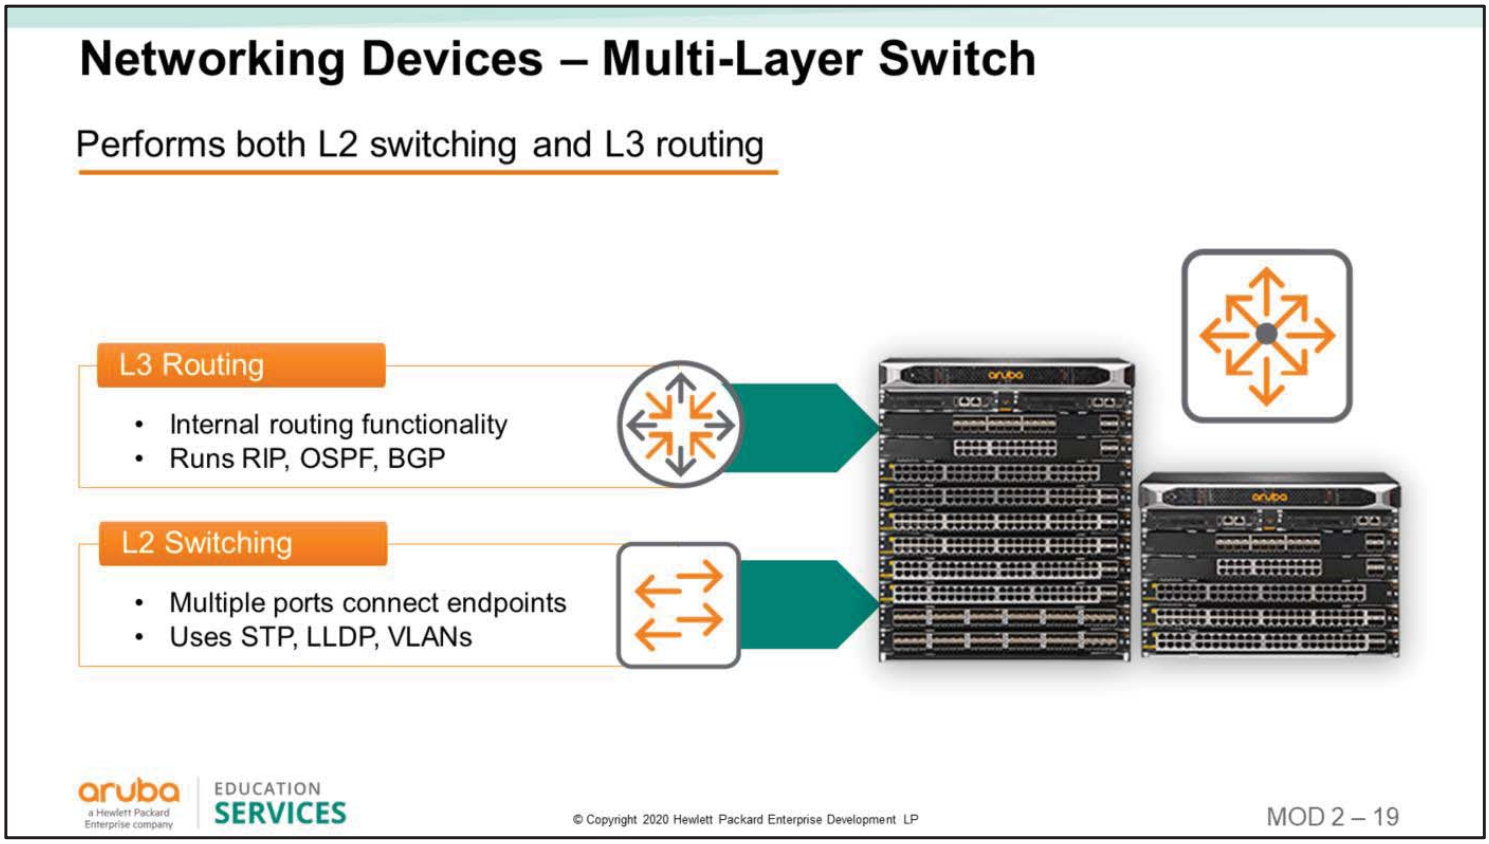

The OSI model sets clear distinctions between a Layer 2 switch and a Layer 3 router. However, networking professionals know that it is often good to have a single device that can perform both functions - a hybrid between a router and a switch.

A multi-layer switch has all the functionality of a Layer 2 switch.
|They support typical Layer 2 functionality, such as STP, LLDP, and VLANs. Multiple ports connect various endpoint devices together into one or more networks.

|A multilayer also has all or most of the functionality of a Layer 3 router, so it can route between those separate LANs internally - an internetwork in a box! Like any router, it supports routing protocols like RIP, OSPF, and BGP.

Of course, multilayer switches support Layer 1 to actually transmit and receive data, and it can even perform some Layer 4 functions – especially as relates to certain security features. These devices help to build a secure and flexible network.

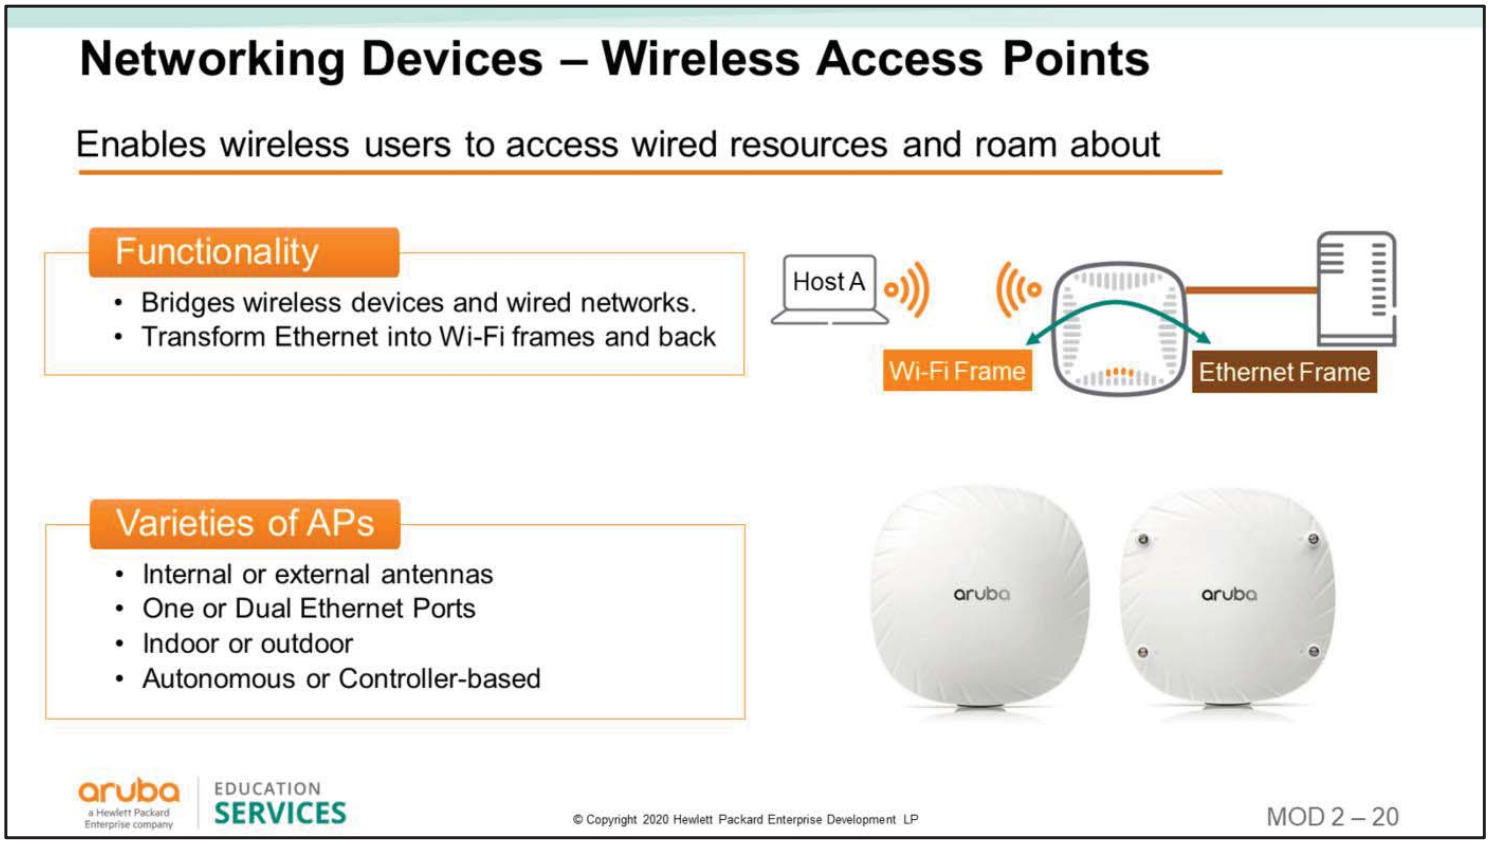

Access Points (AP) enable wireless users to access wired resources. Without being tethered to an Ethernet cables, users are free to roam about.
|APs are translational bridges – they accept Wi-Fi frames from endpoints, translate them into Ethernet frames, and forward them on to some wired resource. This could be an internal corporate server, storage, application, or an Internet resource. When that wired resource responds, the AP accepts that Ethernet frame, translates it to a Wi-Fi frame, and then transmits it over the air to the wireless host, as shown in the figure.

|There are many varieties of APs to meet your technical and budgetary needs. They may be built with internal or external antennas, with one or dual Ethernet ports. Some may be for indoor use only, while others may can be mounted outside.

Wi-Fi systems may be designed to use Autonomous APs or Controller-based AP. Autonomous APs can perform all necessary functions to create a functional system. They process wireless and ethernet frames and provide a certain amount of manageability and control. As the name implies, they work autonomously, without control from some external device. This type of deployment is relatively simple and easy, but it is not very scalable. It is suitable for smaller environments.

With Controller-based AP solutions, the APs send and receive Wireless and Ethernet frames, much like Autonomous APs. However, much of the processing and management functions are off-loaded to one or more centralized controllers.

These centralized configuration and control features might increase the initial complexity of the deployment, but once in place they provide far superior visibility into network health, enable more proactive network management, and can scale from medium to very large deployments.

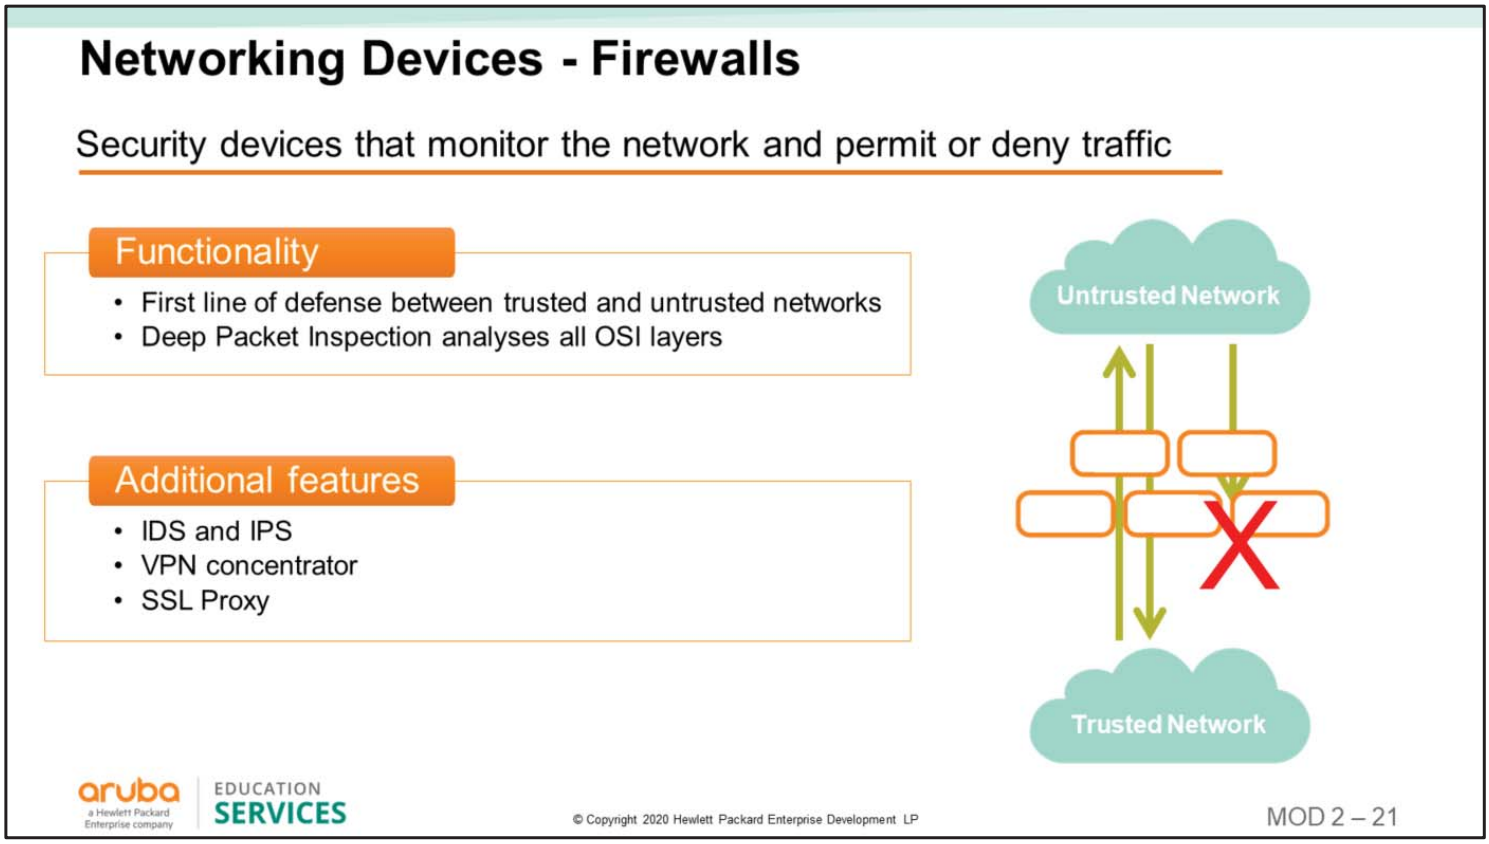

Firewalls are security devices that monitor and control network traffic, based on security rules.

|Appropriate traffic is permitted, while suspicious traffic is denied. and determine whether each packet should be permitted or denied, based on security rules.

|Firewalls are commonly deployed as the first line of defense between trusted networks (managed by a corporation) and untrusted networks such as Internet, permitting only the valid connections.

Many firewalls can inspect all OSI Layers, permitting engineers to create elaborate rules based on applications. Compared with older firewalls that only understand up to OSI model Layer 4, these newer firewalls can get deeper into the packet headers – all the way to layer 7. This is known as Deep Packet Inspection.

|Some firewalls can integrate other functions such as:
Intrusion Detection Systems (IDS)
Intrusion Prevention Systems (IPS)
VPN Concentrators
SSL proxy

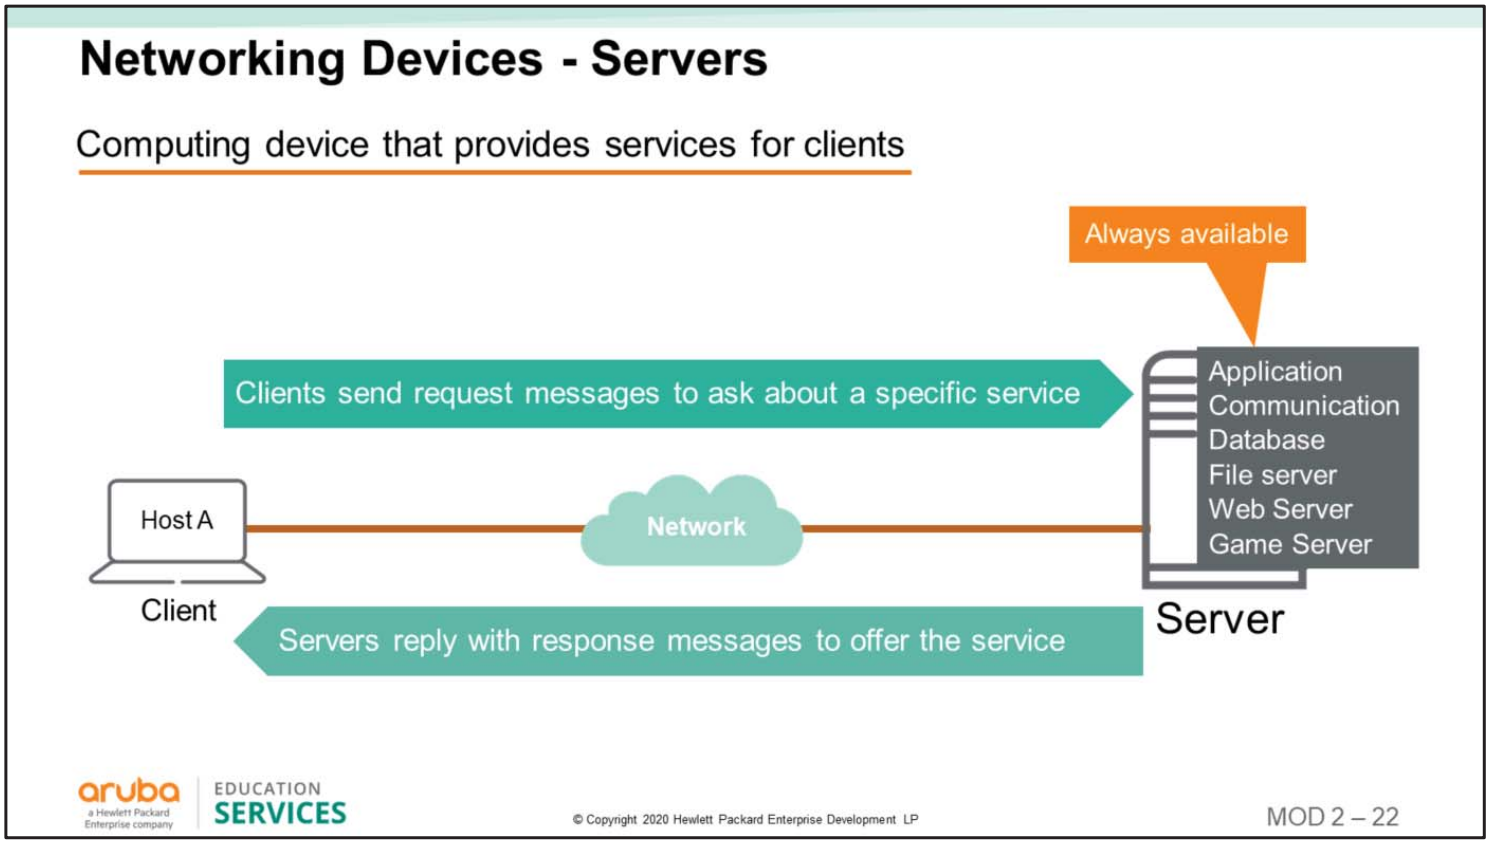

A server is a computing device that provides services for other programs or devices called clients.
|The underlying switches, routers, APs and firewalls facilitate this client-server communication.

|Clients often send request messages to servers to ask about a specific service. Servers reply with response messages to offer the service.

|Clients can contact a server anytime, and servers use hardware and software that is designed to be available all time.

Servers are often classified based on the services that they provide.
|Some examples include the following:
Application Servers
Communication Servers
Database Servers
File Servers
Web Servers
Game Servers

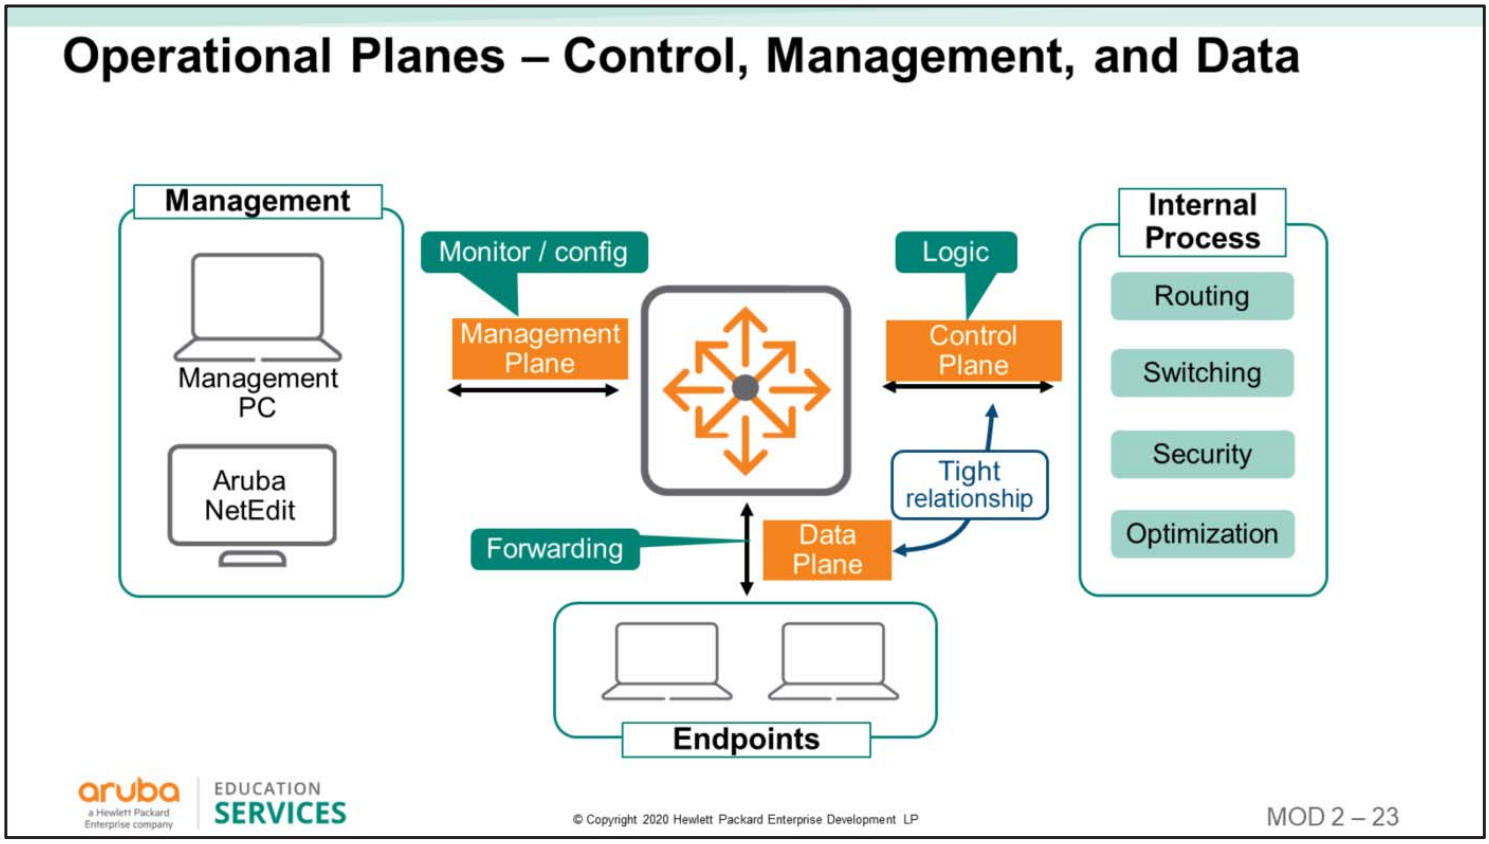

A network device is logically composed of three operational planes, and each plane performs specific tasks.

Data Plane. The data plane receives and sends frames using specialized hardware called Application-Specific Integrated Circuits (ASIC), which is much faster than using software. ASICs modulate and demodulate data, and handle other functions related to frame transmission and receipt.

Control Plane. The control plane logic determines what to do with the data that has been received. These decisions are made with internal processes - routing, switching, security, and flow optimization. Data plane and Control planes have a tight relationship, to process any data as fast as possible.

Management Plane. You use the management plane to monitor and configure the device. This plane must be separate from the data plane, for security and accessibility reasons. You must be able to access the device even if the control and data planes fail. Also, you don't want end users to gain access to the management plane, this could be an egregious security issue.

Note. AruabOS-CX devices have a specific Interface and Virtual Routing and Forwarding (VRF) instance that is used for Out-of-Band Management which maintains a total separation from the data plane. You will learn more about VRFs in future modules.

Question #2

Networking Devices

Which components and concepts are directly focused on Layer 2 communications?

a. Switch
b. Router
c. Multi-Layer Switch
d. MAC addresses
e. IP addresses
f. Access Points

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 24

Correct answers:
a.
c.
d.
f.

Question #3

Networking Devices

Which components and concepts are directly focused on Layer 3 communications?

a. Switch
b. Router
c. Multi-Layer Switch
d. MAC addresses
e. IP addresses
f. Access Points

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 25

Correct answers:
b.
c.
e.

Common Networking Services

DHCP, DNS, HTTP, Telnet, SSH, FTP

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 26

This section will introduce you to several common network services.

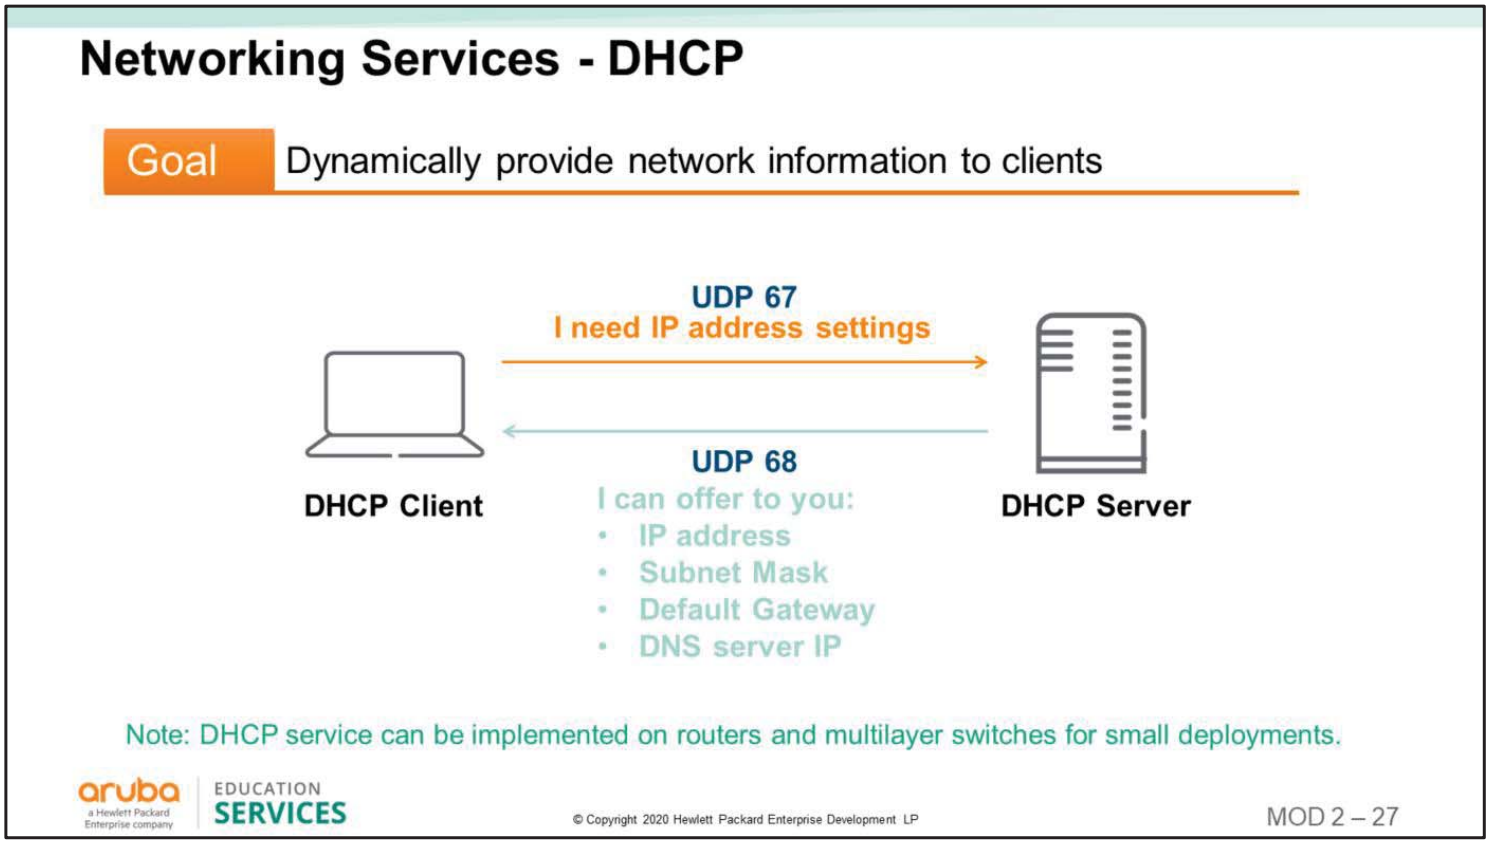

Challenge
Every device on a TCP/IP network requires an IP address, which can be manually configured. But imagine the overhead of manually assigning IP addresses to hundreds or thousands of hosts. Even worse, if you're laptop has a statically assigned address appropriate for work, it will unlikely be the correct information when you then connect at home, or at your favorite café.

Solution
DHCP dynamically assigns network information to clients. Thus, when a host boots up, it automatically acquires its IP address, subnet mask, default gateway and other configurable parameters.

|Clients typically initiate the communication by broadcasting a DHCP query on the network, "Attention everyone! I need an IP address".
|The DHCP server receives this query and replies with an offer. If the offer is valid, the client uses the provided information.

Recall that both UDP and TCP headers include a destination port number, which indicates the type of application data that is received. Clients use **UDP port 67** for the DHCP query and servers reply with an offer on **UDP port 68**.

For larger, enterprise-class deployments, this DHCP service is often implemented on a server.
|For smaller deployments, DHCP can be implemented on routers and multilayer switches.

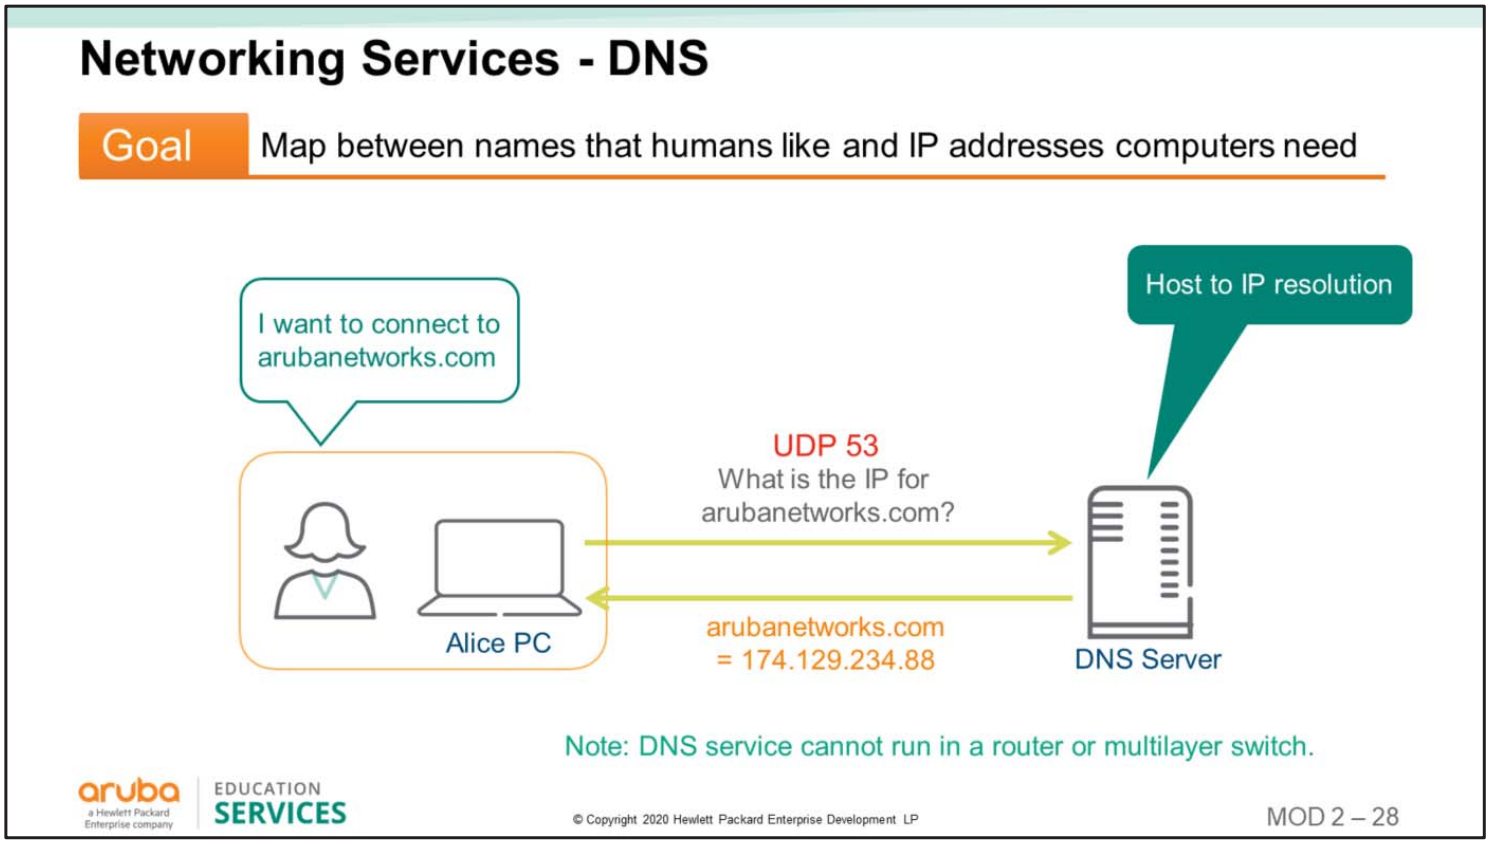

Challenge

To send or receive data, all devices in a computing network must have a valid IP address. In a typical internal network with hundreds or thousands of devices, two problems arise.

First, how can users identify the target computing device to which they must connect? And secondly, how can users know the IP address of that target destination device?
To solve the first problem, endpoints can be assigned an intuitive name, like “fileserver1” or “PC-of-John”. Also, groups of devices can exist in a common domain, like www.arubanetworks.com or www.hpe.com.

This naming makes computer systems more user friendly for humans, but the underlying network devices still require Layer 3 IP addresses to communicate. This brings us to the second problem – how can users identify a device by its name, and have that name translated to the correct IP address?

Solution

Domain Name Service (DNS) maintains a mapping between the names that humans like and the IP addresses that computers need.

|The client device sends a unicast request to the DNS server on **UDP port 53**, “What IP address is associated with www.hpe.com?” The server performs a name lookup in its database, finds a matching record, and responds to the client with the associated IP address.

DNS uses UDP for smaller, more common operations like name queries. However, DNS

occasionally use TCP port 53, for larger operations like a “Zone Transfer”.

|Unlike DHCP service, DNS cannot run in most network devices like routers or multilayer switches. It is recommended to use a dedicated name server for this purpose.

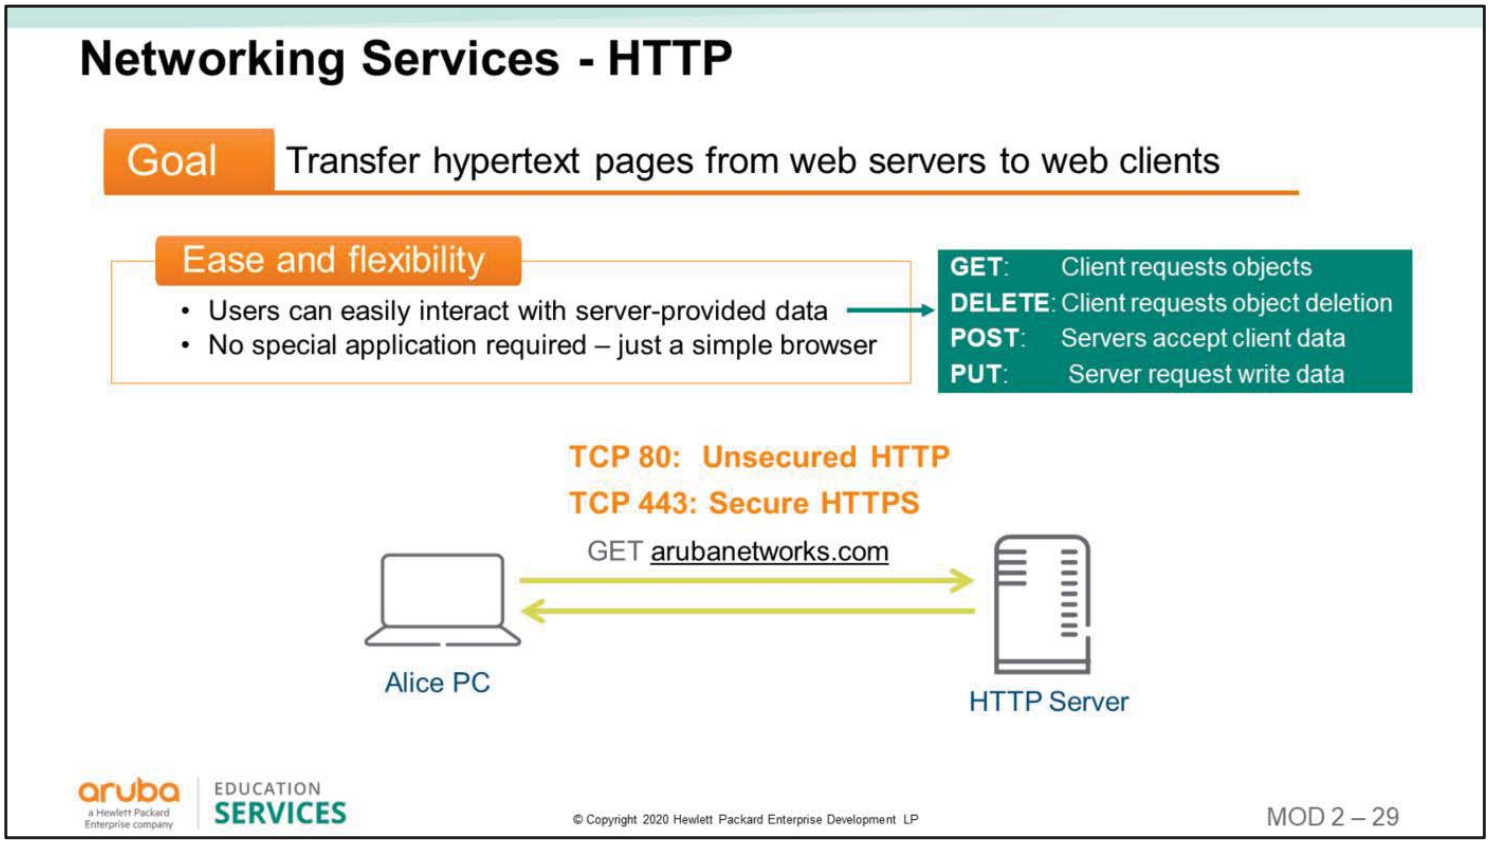

The HTTP protocol is used to transfer hypertext pages from web servers to web clients. Hypertext pages are documents with tags or links. When users click on these tags they gain access to new pages. Typically, tags are defined by a language such as the Hyper-Text Markup Language (HTML) or the eXtensible Markup Language (XML).

|HTTP access methods provide a flexible communication mechanism. Users can easily interact with server-provided data, using one or more access methods, including the following:

- GET: Clients request specific objects from the server
- DELETE: Clients request deletion of a specific object.
- POST: Servers accept the data enclosed in the body of the request message, mostly to store it.
- PUT: Client request write specific data on the server.

Clients do not require specialized applications – only a simple web browser is used to establish HTTP sessions.

Although HTTP is a well-known protocol, it can be dangerous to use it when browsing public internet sites. This is because HTTP provides no security mechanisms – your activity can be monitored and manipulated by hackers. Are you really accessing your bank's web site, or a fake site stood up by hackers? Is someone viewing or copying the data that I transmit or receive?

These issues are addressed by using HTTPS, the secure version of HTTP.

HTTP sessions are established on **TCP port 80** while HTTPS uses **TCP port 443**.

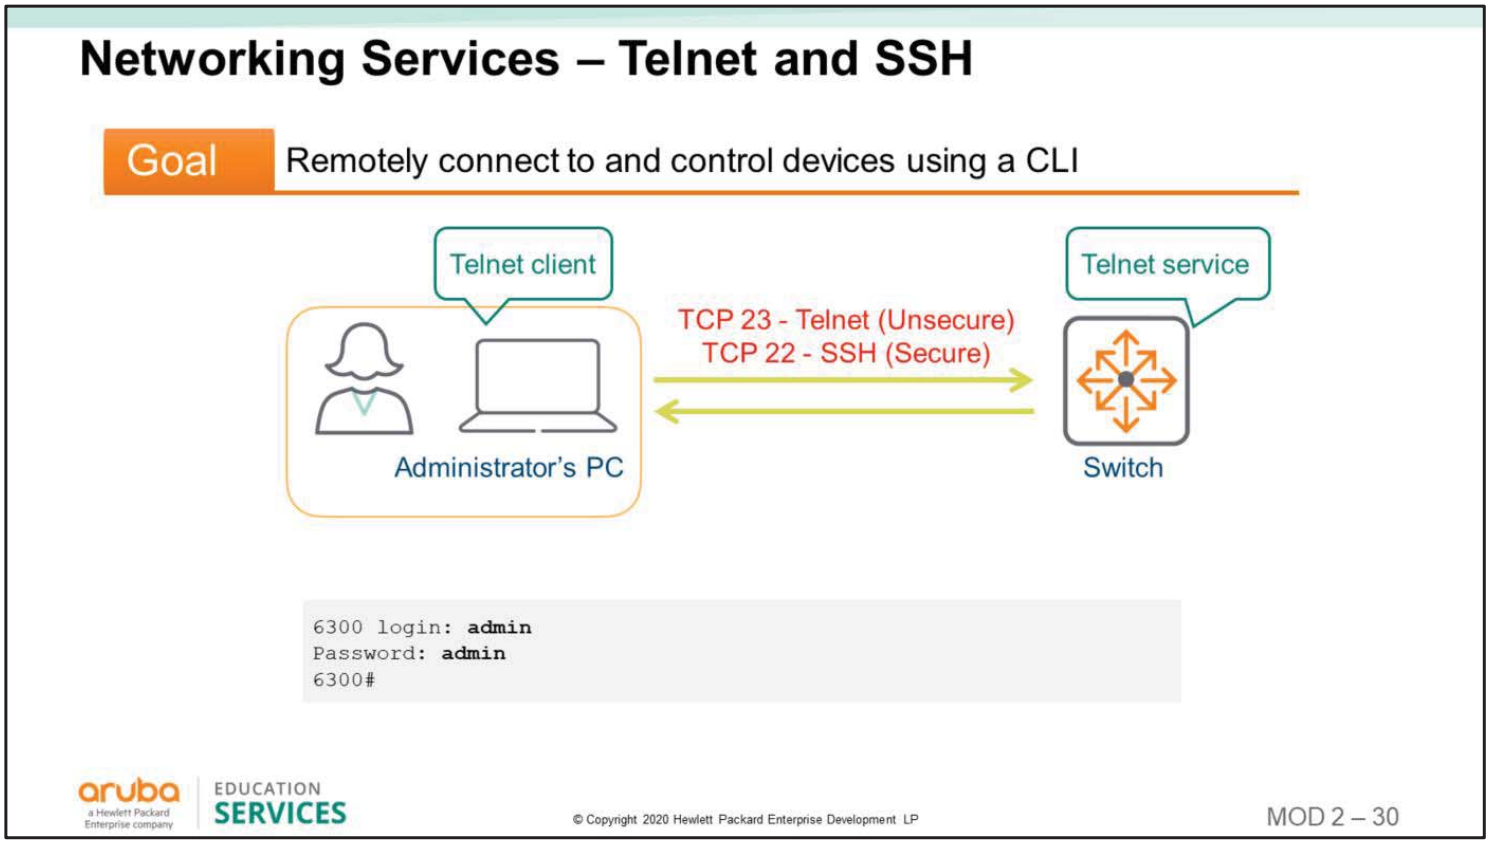

Challenge

Network administrators may need to configure and troubleshoot dozens, hundreds, or even thousands of network devices. To directly, physically connect to them might require a hike across campus or flying to another city. Perhaps you are in your office, and you need to configure or view the status of a router or switch. Instead of walking across the campus, going up to the 12th floor, and sitting in a small network closet, you can simply Telnet to the device.

Solution

|The Telnet protocol enables you to remotely connect to and control other devices, using that device's Command Line Interface (CLI). This protocol does not provide for the use of a graphic interface.

First you establish a telnet session, using a telnet client such as the common PuTTY application. Then you send commands that are executed on the remote device. The remote device must be running a telnet service. In the example shown, the administrator has remotely attached to the switch and issued the command show mac-address-table. This reveals the Layer 2 MAC addresses of all devices connected to the switch, and to which ports they are connected. You will learn more about "show" commands later in this module.

Like HTTP, Telnet has no security mechanisms. Hackers can intercept your sessions and harvest your data – including usernames and passwords. For this reason, many people avoid using Telnet, and use the Secure Shell (SSH) instead. Other than SSH's security and encryption mechanisms, the functionality is about the same as telnet.

Telnet uses TCP port 23 while SSH uses TCP port 22 to establish the session.

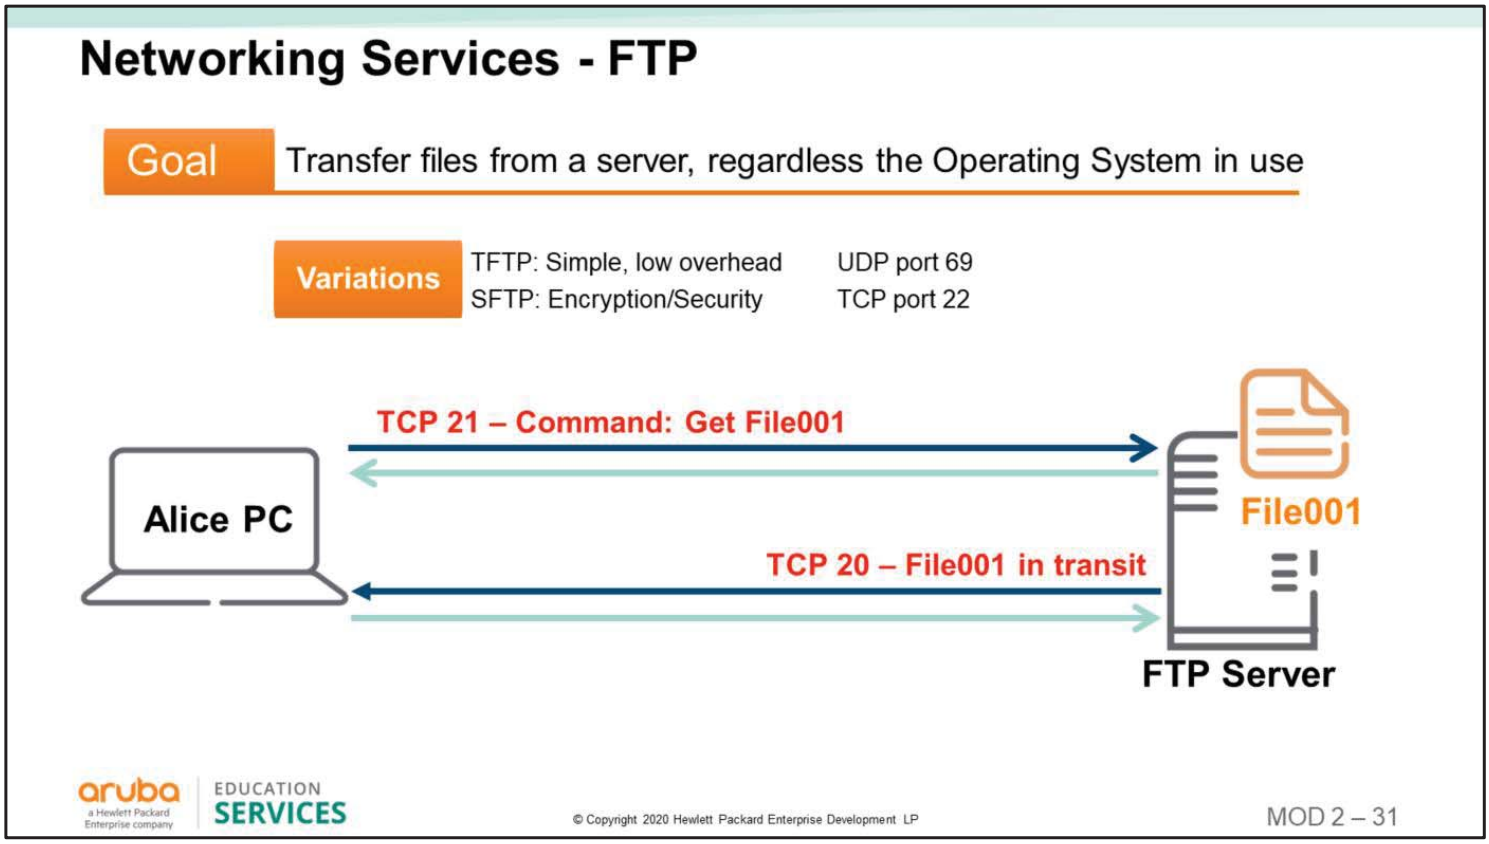

FTP enables you to upload and download files from a server, regardless of the operating system they are using.

FTP uses TCP as the control protocol – a reliable, flow-controlled mechanism that ensures complete file transmission. FTP uses TCP port 20 for Data and TCP port 21 for Control.

FTP provides for user authentication, where you enter valid credentials to access the FTP server.

Clients could use dedicated applications to establish FTP connections, however a simple web browser also works. If you want to use a web browser, make sure you use the correct URL. When you browse to a web site, the URL always starts with http://. When you connect via FTP, the URL always starts with ftp://.

In the figure, Alice connects to and authenticates the FTP server. Her username is alice@arubalab.com. She sees a list of directories and files. She selects the file named File001, and chooses to download it. This results in an FTP GET action.

|In essence Alice's PC connects on the FTP control port 21 and makes a file request, "I need File001".
|Then the actual file is transferred from server to PC using the FTP data session – TCP port 20.

|There are two common FTP variations. The Trivial File Transfer Protocol (TFTP) works like FTP, except it relies on UDP instead of TCP. UDP has few headers and lower overhead, so the transfer will be faster. However, since UDP has no reliability mechanisms, required

reliability is built into the TFTP application. TFTP uses UDP port 69.

Another variation is Secure File Transfer Protocol (SFTP), which offers enhanced security and encryption not available with FTP. SFTP runs over an SSH session, usually on TCP port 22.

Question #4

Network Services

Which of the following statements accurately describe common networking services?

a. Use DNS to dynamically provide network addressing information to clients

b. Use DHCP to map names to IP addresses

c. All common network services use either TCP or UDP.

d. The advantage of HTTPS over HTTP is that it is more secure.

e. You can use SSH to gain remote access to a switch CLI interface.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 32

Correct answer:
d.

Lab Activity

Lab 2: Packet Exploration

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Additional Content

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Wi-Fi Frames

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Wi-Fi Frame

Wi-Fi is a technology used in Layer 1 and 2 for wireless communications.

Wi-Fi uses 2 extra frames besides data: Control and Management

Wi-Fi Header								Wi-Fi Trailer
Frame Control	Duration ID	Address 1	Address 2	Address 3	Sequence Control	Address 4	Payload	FCS
2 Bytes	2 Byte	6 Bytes	6 Bytes	6 Bytes	2 Bytes	6 Bytes	46 – 1500 Bytes	4 Bytes

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 36

Wi-Fi technology is used at Layer 1 and Layer 2 for wireless communications. Remember, Ethernet only works for wired connections. One major difference between these technologies is that Wi-Fi introduces two extra frames besides of the data frame: control and management. This is because there is special communications between an end system and a Wireless Access Point (AP) that is not directly related to the transfer of payloads. These frames are used to manage and control the wireless network itself.

Note: This text will only discuss the wireless data frame. To learn more about the other wireless frames, refer to the Aruba Mobility Fundamentals training.

The fields shown in the figure are described below:

Frame Control. Indicates the type of frame (Control, management or data) also includes fragmentation and privacy information.
Duration/Connection ID. If used, indicates the time in microseconds the channel will be allocated for successful transmission of a control frame.
Addresses. MAC addresses for the devices that are participating in the communication. The number of addresses depends on frame type and on context. Typically, only the first three address fields are used.
Sequence Control. This field includes information about the fragmentation and reassembly.
Payload. Contains the data to be transmitted.
FCS (Frame Check Sequence). Allows detection of corrupted data using the CRC or checksum method.

Understanding the Wi-Fi Header

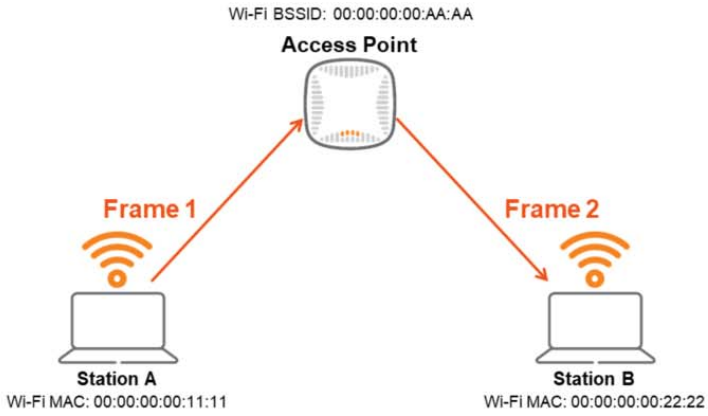

	Address 1 (receiver)	Address 2 (transmitter)	Address 3	Address 4
Frame 1	BSSID (AP) 00:00:00:00:AA:AA	MAC of Station A 00:00:00:00:11:11	MAC of Station B 00:00:00:00:22:22	Not Used
Frame 2	MAC of Station B 00:00:00:00:22:22	BSSID (AP) 00:00:00:00:AA:AA	MAC of Station A 00:00:00:00:11:11	Not Used

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 2 – 37

To understand how the address fields are used, consider the following example.

Stations A and B want to communicate each other. Each station has a specific MAC address assigned to their Wireless NIC – Station A has 00:00:00:00:11:11 while Station B has 00:00:00:00:22:22. This is a typical scenario, where an Access Point (AP) acts as a central device. Stations do not communicate directly with each other – they only communicate via the AP.

Therefore, 2 frames are generated - one from station A to the Access Point and other from the Access Point to Station B. The table summarizes how the MAC addresses are used in this scenario.

You see that Station A is the wireless transmitter of this frame, so its MAC address is placed in the frame's Address 2 field. The AP is the wireless receiver of the frame, so its address is placed in the frame's Address 1 field. Station B is the ultimate destination, so its address is in the Address 3 field. The Address 4 field is not used. This is typical. This field is typically only used in special "repeater" scenarios, where frames must pass between multiple APs before arriving at their destination.

The AP receives this frame, creates a new frame, and transmits it to Station B. In this case the AP is the transmitter, so its address is in the Address 2 field. Station B is the wireless receiver of the frame, and so its MAC address is in the Address 1 field. Station A is the original sender of the frame, and so its address is in the Address 3 field.

Note. In Wi-Fi, the MAC address associated to a Wireless LAN and a specific radio band is known as the Basic Service Set Identifier (BSSID).

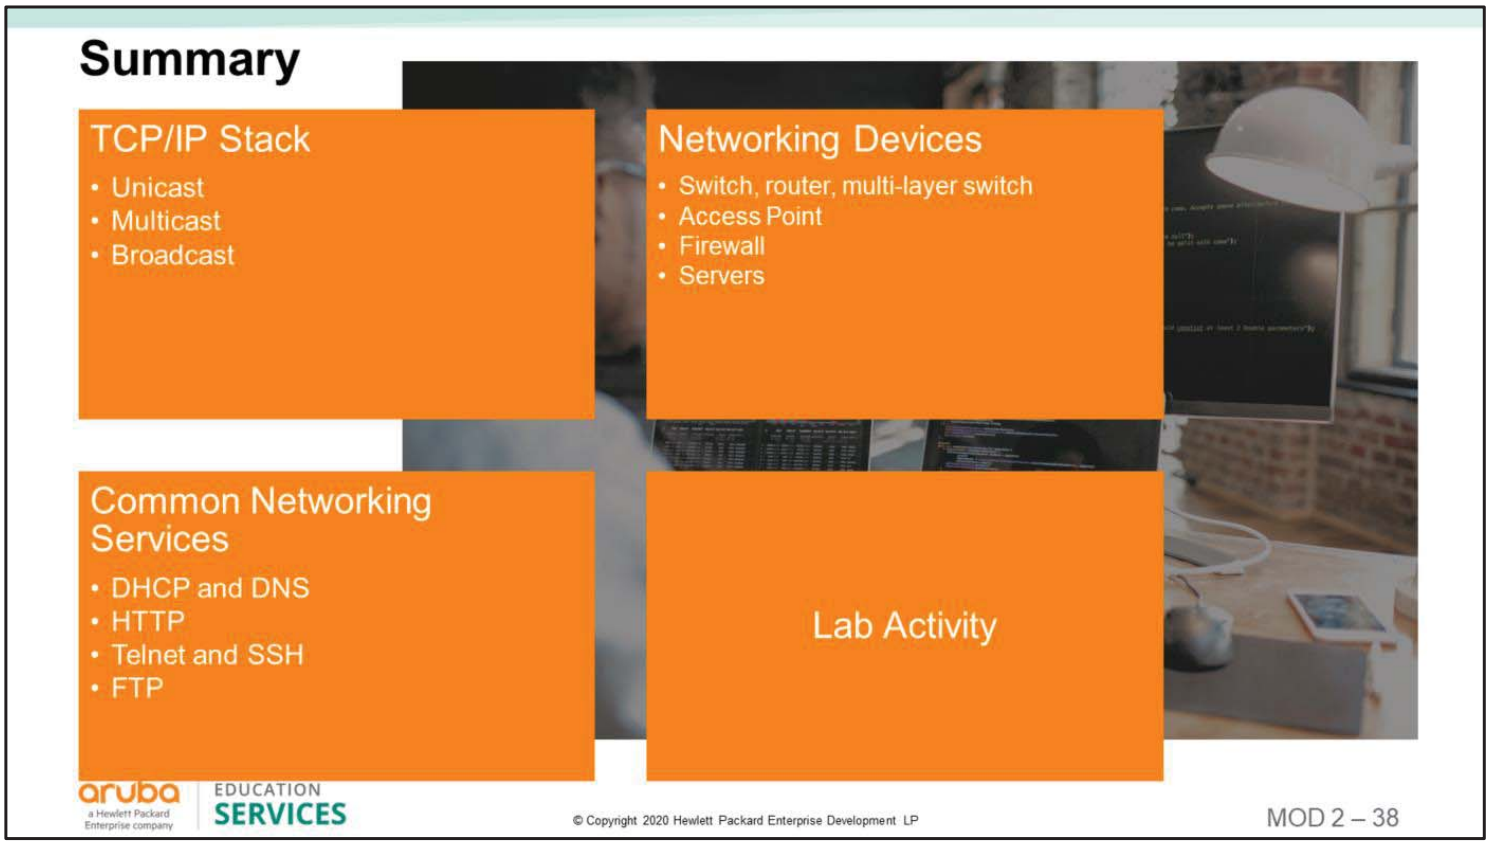

You began this module by exploring the networking devices used to build a typical network. This knowledge can help you to design, implement, and troubleshoot networks. You learned about important network services that are vital for proper network function and end-user interactions, along with exploring 2-tier and 3-tier network design models. You reviewed Aruba's portfolio of switches, to help ensure that you can select the proper device for a given scenario.

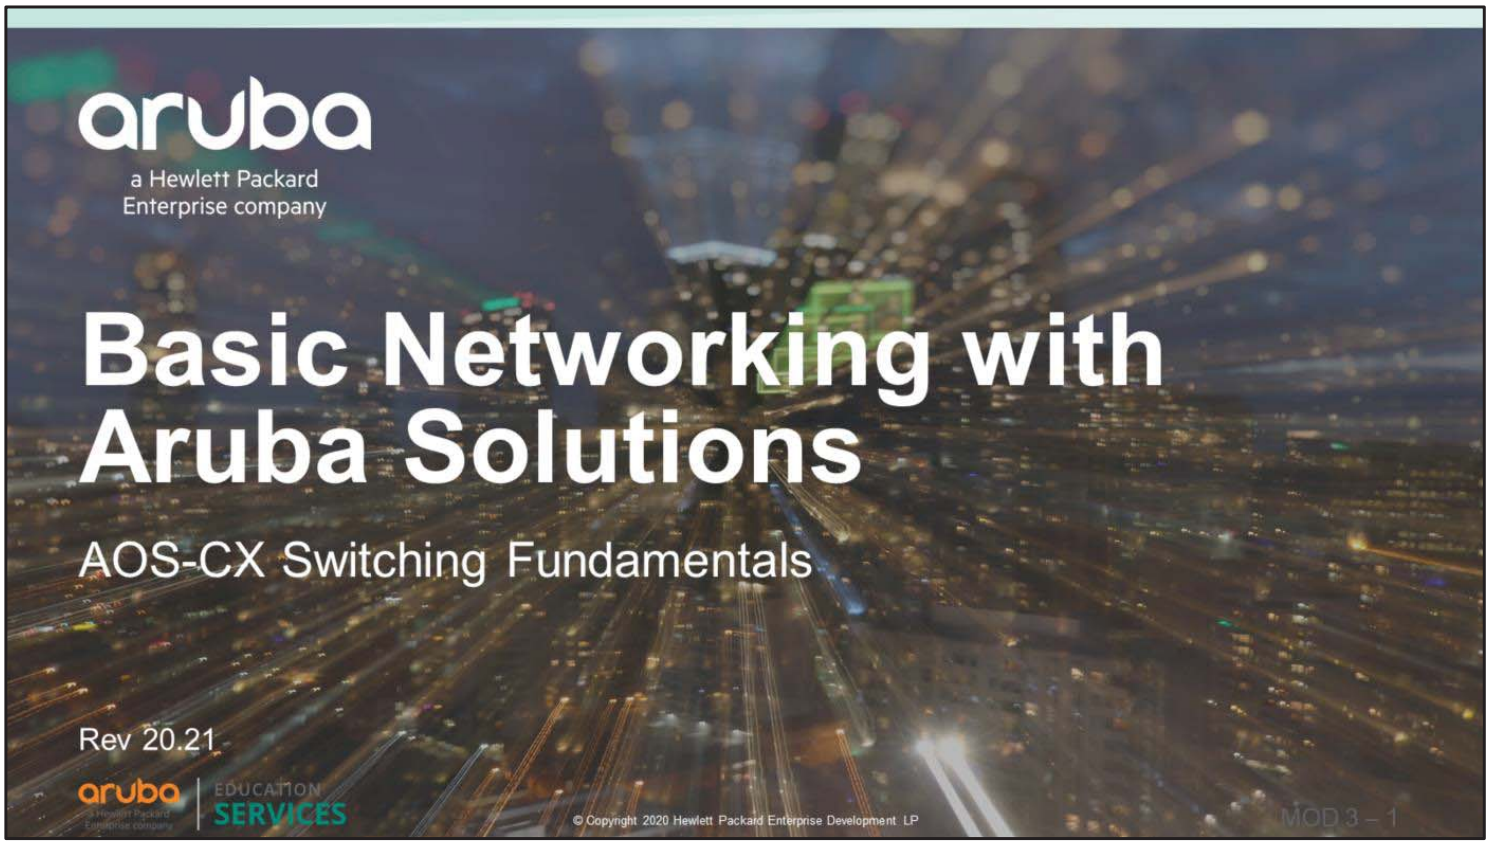
aruba
a Hewlett Packard Enterprise company
Basic Networking with Aruba Solutions
AOS-CX Switching Fundamentals
Rev 20.21
aruba EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 3 – 1

Objectives

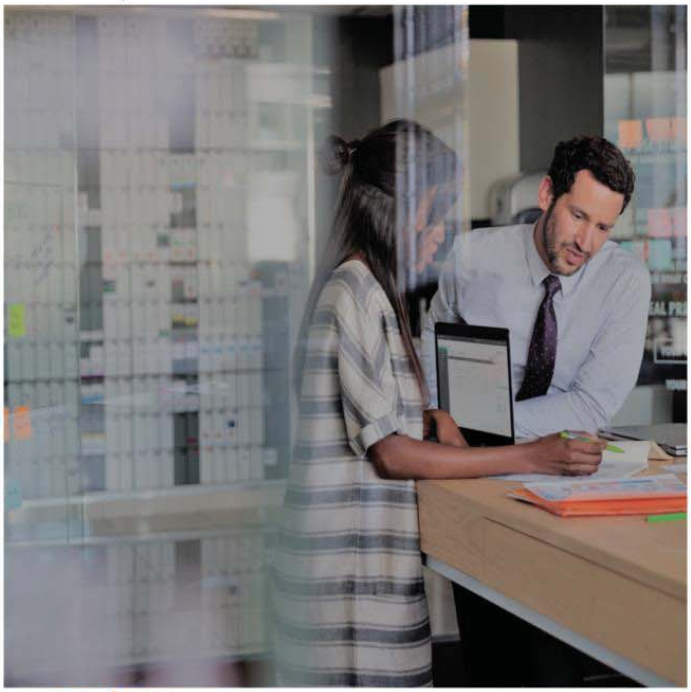

–Learning Objectives

- Explain network hierarchical models
- Compare the Aruba switch portfolio
- Describe Aruba OS-CX Operating system
- Discover switch information with CLI show commands
- Perform basic configuration tasks
- Discover network status and connectivity

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 2

You will be able to explain basic concepts related to hierarchical network design models, including 2-tier and 3-tier designs. Then you will explore the Aruba switch portfolio – models, specifications, and use cases. You will be able to describe the AOS-CX Operating System connectivity, port numbers, and navigating the Command Line Interface (CLI).

You will expand on this knowledge by understanding how to use basic CLI commands to gather system health, status, and connectivity information, perform basic configuration, and discover enough information to make detailed network documentation.

You begin this module by understanding 2-tier and 3-tier network design models. This will give you a good network design and troubleshooting foundation on which to build.

Next you will review Aruba's portfolio of switches. This will help you to select the proper device for a given scenario.

Then you will learn how to use Aruba switches – how to connect to them and use the CLI to glean basic network health and status information, perform basic configuration, and discover how an unknown network is connected. This ability to document a network will serve you well when you must diagnose network outages.

.

Network Design

2-Tier and 3-Tier Hierarchy

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

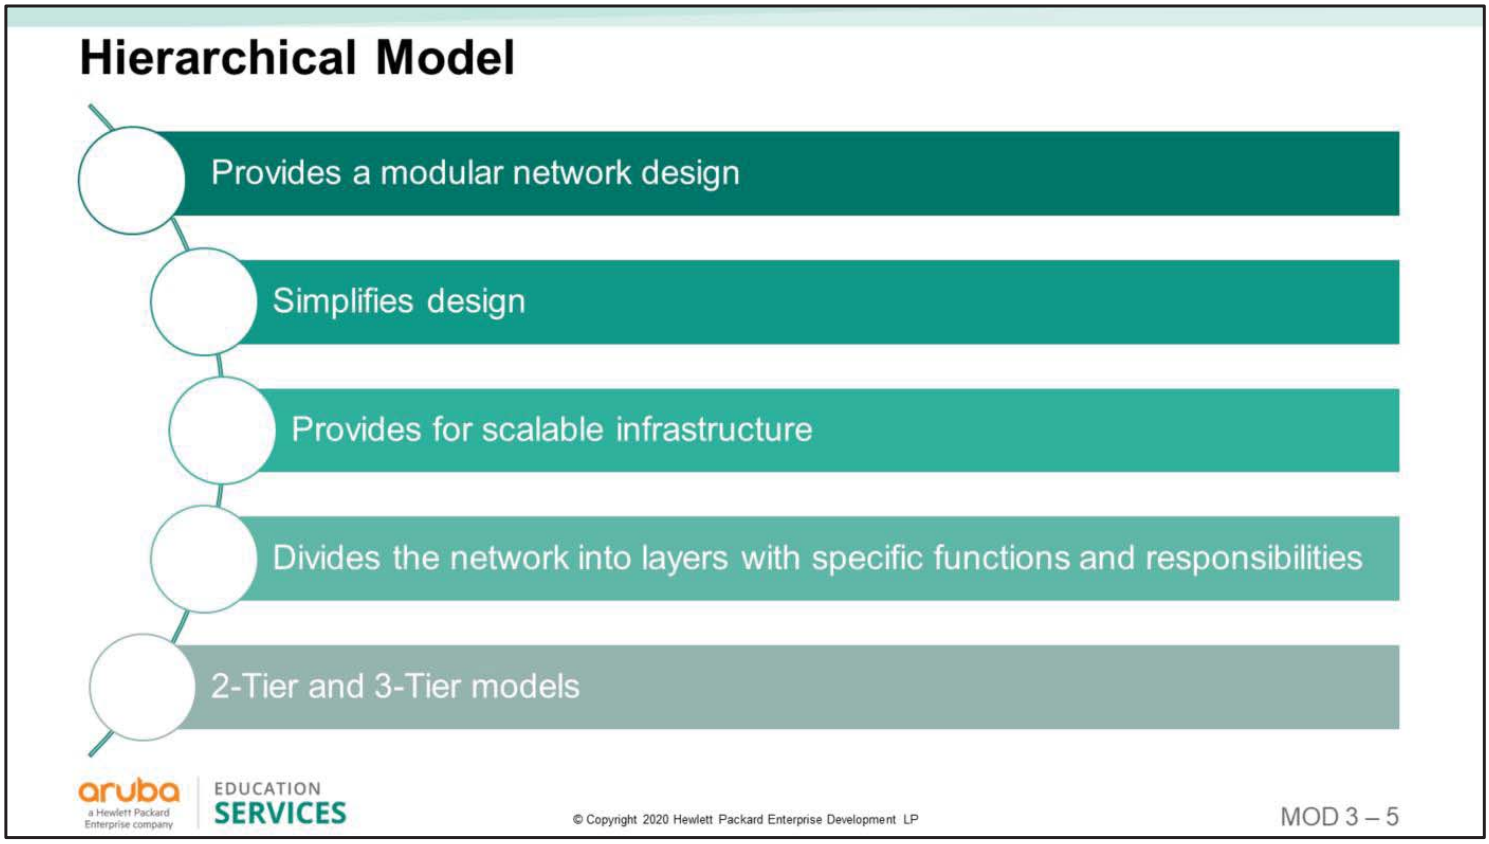

A hierarchical model provides a modular view of the network, simplifies the design, and provides a scalable infrastructure.

The hierarchical model divides the network into layers, and each layer has specific functions and responsibilities. This helps to ensure that you have an available, fault tolerant, flexible, and secure network.

Some engineers may choose a two-tier hierarchy, while others might choose a three-tier hierarchy. This decision is based on various factors:

- Number of wired and wireless users
- Mobility requirements
- Number of wireless Access Points deployed
- Cabling
- Number of buildings in the campus
- Desired security level
- Features and protocols needed

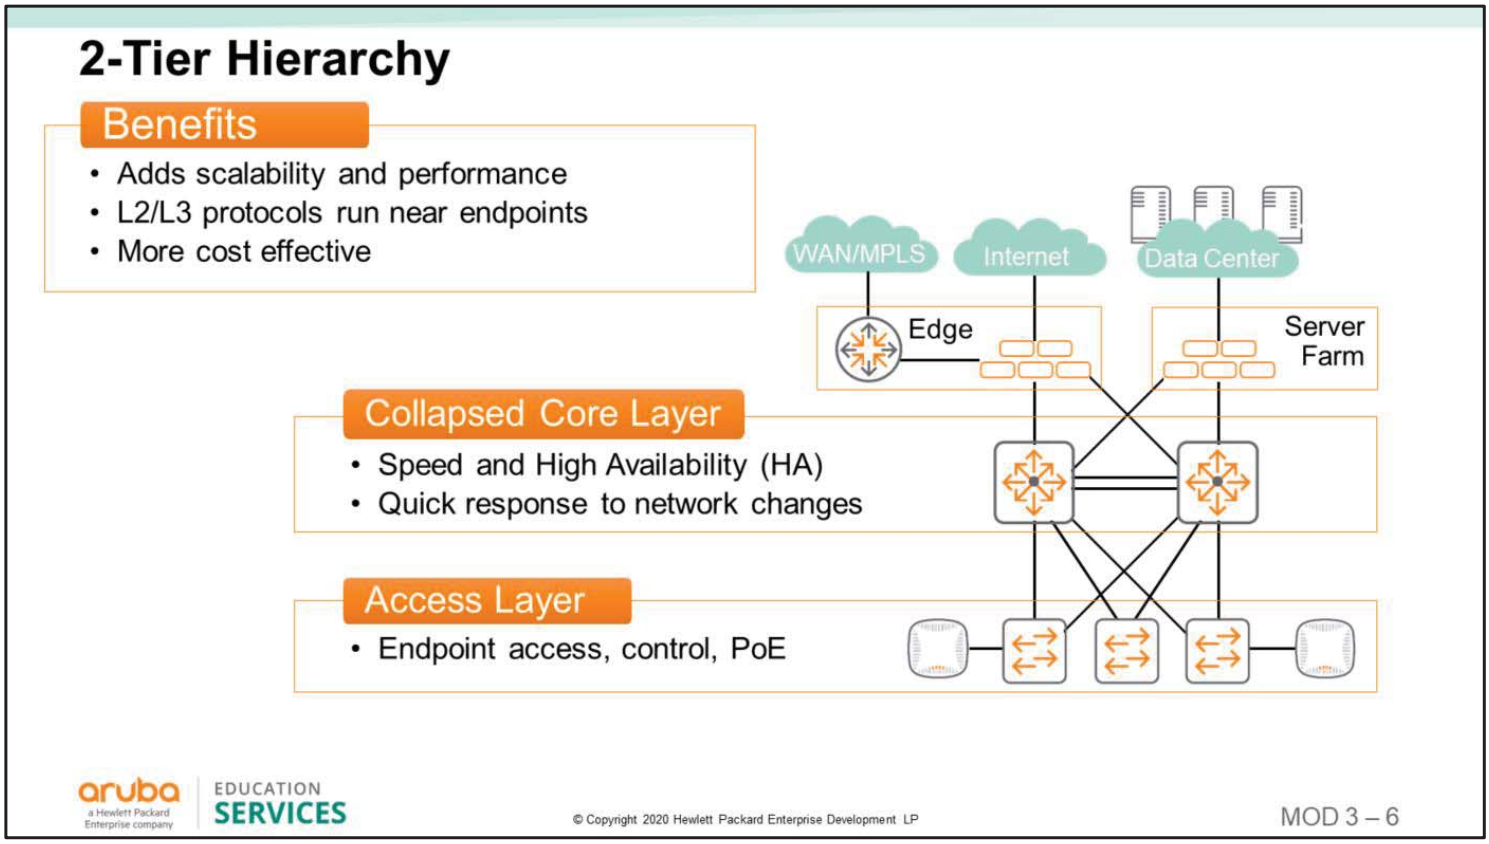

A two-tier design divides the network in 2 layers: Access and Core. This means that Layer 2 and Layer 3 protocols run close to the endpoints. This design can secure the network by deploying Access Control Lists (ACLs). Quality of Service (QoS) can analyze and prioritize traffic. This ensures the best performance for the most business-critical and delay-sensitive applications.

Wireless Networks can also benefit from this two-tier design. Multilayer access switches receive and process critical wireless traffic from APs. Appropriate security and QoS policies are applied, and the traffic is routed toward its destination.

|You want three key things at the Core Layer – speed, speed, and speed! This core must switch and route packets as fast as possible. This is why most of the more processing-intensive end-user services are off-loaded to the Access layer. You also need reliability at the core, so High Availability (HA) is also critical, as is the ability to quickly respond to changes in the network.

|Access Layer switches have redundant connections to the core to improve resiliency and mitigate outages. The Access Layer provides endpoint device connectivity, with Power over Ethernet (PoE), and several Layer 2 and Layer 3 features.

Other sections of the network may also connect to the core.
|For example, the Edge Network contains network devices that provide secure connectivity to WAN links and the Internet.
|The Server Farm contains racks of physical and/or virtual servers and their associated switches and hypervisors. Note: ACLs and QoS are out of the scope for this training.

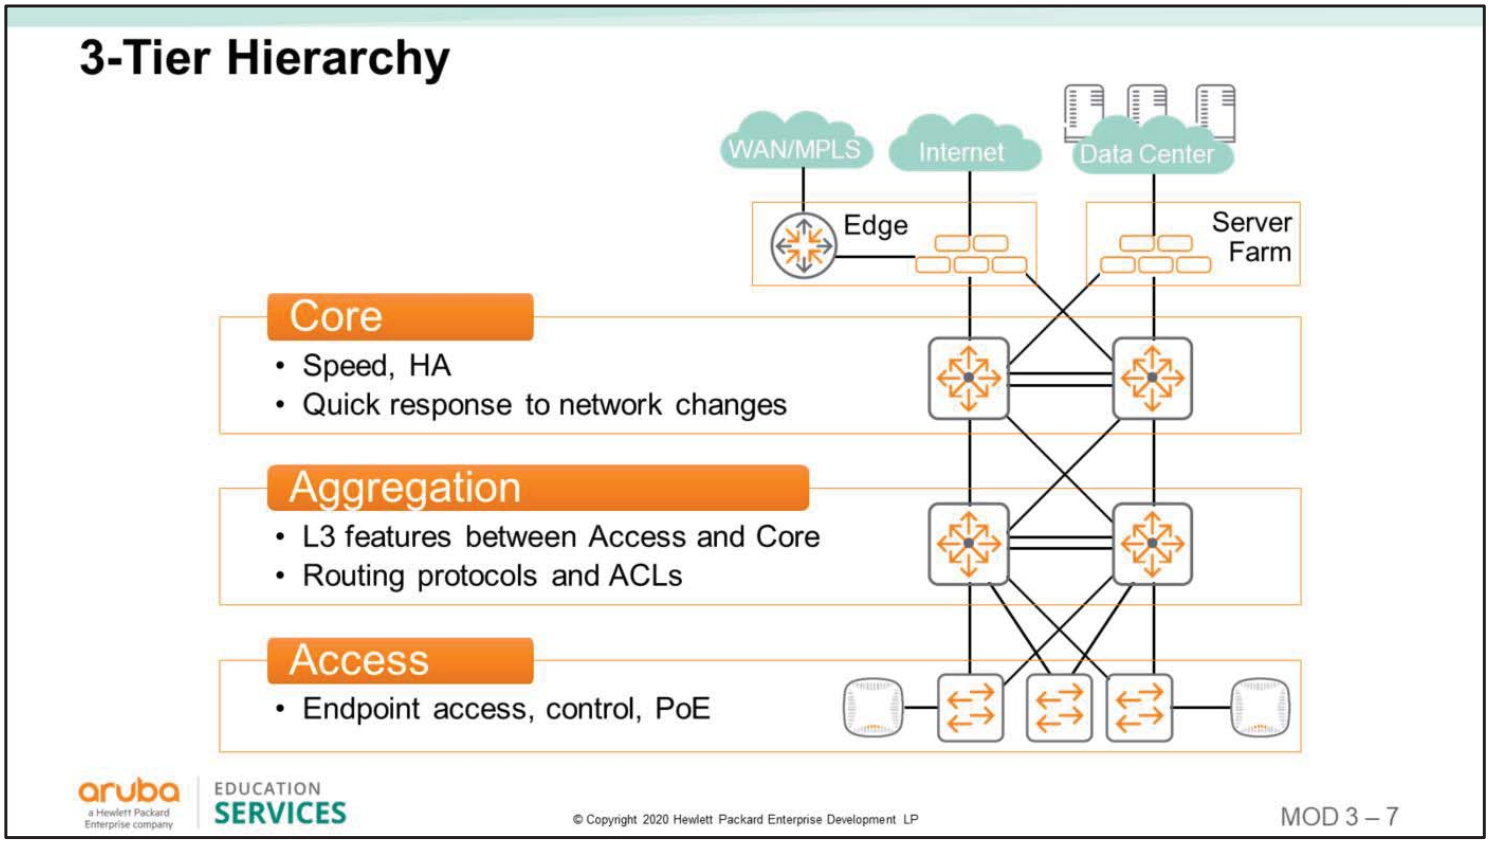

A Three-tier design divides the network in three layers - Access, Distribution and Core.

|The Core Layer is as described for a two-tier design. It must switch and route packets as fast as possible, maintain high availability, and quickly respond to changes in the network. The aggregation layer, Edge Network, and Server Farms may connect to this layer.

|The Access Layer provides for endpoint access, with PoE and Layer 2 features, as with a 2-Tier design.

|The aggregation Layer provides Layer 3 features for traffic to and from the access layer. Routing protocols and ACLs ensure that the correct traffic is routed along the best path. Essentially, you have expanded the core into two separate layers. This can increase cost, and perhaps complexity. However, it may make the network more scalable in certain larger deployments.

Question #1

Network Design

Which options below describe differences between 3-tier and 2-tier network designs?

a. In a 3-tier design, the core is must quickly respond to network changes.

b. The 2-tier Access Layer must often provide PoE for end systems.

c. The 2-tier aggregation layer is less robust than a 3-tier design.

d. A 3-tier design might be more scalable than a 2-tier design.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 8

Correct answer:
d.

Aruba Switching Portfolio

AOS-S and AOS-CX

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

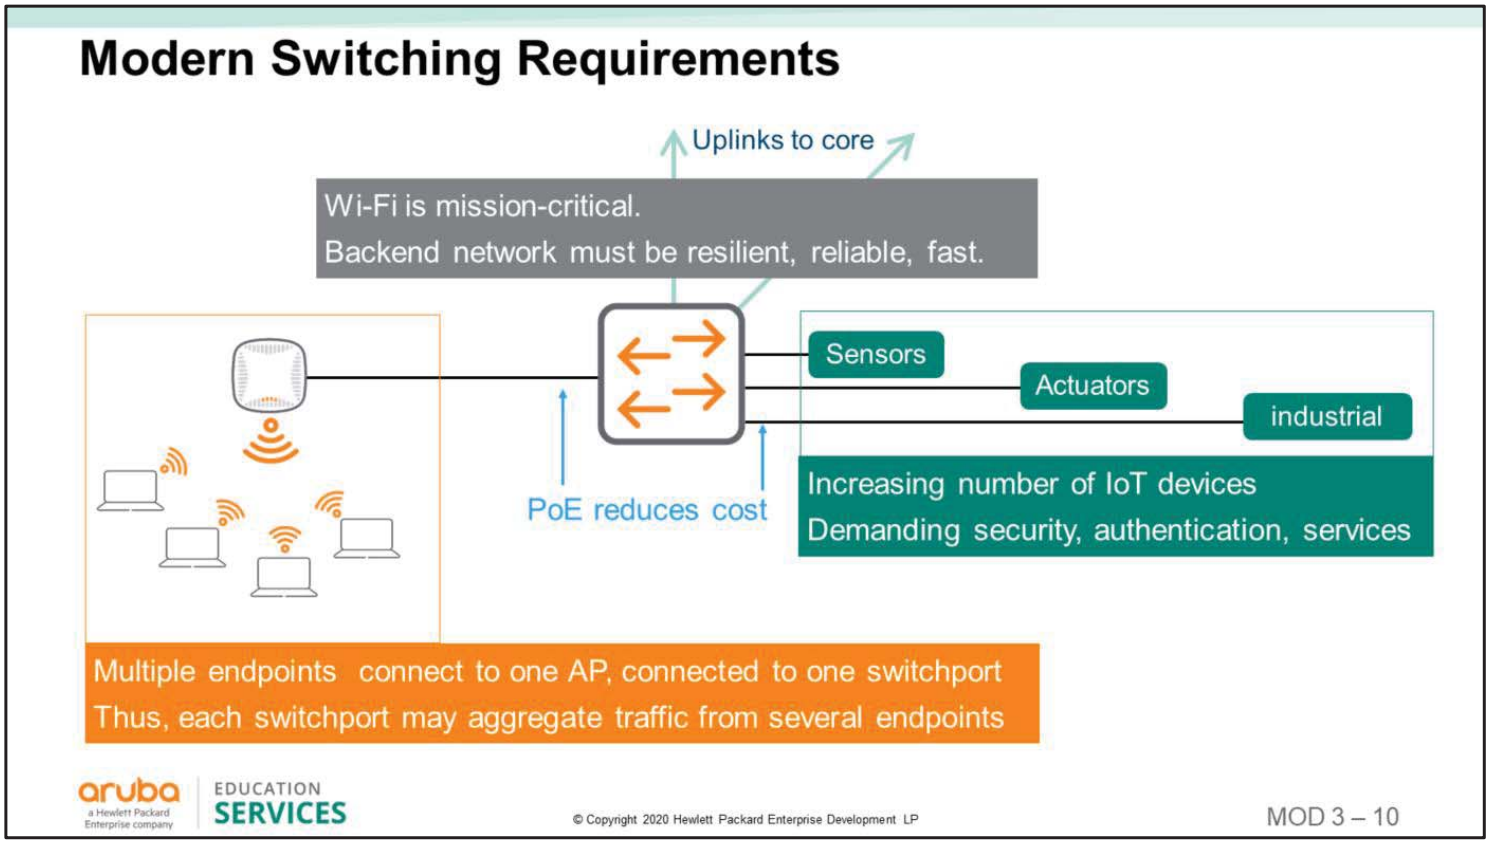

Traditionally the network switch had one job - provide wired access to computing devices. These legacy switches had low port bandwidth utilization and minimal power requirements. Each wired user connected to a separate switch port – each port handled traffic from a single user.

Today, the role of the switch has dramatically changed. Many users prefer a wireless connection.
|They connect to APs, which in turn connect to switches. Thus, the switch has become an aggregation device - each port handles traffic from multiple endpoints.

|For many organizations, good wireless communication is mission-critical. Resilient, reliable, and high-performance aggregation must ensure the best mobile user experience.

|Also, there is an increasing number of Internet-of-Things (IoT) devices – often deployed throughout the organization. To support these devices, switches must offer security, authentication, and services such as Power over Ethernet (PoE) to devices that request it.

|The Power over Ethernet (PoE) feature enables a switch to provide device power, in addition to data connectivity. This eliminates the need to run separate power cables to every device, and drastically reduces cost. However, this means that the switch must have big enough power supplies to not only power itself, but all connected PoE devices.

Aruba networks has created a switching portfolio that meet today's requirements in terms of performance, security, manageability, automation and wireless optimization. The portfolio accommodates any organization, regardless their size and complexity.

To meet the objectives of modern network deployments, Aruba Networks has developed two Network Operating Systems (NOS) - Aruba OS-S and AOS-CX.

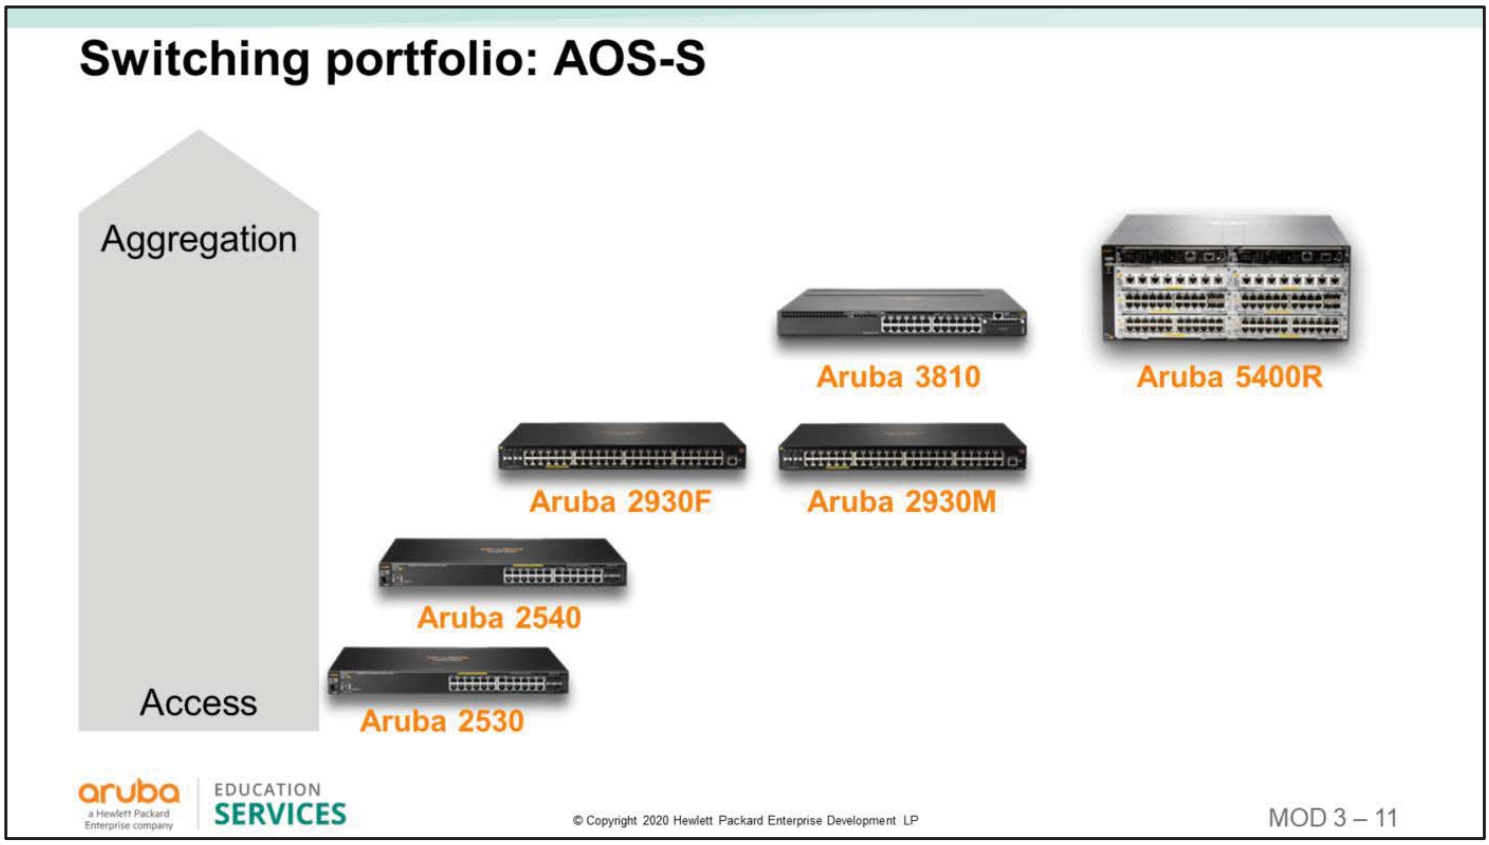

This family switches offer excellent performance in Campus LANs, offering Layer 2 and Layer 3 features.

|Aruba 2530 switches series

These are fixed switches available in 8, 24, and 48-port models - suitable for the Access layer. They have typical Layer 2 features, along with simplified Quality of Service (QoS) and security features. This series supports 10/100 Mbps only ports and Gigabit Ethernet ports, with up to 4 Small Form-factor Pluggable (SFP) uplinks. Full PoE support provides up to 370 Watts to power Access Points, IP Cameras, IoT devices, and more.

Multiple authentication methods are available in this series. This includes the very robust 802.1X authentication and the simpler, but less secure MAC authentication. Web-based authentication is available for clients that do not support 802.1X. This enhances security by controlling access to the network.

This series can be integrated into the Aruba Central cloud-based management platform. This offers a simple, secure, and cost-effective way to manage switches anywhere and anytime.
Zero-Touch-Provisioning (ZTP) simplifies installation of the switch infrastructure in remote locations.

Note. Small Form-Factor Pluggable (SFP) is a compact, hot-pluggable network interface module which gives the switch the interface port to be equipped with any suitable type of transceiver as needed. For example, some SFP interfaces support copper media, while

others may support multi-mode (MM) or Single Mode (SM) fiber optic connections.

|Aruba 2540 switches series

These are fixed switches available in 24 and 48-port models - suitable for the Access layer. This series differ from 2530 because the ports do not support 10/100mb Ethernet. The minimal port rate is Gigabit Ethernet. The Uplink ports that connect to aggregation switches support 1 or 10 Gigabit Ethernet.

Full Power over Ethernet (PoE+) is supported. This series supports the typical Layer 2 protocols plus basic Layer 3 features such as: RIP routing, ACLs, and robust QoS. Security is enhanced by authentication methods and Access Control Lists (ACL). Source-port filtering can limit certain ports to communicate only with each other. Simple deployment with Zero Touch Provisioning is available, as is the automation of network operations, monitoring, and troubleshooting. This is enabled using REST APIs.

|Aruba 2930F switches series

These fixed switches are available in 8, 24, and 48-port models - suitable to be deployed in the Access Layer. This series uses Gigabit Ethernet ports with 1 and 10 Gbps uplinks. Full PoE+ support provides up to 740 watts of power.

The 2930F series offers Layer 3 capabilities including OSPF, Dynamic segmentation, ACLs, IPv6, and robust QoS - all without requiring an additional software license.

This series also supports Virtualized switching or stacking. When switches are stacked, they appear as a single chassis, to simplify management. Up to 8 members can be stacked in a ring topology. Zero Touch Provisioning (ZTP) simplifies installation of the switching infrastructure using Aruba Activate or a DHCP-based process. Aruba Central management is also available. Security is enhanced by running user authentication methods, ACLs, STP protection, and Private Virtual LANs (VLAN).

|Aruba 2930M switches series

This series is available in 24 and 48-port models and has similar software capabilities to the Aruba 2930F switches series. The key difference is modularity. A modular switch allows you to choose the right module for the uplink ports, add a secondary power supply, use modular 10 or 40 Gigabit

Ethernet uplinks, and select models with Smart Rate ports.

HPE Smart Rate multi-gigabit Ethernet delivers faster connectivity than a regular Gigabit Ethernet port, along with PoE, using existing campus cabling. Full PoE+ is supported with the newer 802.3bt Class 6 cabling standard. This means that you can power devices that require up to 60W. A redundant power supply is available, to provide up to 1440 total watts of power.

|Aruba 3810M switches series

These fixed switches are available in 16, 24, and 48-port models, with or without full PoE+ and Smart Rate2 ports. These Smart Rate ports can deliver high speed rates and power for 802.11ax and 802.11ac devices using existing CAT5e AND CAT6 twisted pair wiring.

The 3810 series supports typical Layer 2 features, and Layer 3 features such as OSPF, BGP. They also support the Virtual Router Redundancy Protocol (VRRP), as well of QoS and security features.

The switch series supports Virtualized switching or stacking, and an available slot accepts both 10 and 40 Gigabit Ethernet modules. Stacking has been designed for high performance and provides up to 336Gbps of stacking throughput. Using a ring topology, you can stack up to 10 switches. In a mesh topology, the maximum is 5.

|Aruba 5400R ZL2 switches series

These modular switches are available in 6-slot or 12-slot chassis and can be deployed at the Access or Core of a campus network, depending on the size of the network. The 5400R ZL2 switches support the most demanding network features , including QoS and security. Redundant management modules and redundant power supplies providing high availability for environments that cannot tolerate down time.

The switch supports any combination of 10/100Mbps Ethernet and 1/10/40 Gigabit Ethernet with full PoE+ on all ports. Smart Rate2 ports are also supported.

5400 ZL2 switches can run Layer 3 features such as: OSPF, BGP, VRRP, and Protocol Independent Multicasting (PIM). They also support the Virtual Switching Framework (VSF) feature, which enables you to combine two switches into one virtual switch, called a fabric.

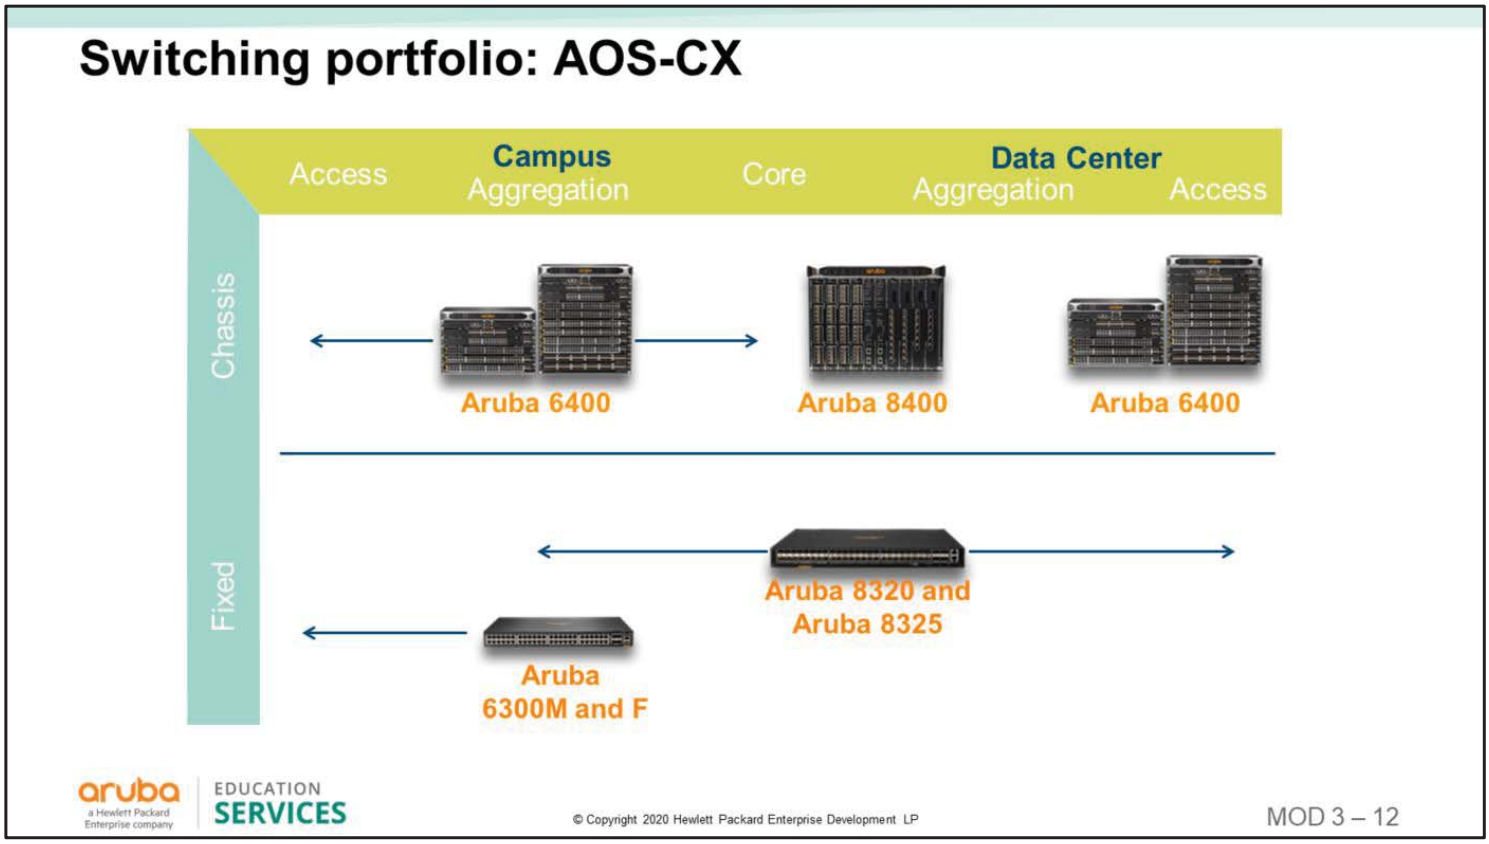

AOS-CX is a modern software system for the enterprise core and aggregation that automates and simplifies many critical and complex network tasks. It delivers enhanced fault tolerance and facilitates zero-service disruption during planned or unplanned control-plane events. The key innovations in AOS-CX are its micro-services style modular architecture, REST APIs, Python scripting capabilities and the Aruba Network Analytics Engine (NAE). AOS-CX is based on modular architecture that allows individual processes to re-start and be upgraded. Its REST APIs and Python scripting enables fine-grained programmability of the switch functions and its unique Aruba Network Analytics Engine (NAE) provides for easy network monitoring and troubleshooting.

A Data Center is a facility that centralizes an organization's IT operations and equipment, as well as where it stores, manages and disseminates its data. Data centers house a network's most critical systems and are vital to the continuity of daily operations. The security and reliability of data centers and their information is top priority for organizations.

A Data Center architecture could include a three-tier layer approach. Access, Aggregation and Core Layers help to create a resilient network. AOS-CX switches were design to work in the Campus LAN network and as well in the Data Center, offering a rich set of features and high performance.

Aruba 8400 switches series

The Aruba 8400 switch is suitable for deployments as a Campus Core switch or as a Core/Aggregation switch in the Data Center. This series switch has been designed to provide high availability and resiliency in every part of the hardware, supporting 99.999%

availability.

These modular switches are available in 8-slot chassis for interface modules or Line Cards (LC). Options for Line Cards include 32-port 10GbE modules, 8-port 40GbE and 6-port 100GbE modules. Fabric modules provide the ability for traffic to flow between the line cards. The 8400 switch supports up three fabric modules which provide redundancy and keep line rate speeds in case of a failure. The switch offers up to 19.2 Terabit per second switching capacity.

The 8400 series switch supports redundant Management Modules (MM), fan assemblies, and power supplies. All modules are hot-swappable, permitting upgrade or replacement without powering off the chassis. These switches also include the Aruba Network Analytics Engine (NAE) - a framework for monitoring, troubleshooting, and capacity planning.

Advanced Layer 2 and Layer 3 features include BGP, OSPF, VRF, Multicasting, and IPv6. Also, this switch supports VSX and Multi-chassis Link Aggregation (MLAG) which enables you to create one virtualized switch composed by two individual physical switches.

Aruba 6400 switches series

The Aruba 6400 switch series is a modern, flexible and intelligent family of modular switches ideal for access, aggregation, and core in enterprise Campus LAN and Data Center deployments. The switches provide the foundation for high-performance networks supporting IoT, mobile, and cloud applications.

This switch series comes with high quality next generation Aruba Gen7 ASICs to support multi-Gigabit requirements, automation, security, high availability, and PoE - all running on a next generation modular AOS-CX operating system.

This switch series introduce the models 6405 - a 5-slot chassis for Line Cards, and the 6410 switch - a 10-slot chassis. Both chassis models support up to two management modules, four modular AC power supplies, and two or four fan trays (6405/6410).

Depending on the line card used the 6400 series supports a variety of interfaces: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, and 100GbE. This switch offers 24 Terabit per second switching capacity.

These switches support BGP, EVPN, VXLAN, VRF and OSFP with robust

security and QoS. High availability is accomplished with VSX redundancy and redundant power supplies and fans.

Aruba 8325 switches series

The Aruba 8325 switch series offers a flexible and innovative approach to addressing the application, security and scalability demands of the mobile, cloud and IoT era. These switches serve the needs of the core and aggregation layers, as well as Top of Racks (ToR) and End of Row (EoR) Data Center requirements.

This switch series provide over 6.4 Terabit per second switching capacity with line-rate Gigabit Ethernet interfaces including 1Gbps, 10Gbps, 25Gbps, 40Gbps and 100Gbps.

This 1 Rack Unit switch supports advanced Layer 2, 3 features that includes BGP, OSPF, VRF-Lite and IPv6. Dynamic VXLAN with BGP-EVPN for deep segmentation in Data Center and Campus networks.

This switch series offer two models with 48 and 32 ports, both models include 6 fans and 2 power supplies. Also, each switch model offers the choice to select the airflow front-to-back and back-to-front.

Aruba 8320 switches series

The Aruba 8320 switches are powerful 1 Rack Unit (RU) devices designed to provide core and aggregation services in mid-sized Campus networks. These switches offer three fixed-port models with 10GbE and 40GbE interfaces offering up to 2.5 Terabits per second switching capacity.

This model offers similar features as the 8325 switches series including BGP, OSPF, VRF, IPv6, VSX and Multi-chassis link aggregation.

The 8320 switches are also built for resiliency. They offer redundant power supplies and fans.

Aruba 6300 switches series

The Aruba CX 6300 switch series is a modern, flexible and intelligent family of stackable switches ideal for enterprise access and aggregation layers. This switch series is built around the new Aruba Gen7 ASIC, a highly capable CPU

and the next generation modular AOS-CX switch software platform.

The Aruba Virtual Stacking Framework (VSF) allows for stacking of up to 10 switches, providing scale and simplified management. This flexible series has built-in 1GbE, 10GbE, 25GbE and 50GbE uplinks and support high density High-Power PoE interfaces, offering up to 880 Gbps system switching capacity.

This series support one touch deployment with the Aruba CX Mobile App. Aruba Dynamic Segmentation extends Aruba's foundational wireless role-based policy capability to Aruba wired switches. What this means is that the same security, user experience and simplified IT management can be enjoyed throughout the network. Regardless of how users and IoT devices connect, consistent policies are enforced across wired and wireless networks, keeping traffic secure and separate.

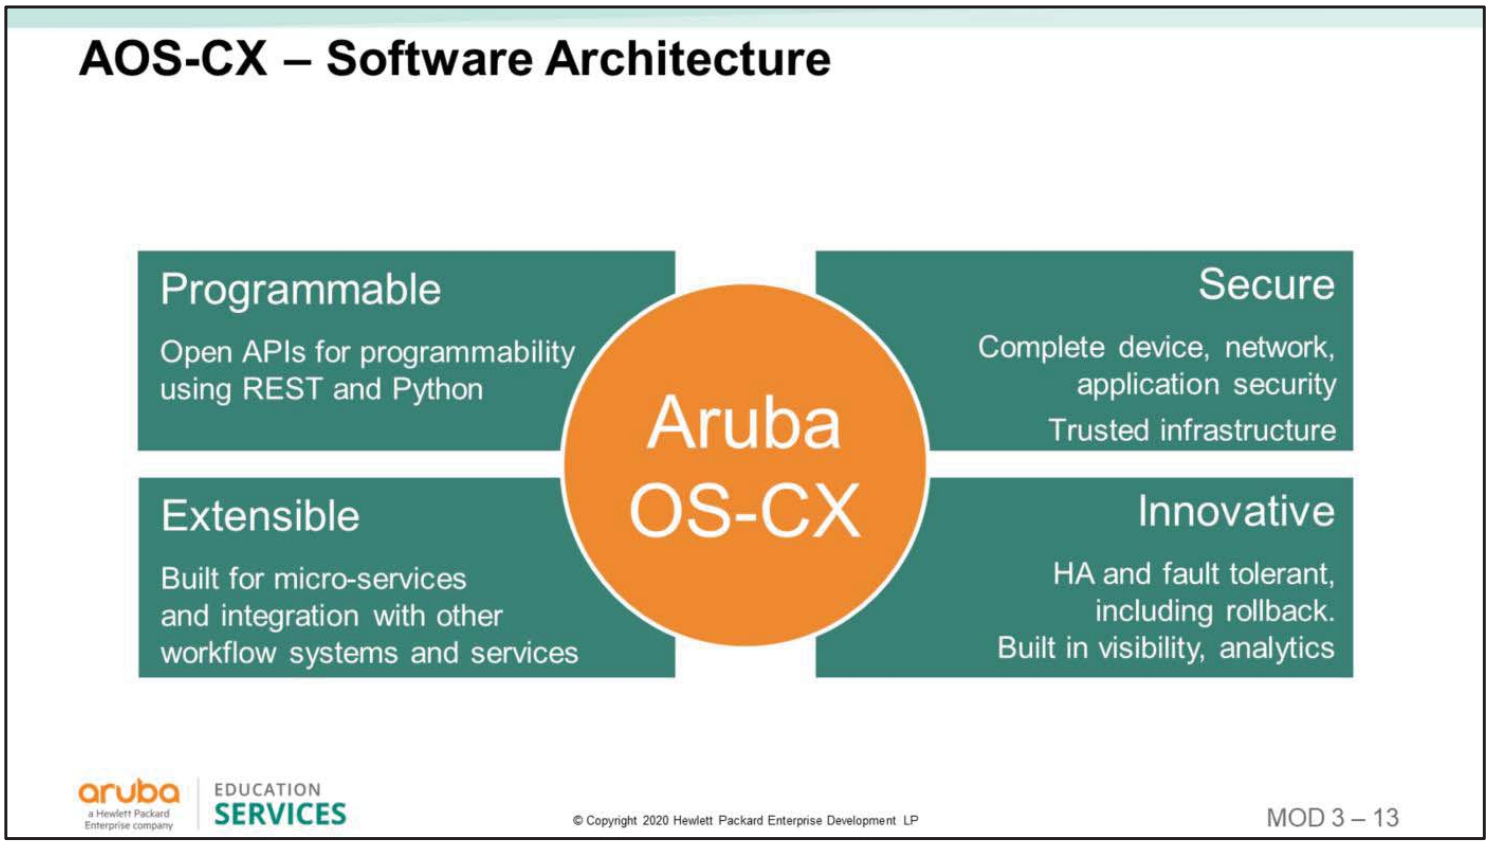

|The AOS-CX software is a modern, database-driven operating system that automates and simplifies many critical and complex network tasks. A built-in time series database enables customers and developers to use software scripts for historical troubleshooting and to analyze past trends. This helps to predict and avoid future problems due to scale, security, and performance.

This network operating system is built on a modular Linux architecture with a stateful database which helps to offer the following unique capabilities:

Easy access to all network state information allowing unique visibility and analytics.
REST APIs and Python scripting for fine-grained programmability of network tasks.
A micro-services architecture that enables full integration with other workflow systems and services.
Continual state synchronization that provides superior fault tolerance and high availability.
Security best practices were applied to completely create a trusted platform.

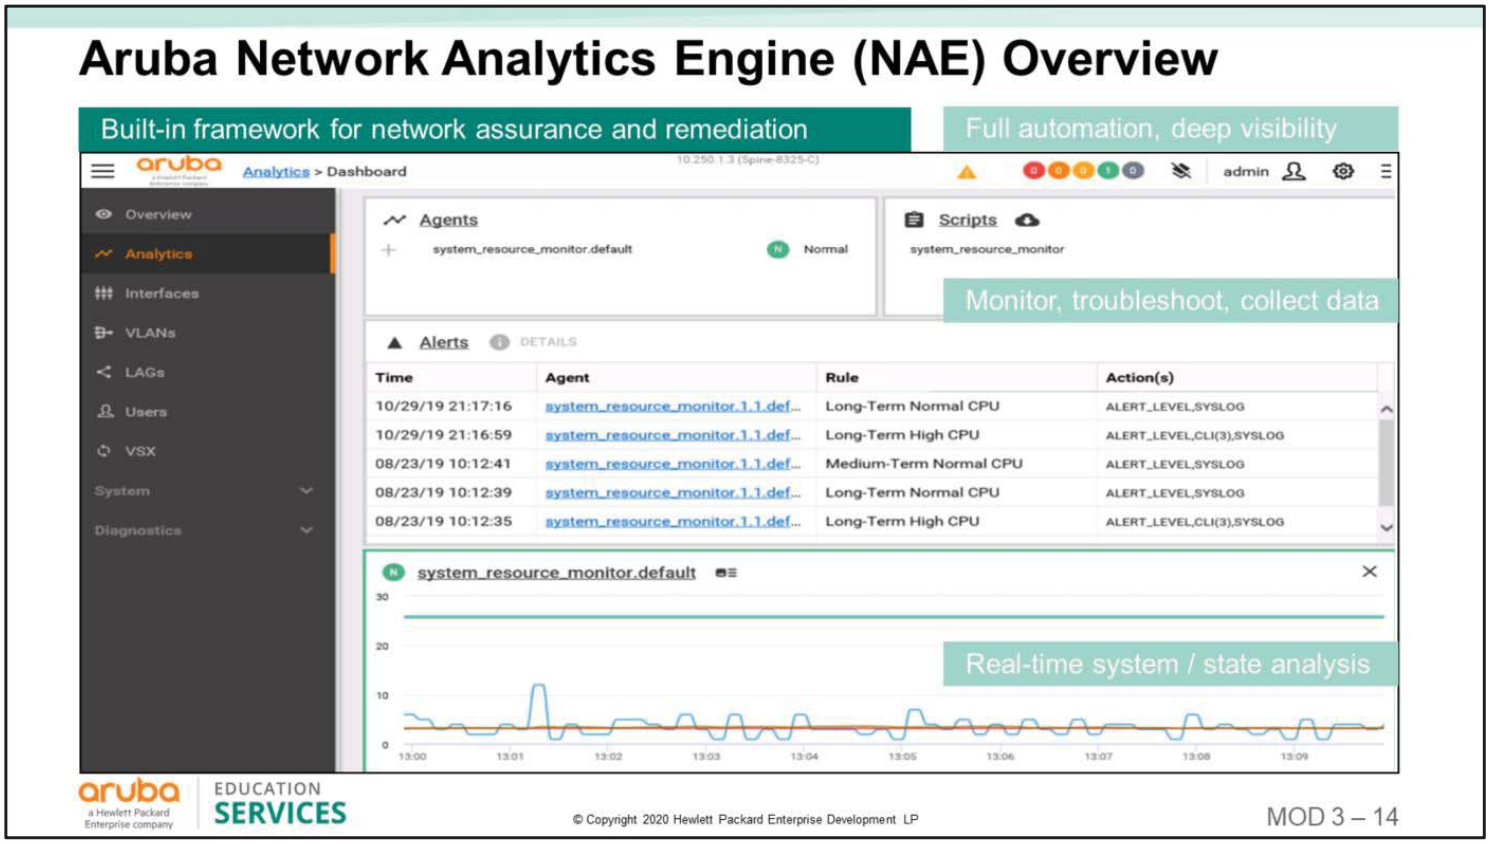

|The Aruba Network Analytics Engine (NAE) is a built-in framework for network assurance and remediation. Combining the full automation and deep visibility capabilities of the AOS-CX, this framework allows monitoring, troubleshooting and easy network data collection using simple scripting agents.

This framework analyzes a problem in real time giving you the insight you need to resolve the issue or take corrective action based on established policies. When an anomaly is detected, it can proactively collect additional statistics and data.

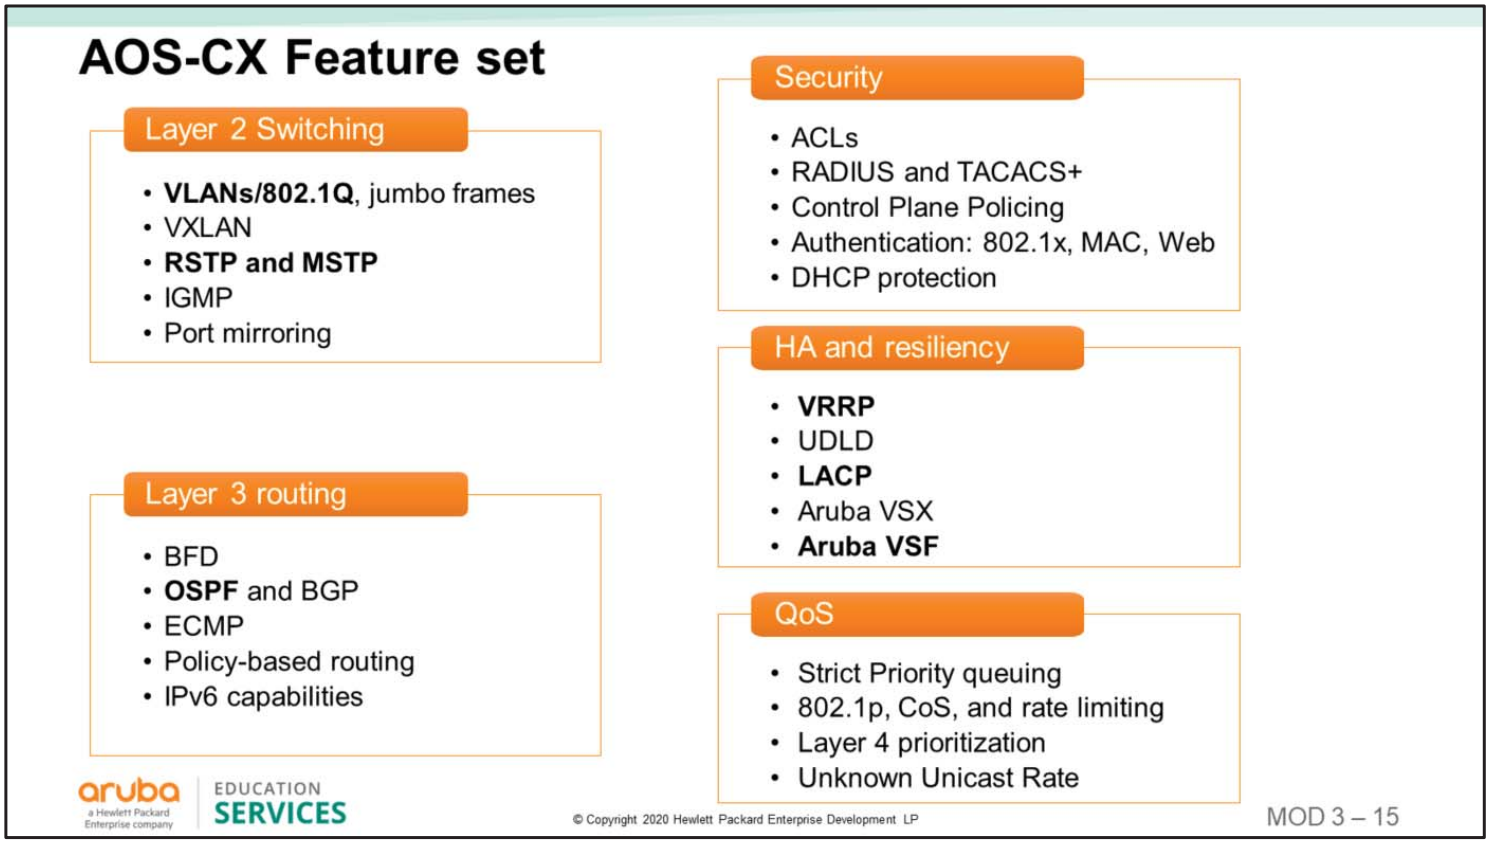

The AOS-CX has the following feature set:

|Layer 2 switching
VLAN support and tagging for 802.1Q.
Jumbo packet support.
VXLAN encapsulation protocol.
Rapid Spanning Tree Protocol (RSTP).
Rapid Per-VLAN Spanning Tree (RPVST+).
Multiple Spanning Tree Protocol (MSTP).
Internet Group Management Protocol (IGMP).
Port mirroring

|Layer 3 routing
Bidirectional Forwarding Detection (BFD).
Border Gateway Protocol (BGP).
Equal-Cost Multipath (ECMP).
Multi-protocol BGP.
Open shortest path first (OSPF).
Static routing.
Policy-based routing.
IP performance optimization.
IPv6 capabilities.
Protocol Independent Multicast (PIM).

|Security

Access Control Lists (ACL) for both IPv4 and IPv6.
RADIUS.
TACACS+.
Control plane policing.
Authentication based on 802.1X, MAC and web based.
DHCP Protection.
Secure encryption for all access methods including: SSH, SSL and SVMPv3.
Switch CPU protection.
Identity-driven ACL.

|High Availability (HA) and resiliency
Virtual Router Redundancy Protocol (VRRP).
Uni-direction Link Detection (UDLD).
Link Aggregation Control Protocol (LACP).
Aruba Virtual Switching Extension (VSX).

Quality of Service (QoS)
Strict Priority queuing.
Deficit Weighted Round Robin (DWRR).
Traffic prioritization (802.1p).
Layer 4 prioritization based on TCP/UDP port numbers.
Class of Service.
Rate limiting.
Unknown Unicast Rate.
Large buffers for graceful congestion management.

Note. Some of these features are to be discussed in this course (those are bold in the slide)

Note. The features listed do not apply to all switch models, please refer to the specific data sheet for each model to verify if the feature is supported.

Note. This text considers the feature set for AOS-CX 10.4 release

Question #2

Switch Platforms

What are some advantages of a modular, chassis-based switch?

a. They are more scalable.
b. They tend to cost less.
c. They use less power.
d. They include optional PoE capabilities.
e. They are more flexible.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 16

Correct answers:
a.
e.

AOS-CX

CLI Access

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

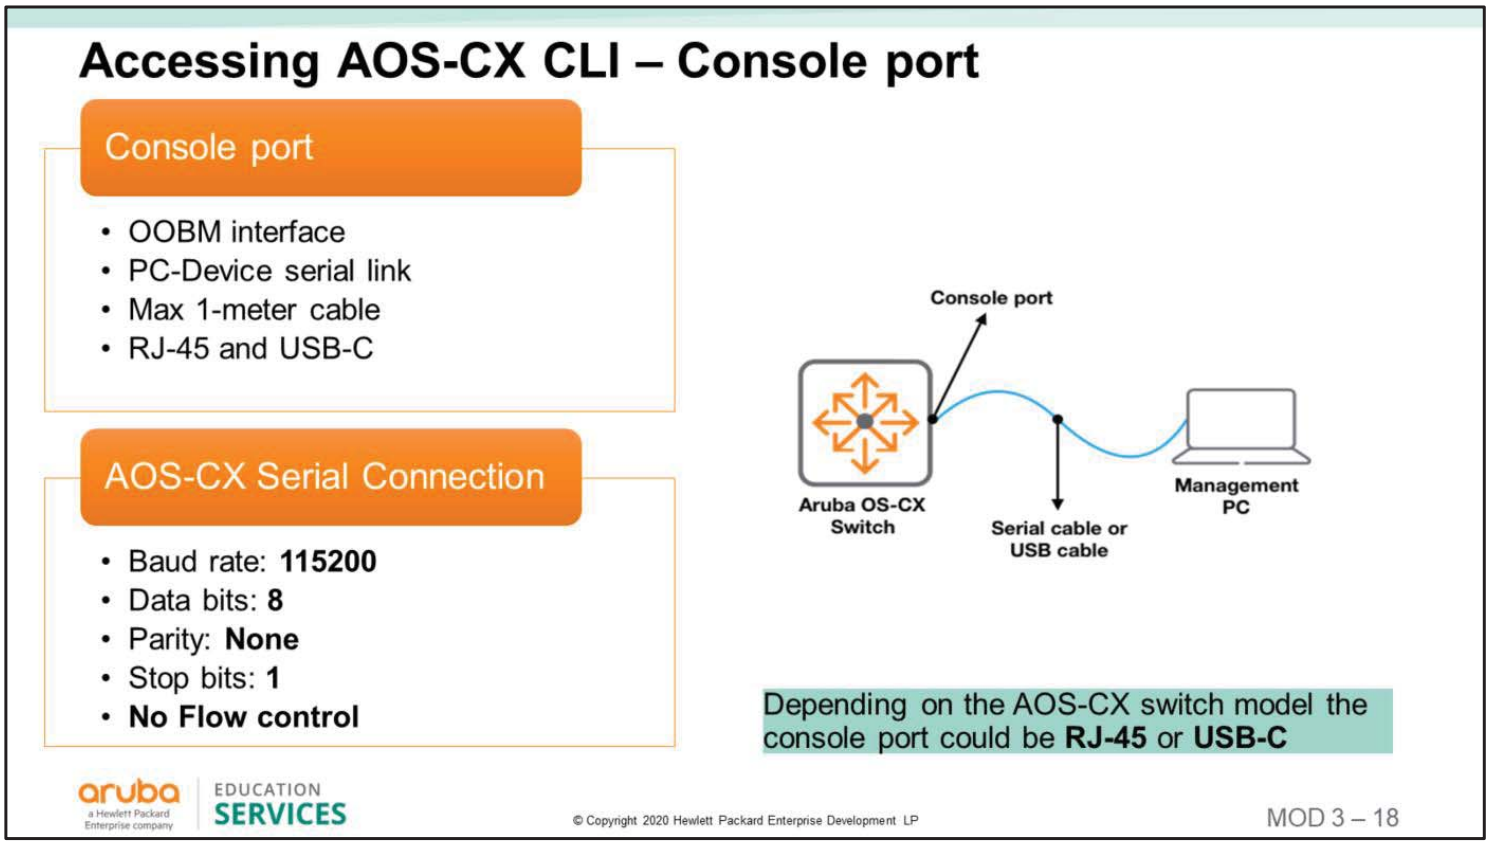

Console port

When you initially configure a switch, you will typically use out-of-band Management (OOBM) on a console port. This special console port is integrated on all switch models and types, to facilitate configuration, troubleshooting, and management.

Depending on the switch model, the console can be a USB-C port or an RJ-45 port, however for both port types you must establish a serial connection between the management station (perhaps your laptop) and the switch. To do this, you need the following items:

If the switch has an RJ-45 port:
Serial cable (shipped with the switch).
A USB-to-Serial converter - modern laptops no longer include a serial port.
|Software on your laptop or management station that emulates the serial session. Some common examples include PuTTY, TeraTerm, and Secure CRT.

If the switch has a USB-C port:
USB-C cable.
Software on your laptop or management station that emulates the serial session. Some common examples include PuTTY, TeraTerm, and Secure CRT.

|Also, the serial session must be configured with the following parameters:

Baud rate: 115200

Data bits: 8
Parity: None
Stop bits: 1
No flow control

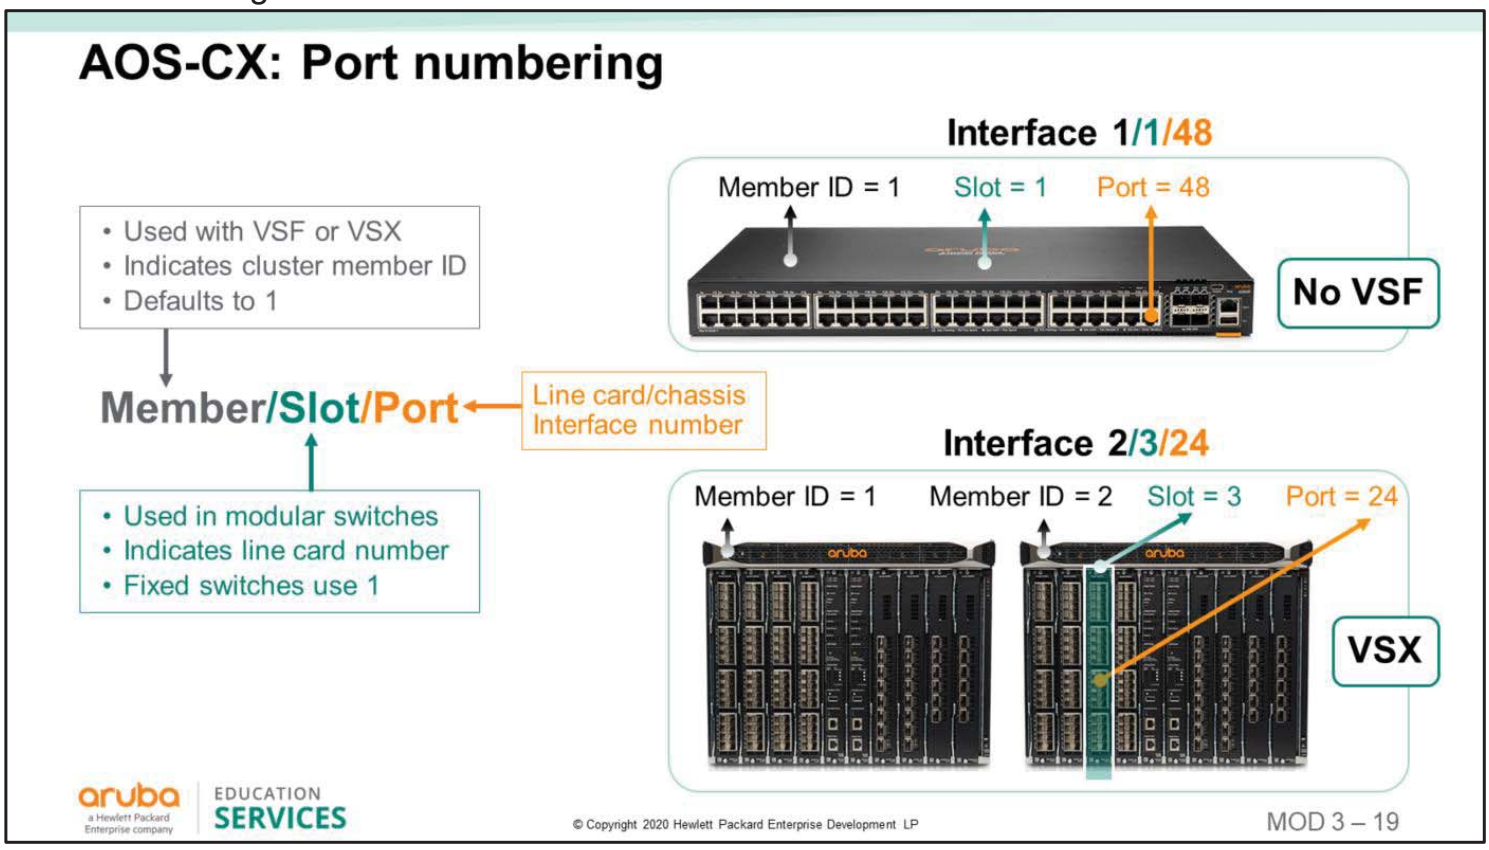

AOS-CX reference interfaces following the member/slot/port notation.

|Member. When Aruba Virtual Switching Framework (VSF) or Virtual Switching Extension (VSX) features are running multiple switches can be seen as one virtual, this number indicates the member ID in the cluster. By default, an AOS-CX switch will be member 1 indicating that these features are not running.
|Slot. In modular switches like 8400 and 6400, this number represents the slot being used by a particular line card. For fixed switches (8320,8325 and 6300) this number will be always 1.
|Port. This number refers to the individual interface in the line card (modular switches) or in the chassis (fixed switches).

|The image shows an example with a fixed switch and with a Modular switch with VSX enabled.

Note. VSF is available for 6300 switch series, VSX is supported on the other platforms. You will know more about this technology later in this course.

AOS-CX: Prompt modes

Prompt	Context	Description
`Switch>`	Operator	View commands, no configuration commands
`Switch#`	Manager	Use commands that need not be saved
`Switch(config)#`	Global config	Change switch configuration
`Switch(config-if)#`	Interface config	Descendants of global config. Use commands that apply to a specific configuration or protocol
`Switch(config-router)#`	Router config	
`Switch(config-vlan-100)#`	VLAN 100 config	

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 20

AOS-CX is organized into different configuration contexts or levels. Each context determines which parts of the switch can be managed and which commands are available to users with the appropriate authority. The operating system defines the following contexts:

Operator context (>)

This operator context enables you to execute commands to view – but not change – configuration. This context requires the least user privilege to execute commands. When in operator context, the CLI prompt is the switch name, followed by a greater than sign (>), as shown in the figure.

|Manager context (#)

This is where you can execute commands that do not require saving changes to the configuration. In command descriptions, this context is listed as the name of the switch followed by a hashtag.

To navigating to the manager context, start at the operator context (>). Then enter the enable command, as shown in the figure. You most have administrator authority to enter the enable command.

```
Switch> enable
Switch#
```

|Global configuration context (config)

This is where you can execute commands that change the switch configuration, as shown. To access this mode, start in Manager mode, then inter the command configure terminal, or shorten it to just config, as shown in the figure. To move back up one level from this or any other context, issue the exit command.

|Other configuration command context

All other configuration command contexts are descendants of the global configuration command context (config). From these command contexts you can execute commands that apply to specific configuration or protocol such as an interface or a VLAN.

Examples:

```
Switch(config-if)#
Switch(config-router)#
Switch(config-vlan-100)#
```

To return back to Manager context from any child or descendent context, enter the end command. For example:

```
Switch(config-vlan-100)# end
Switch#
```

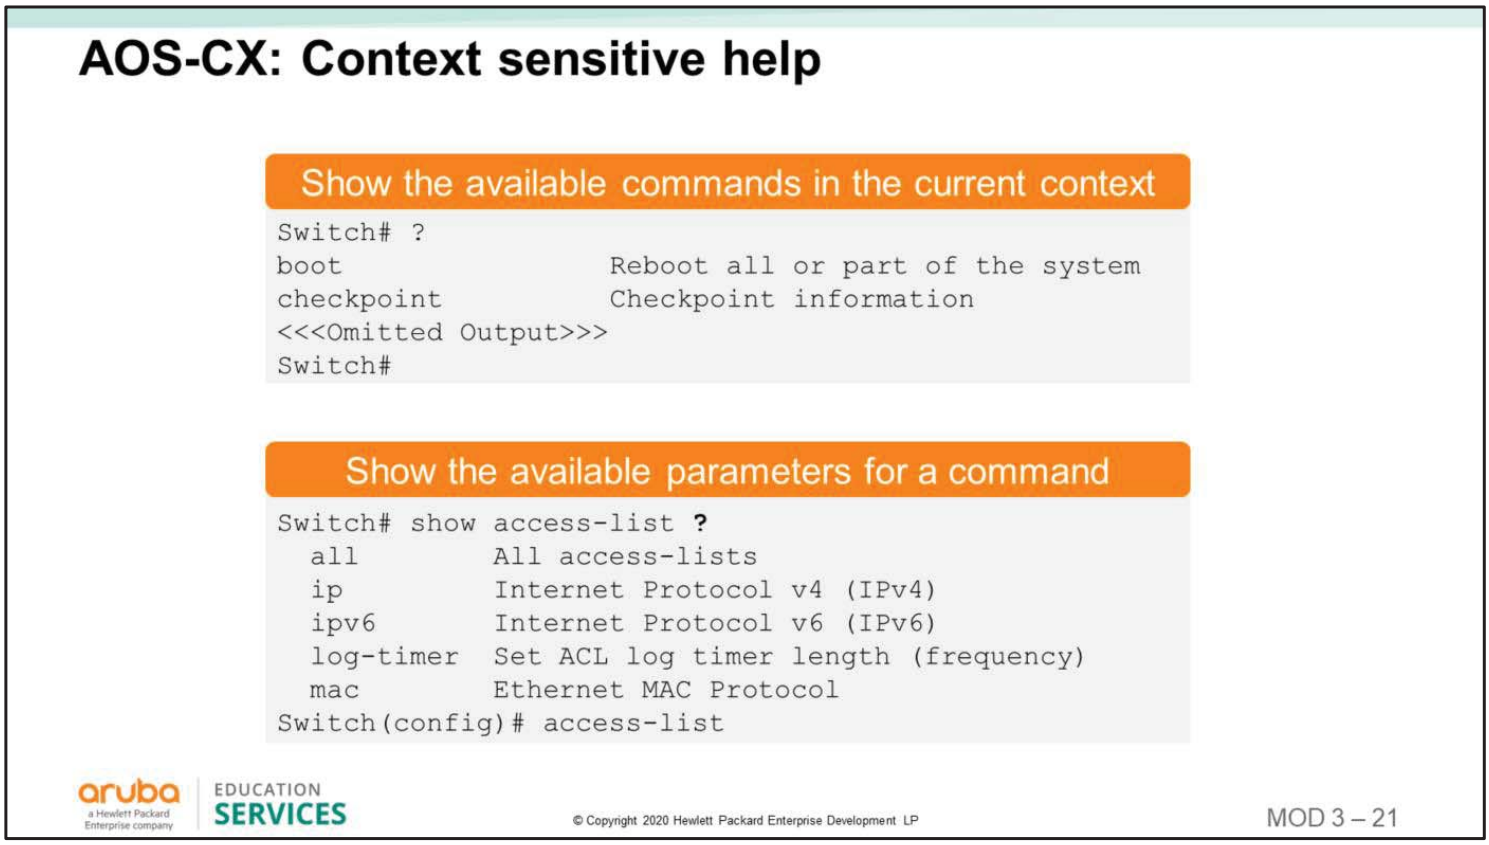

The AOS-CX CLI provides you with built-in help features. For example, to show the available commands that you can execute in the current command context enter the question mark (?) symbol. This is shown in the top example in the figure. The question mark (?) does not display on the screen when you enter it.

|To show the available parameters for a command, enter the command followed by a space and then enter the question mark symbol (?). This is shown in the bottom example in the figure. Please noticed that after the CLI displays the information, it automatically displays the text you entered before without including the help symbol (?).

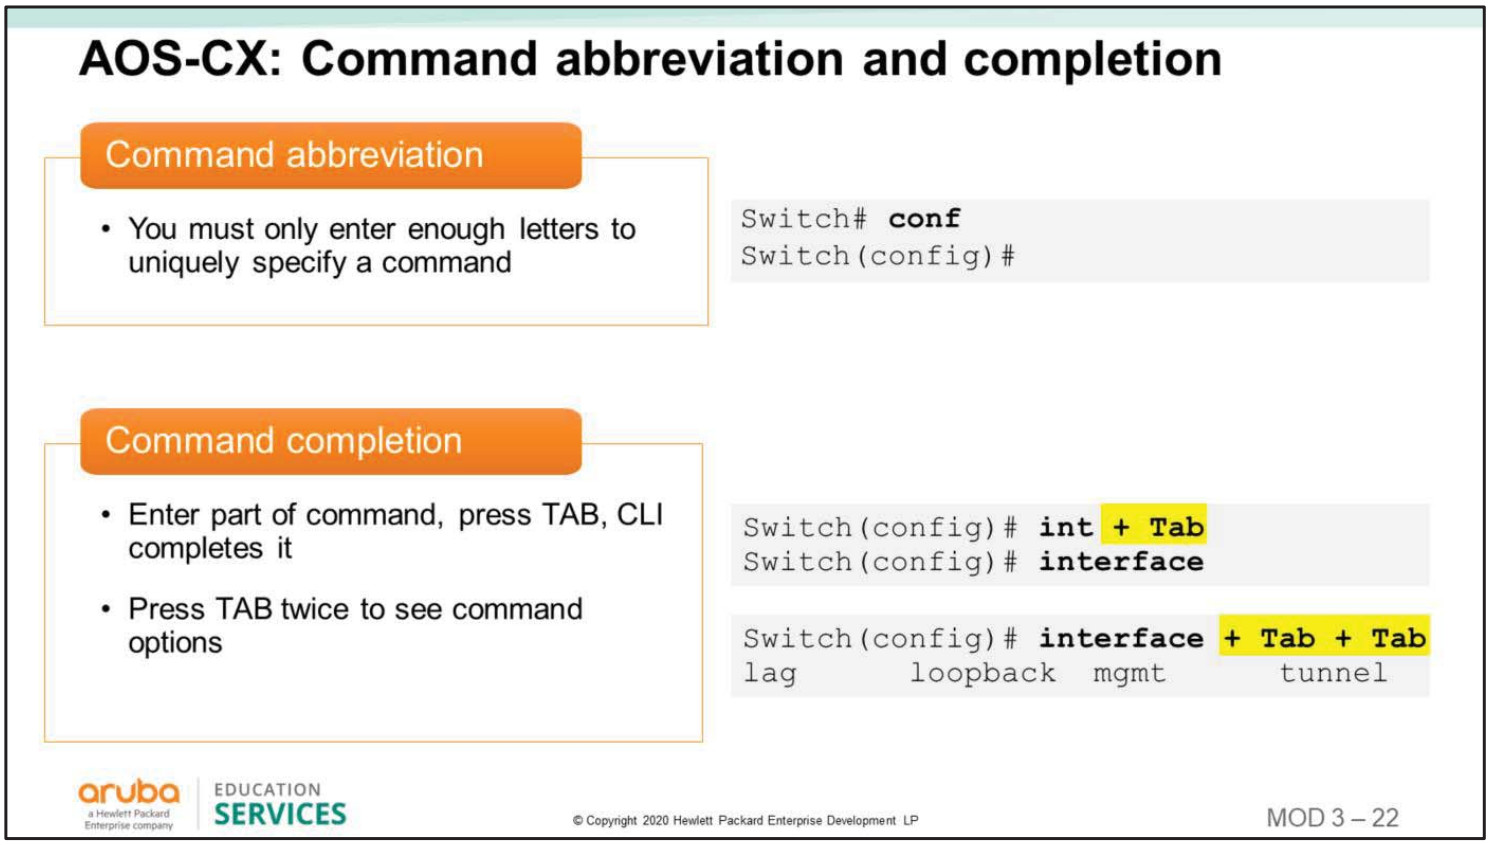

AOS-CX supports both command abbreviation and command completion. To save time by only typing in abbreviated version of the full syntax, enter enough letters to uniquely specify a valid command, and the CLI accepts the command. For example, you can enter con instead of configure to navigate from the manager context to the global configuration context.

|Command completion means that if you enter part of command word and then press the Tab key, one of the following occurs:
If you enter enough letters to match a valid command, the CLI displays the reminder of the word.
If you have not entered enough letters to match a valid command, the CLI does not complete the command.

Example,

Switch(config)# int + Tab
Switch(config)# interface

If you press the Tab key twice after a completed word, the CLI displays the command options. For example.

Switch(config)# interface + Tab + Tab
lag loopback mgmt tunnel vlan vxlan

Question #3

Console Port

What kind of cables might you use to connect to an AOS-CX Switch console port?

a. Ethernet cable

b. Fiber optic cable

c. USB cable

d. Serial cable

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 23

Correct answers:
c.
d.

Getting switch information

Show commands

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

AOS-CX – The value of show commands

Discovery and documentation

- Validate existing documentation
- Create new documentation

Baselining and troubleshooting

- Understand what "normal" is for your network
- The combination of good documentation and good baseline information will greatly ease troubleshooting during an outage.

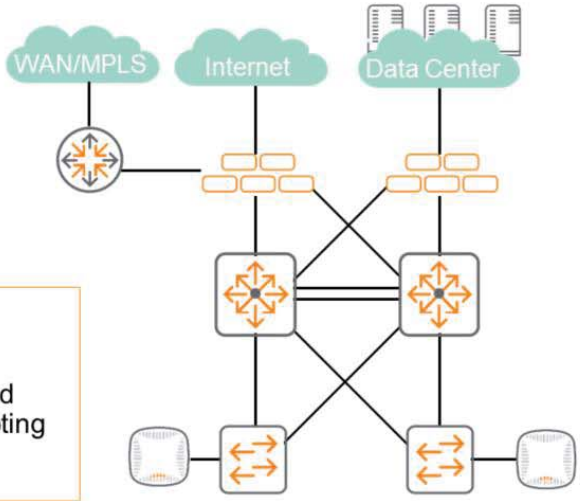

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 25

In this section you will learn about basic commands that will help you to verify your networking equipment, along with its general capabilities. This is some of the most valuable information you can learn during an entry-level networking course. Good network engineers and administrators must be adept at discovering the ACTUAL connectivity of the network. This will help you to verify the accuracy of existing documentation, or to create new network documentation from scratch.

|These show commands are also valuable for baselining and troubleshooting. Using these commands you can view many switch performance and health parameters, including CPU and memory utilization. Suppose that your network is healthy, everyone is happy, and all your switches are generally at about 5% CPU utilization. You have been paying attention to this for months now – you know your network.

During an outage or slowdown, you find a switch that is running at 85% utilization. You might focus your diagnostic efforts on that switch and its directly attached devices. Because you paid attention while the network is healthy, you are more efficient when the network is not healthy. If you had not paid attention, you would have no idea whether 85% utilization is normal or abnormal. Of course, Aruba has very effective management platforms that ease and automate these baseline processes, but it can still be valuable to use the CLI to check network device health and status.

Mastering the commands that you are about to learn, creating good documentation, and having good baseline information is kind of like getting the answers to a test before you take the test. Only this is the real world, and the "test" relates to how effective you are when the network is down.

AOS-CX – Show version

Version info for network OS, service OS, and BIOS

```
Switch# show version
------------------------------------------------------------------
AOS-CX
(c)Copyright 2017-2019 Hewlett Packard Enterprise Development LP
------------------------------------------------------------------
Version      : GL.10.04.0003
<output deleted for clarity>

Service OS Version  : GL.01.05.0002
BIOS Version        : GL-01-0013
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 26

Show version displays version information about the network operating system software, service operating system software and BIOS. In the figure, you see that this system is running the AOS-CX, version GL.10.04.0003. Near the bottom you see the Service OS and BIOS versions. This is useful for your general knowledge, and to help to determine if upgrades are needed.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators.

AOS-CX – Show system

General system status info

```
Switch# show system
Hostname             : Switch
System Description   : GL.10.04.0003
System Contact       :
System Location      :

Vendor               : Aruba
Product Name         : JL635A Aruba 8325-48Y8C 48p 25G 8p 100G swch
Chassis Serial Nbr   : TW98KM003P
Base MAC Address     : 9020c2-bced00
AOS-CX Version       : GL.10.04.0003

Time Zone            : UTC
Up Time              : 1 week, 1 day, 11 hours, 20 minutes
CPU Util (%)         : 4
Memory Usage (%)     :19
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 27

Show system displays general status information about the system. During remote access, the device platform and version information may not be obvious. In the example shown, you are attached to device named “Switch” – An Aruba model 8325-48YC. It has 48 25gbps ports, and 8 100Gbps ports. You see the serial number, base MAC address, and uptime, as well as CPU and memory utilization. This can give you some idea of the general health of the system. By baselining these parameters, you can begin to establish what “normal” is for your network.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

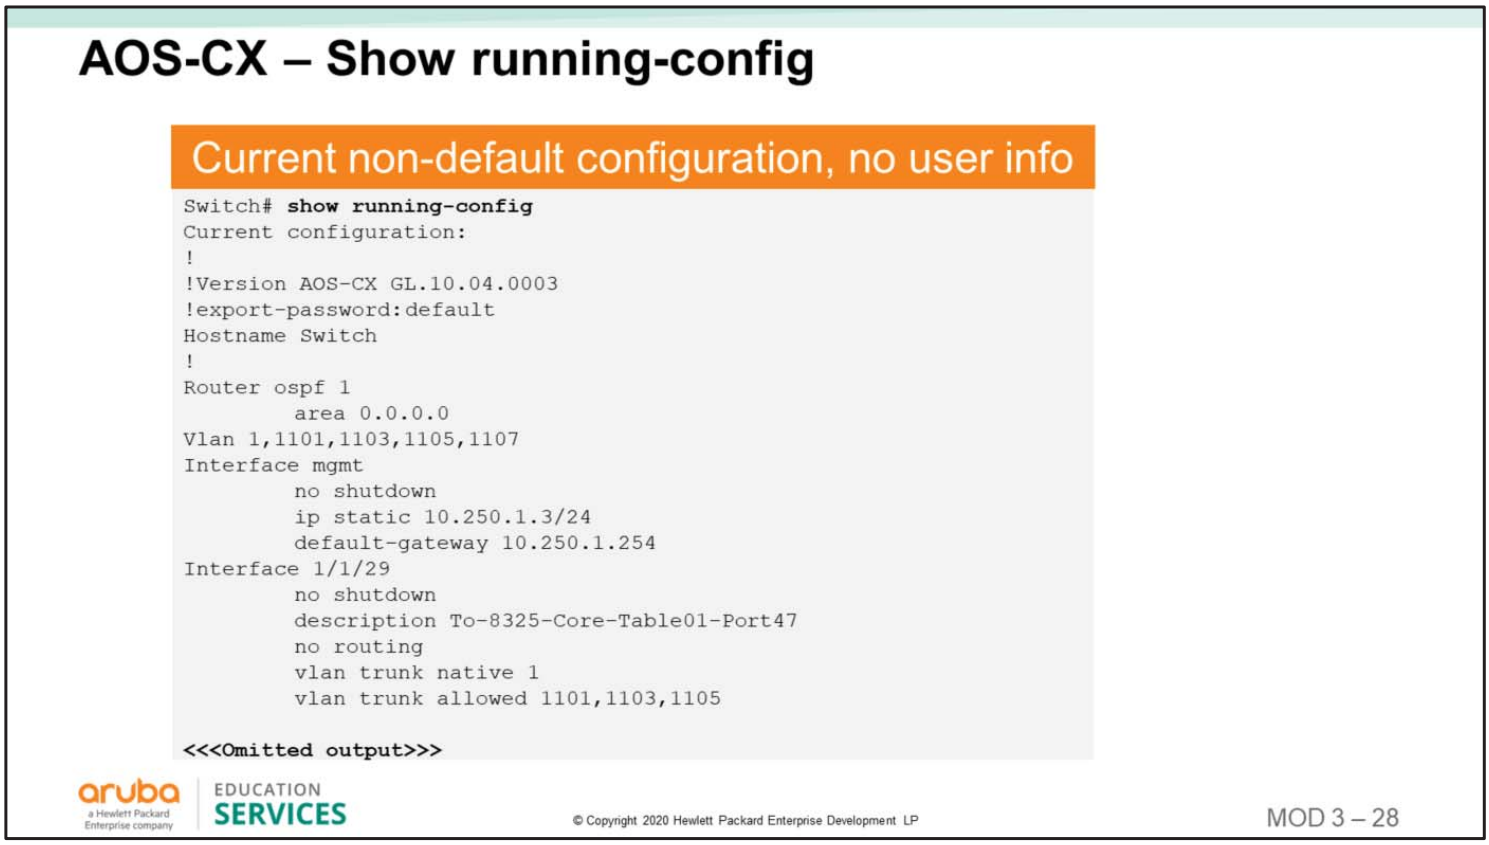

Show running-config displays the current nondefault configuration running on the switch. No user information is displayed.

Command context: Manager (#)
Authority: Administrators

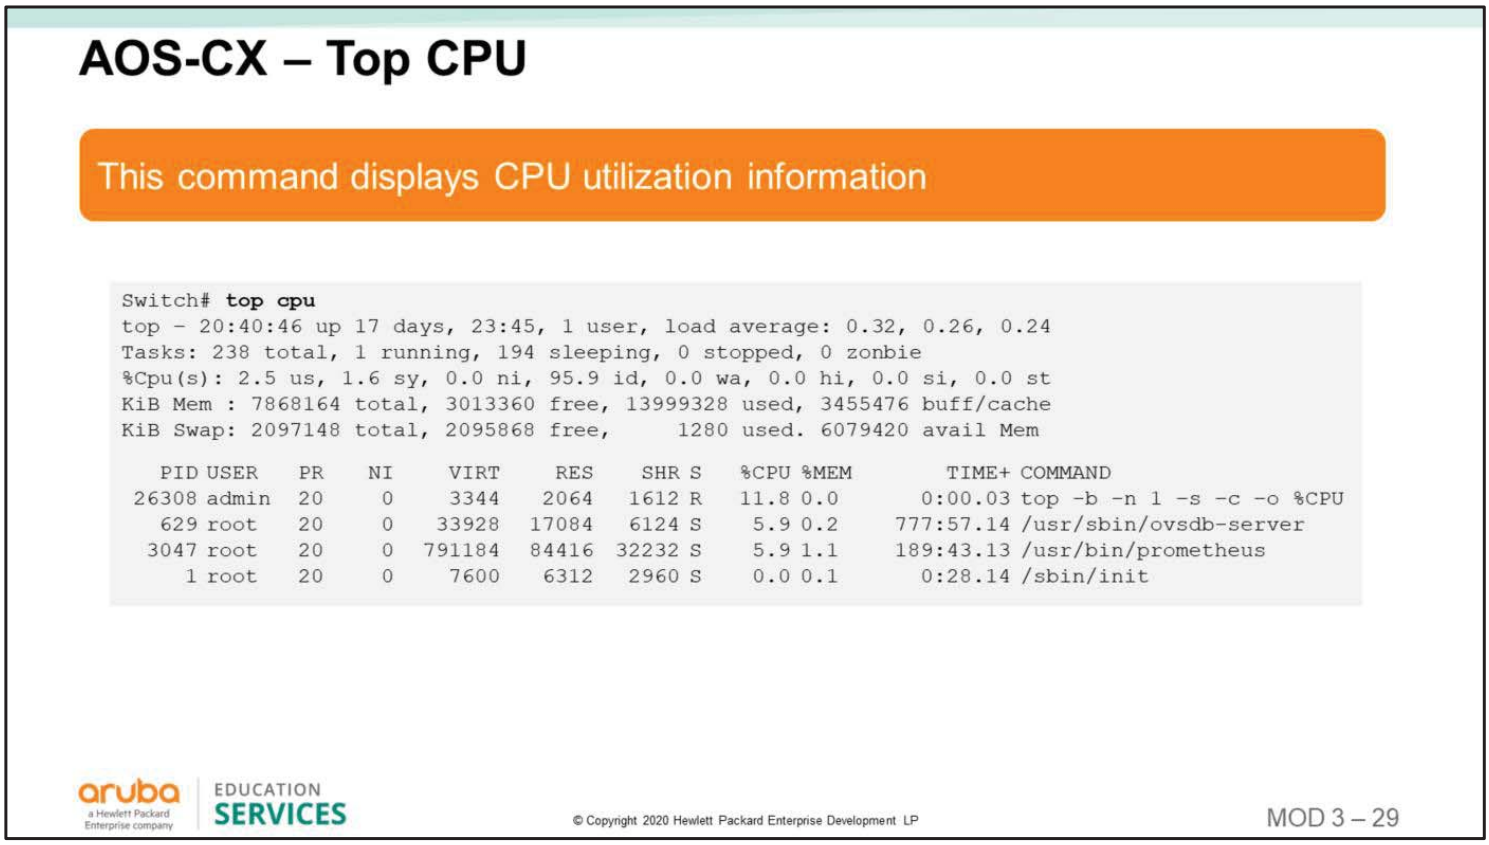

Top cpu. This command shows detailed CPU utilization information. You learned how to see high-level CPU utilization with the command show system. The top cpu command provides more detailed information.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

Top memory. This command shows memory utilization information.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

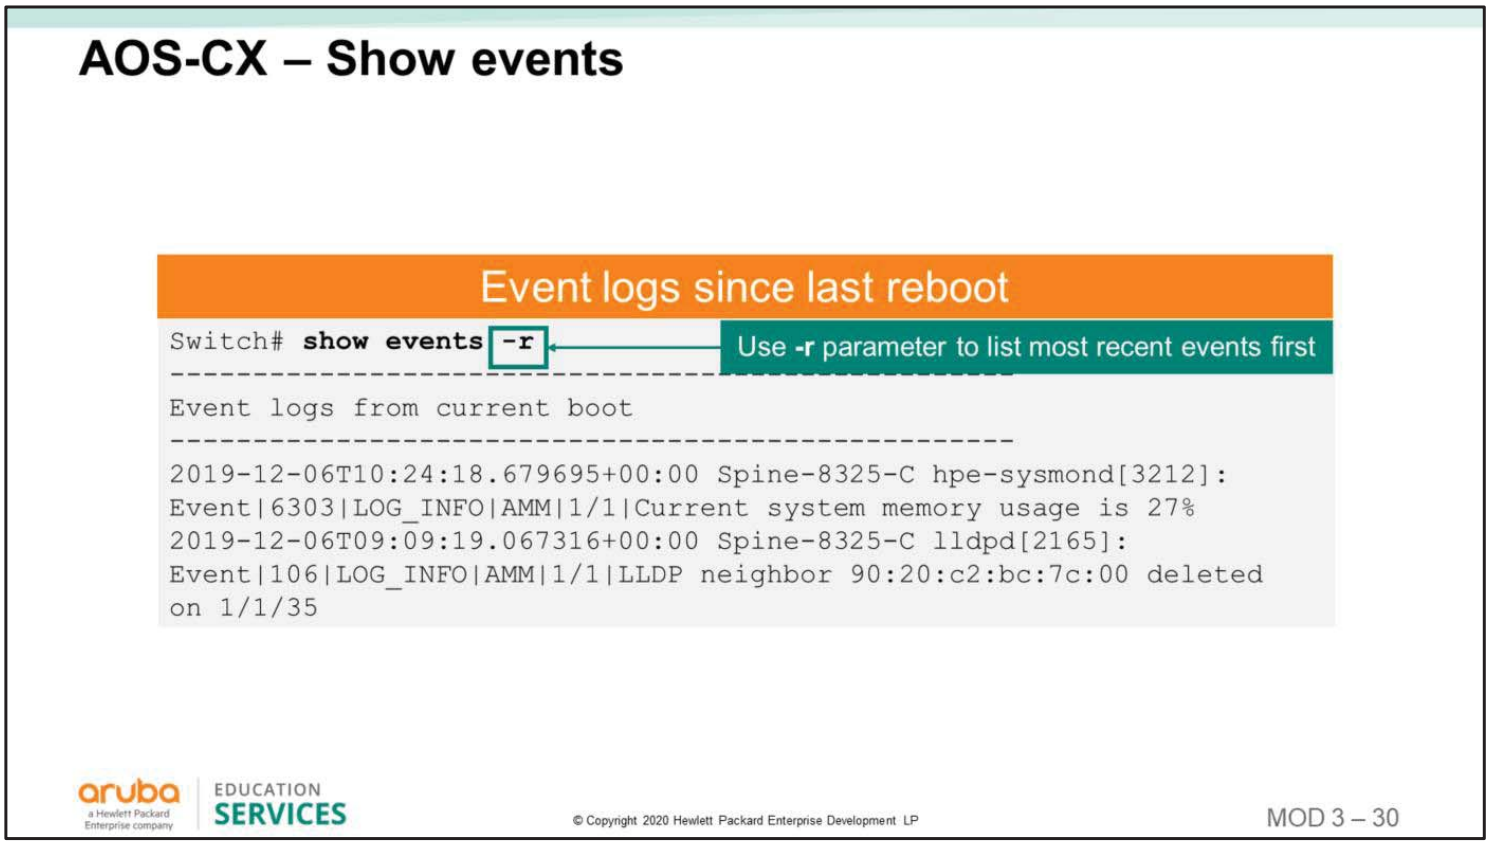

Show events. This command displays the event logs generated by the switch modules since the last reboot. Log analysis is a powerful tool to investigate and troubleshoot system and protocol related problems.

Command Context: Manager (#)
Authority: Administrators or auditors

|Use the -r parameter to list the most recent log events first.

Note: The show event command has other parameters that you can use like -e, -s, -a, -n, -c and -d. You can learn more about how to use these parameters in the Command-Line Interface Guide for AOS-CX document.

AOS-CX – Show interfaces brief

Available interfaces and status

```
Switch# show interface brief
--------------------------------------------------------------------------------
Port      Native  Mode   Type          Enabled Status  Reason          Speed
          VLAN                                                         (Mb/s)
--------------------------------------------------------------------------------
1/1/1     1       routed SFP28DAC3     yes     up                      25000
1/1/2     1       trunk  SFP28DAC3     yes     up                      25000
1/1/3     1       trunk  SFP28DAC3     yes     up                      25000

<<<Omitted Output>>>
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

show interface brief. This command helps the administrator to see what the available interfaces and the current status are. This command also briefly shows Layer 2 and Layer 3 configuration applied to the interface. This is a frequently used command for many network engineers. At a glance, you can tell the status of every interface on the entire device, with a specific focus on the Enabled and Status columns.
Enabled “yes” means that the interface is not disabled in the switch configuration – no administrator has disabled the interface. “Up” in the status column means that something is attached to that interface, and there is at least a good Layer 1/Layer 2 connection between them.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

Question #4

Getting switch information

Which commands could you use to validate network connectivity for an AOS-CX switch?

a. show running-config

b. show system

c. show interface brief

d. show events

e. show interface transceiver detail

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 32

Basic Configuration

Setup hostname and enable an interface

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

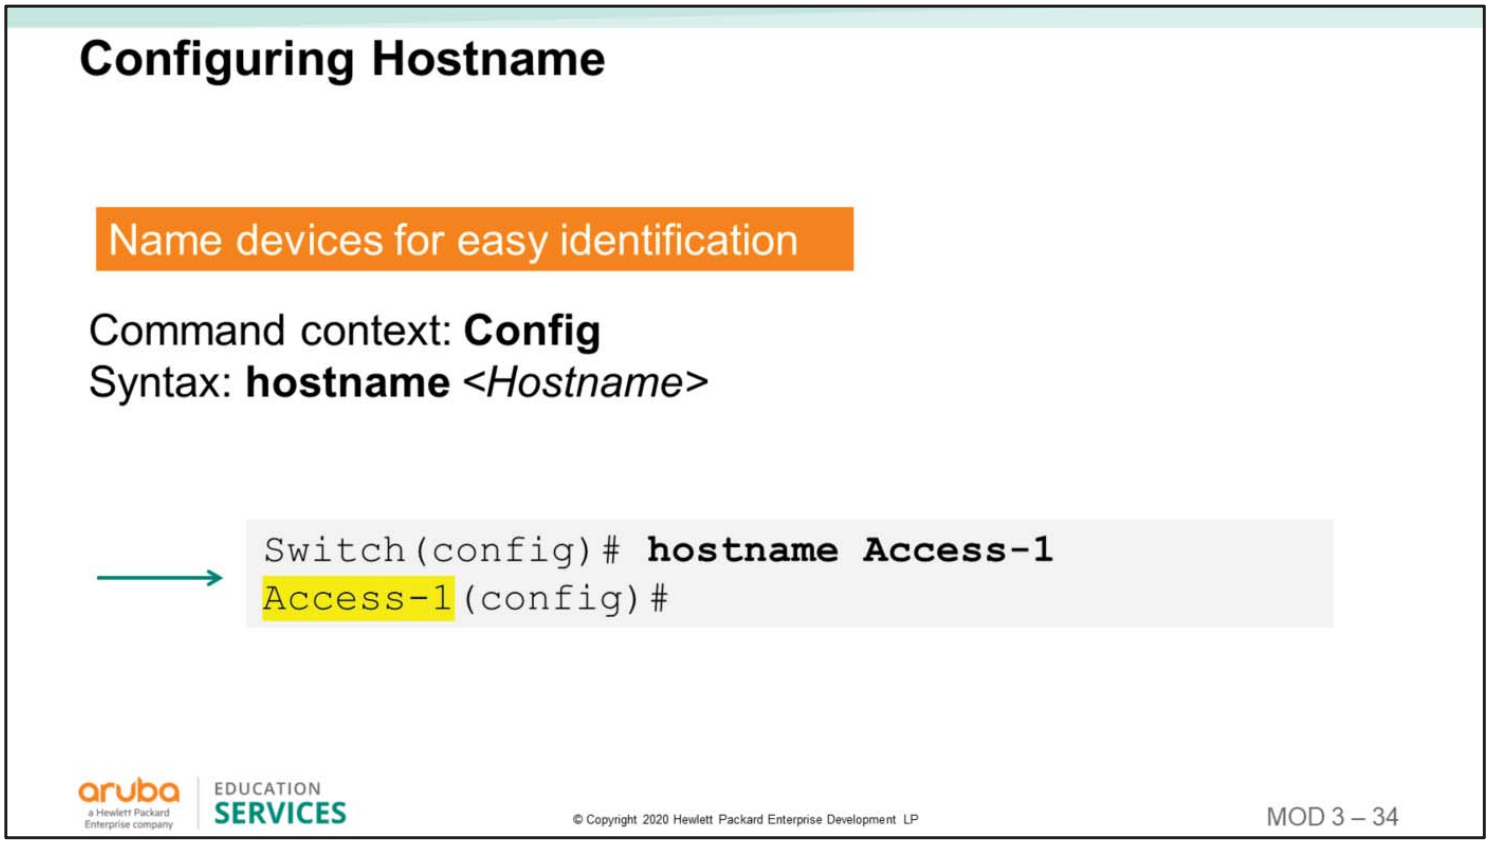

You should name each network device. This ensures that you can easily identify them, especially in large networks. In a real environment this is one of the first configuration steps that you will do.

Many corporations have documented, standardized naming conventions. For example, one standard might be Building-Floor-Rack_Name. So, the switches in Rack 2, on the 2nd floor of building 9 might be named 9-2-2_SWA, 9-2-2_SWB, 9-2-2_SWC, and so on.

The figure shows how to name devices, using the hostname command.

Command context: Config
Syntax: hostname <Hostname>

Example:

Switch(config)# hostname Access-1
Access-1(config)#

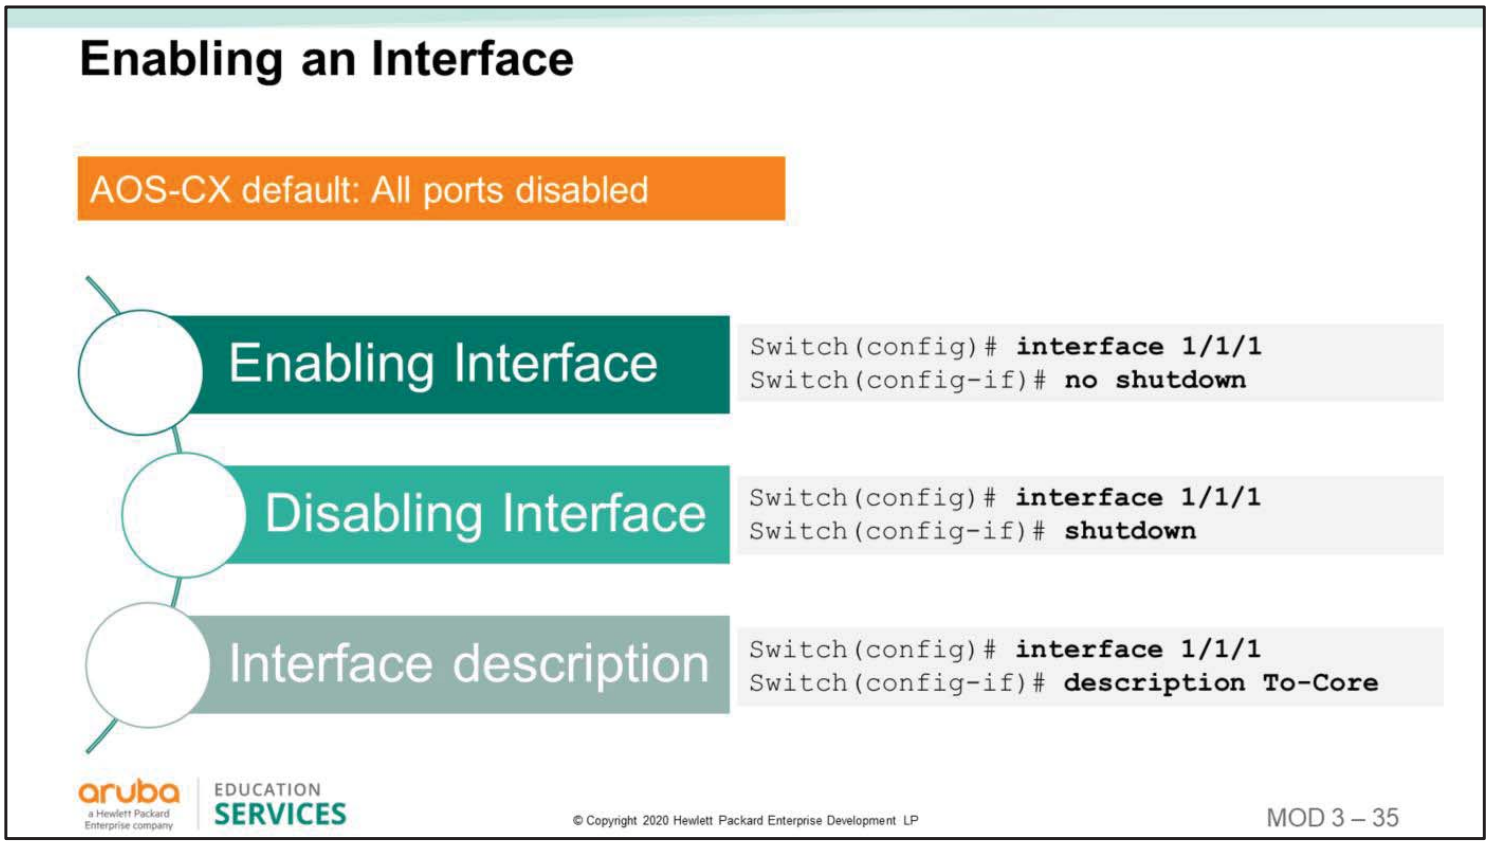

In AOS-CX, all ports are disabled by default. However you can modify this state using the no shutdown command, as shown in the figure.

|Command context: config-if (interface level)
Syntax: shutdown / no shutdown

Example: Enabling Interface

Switch(config)# interface 1/1/1
Switch(config-if)# no shutdown

|Example: Disabling Interface

Switch(config)# interface 1/1/1
Switch(config-if)# shutdown

A description can be configured for each interface, this facilitates management.

|Command context: config-if (interface level)

Syntax: description <Description>

Example:

Switch(config)# interface 1/1/1

Switch(config-if)# description To-Core-switch

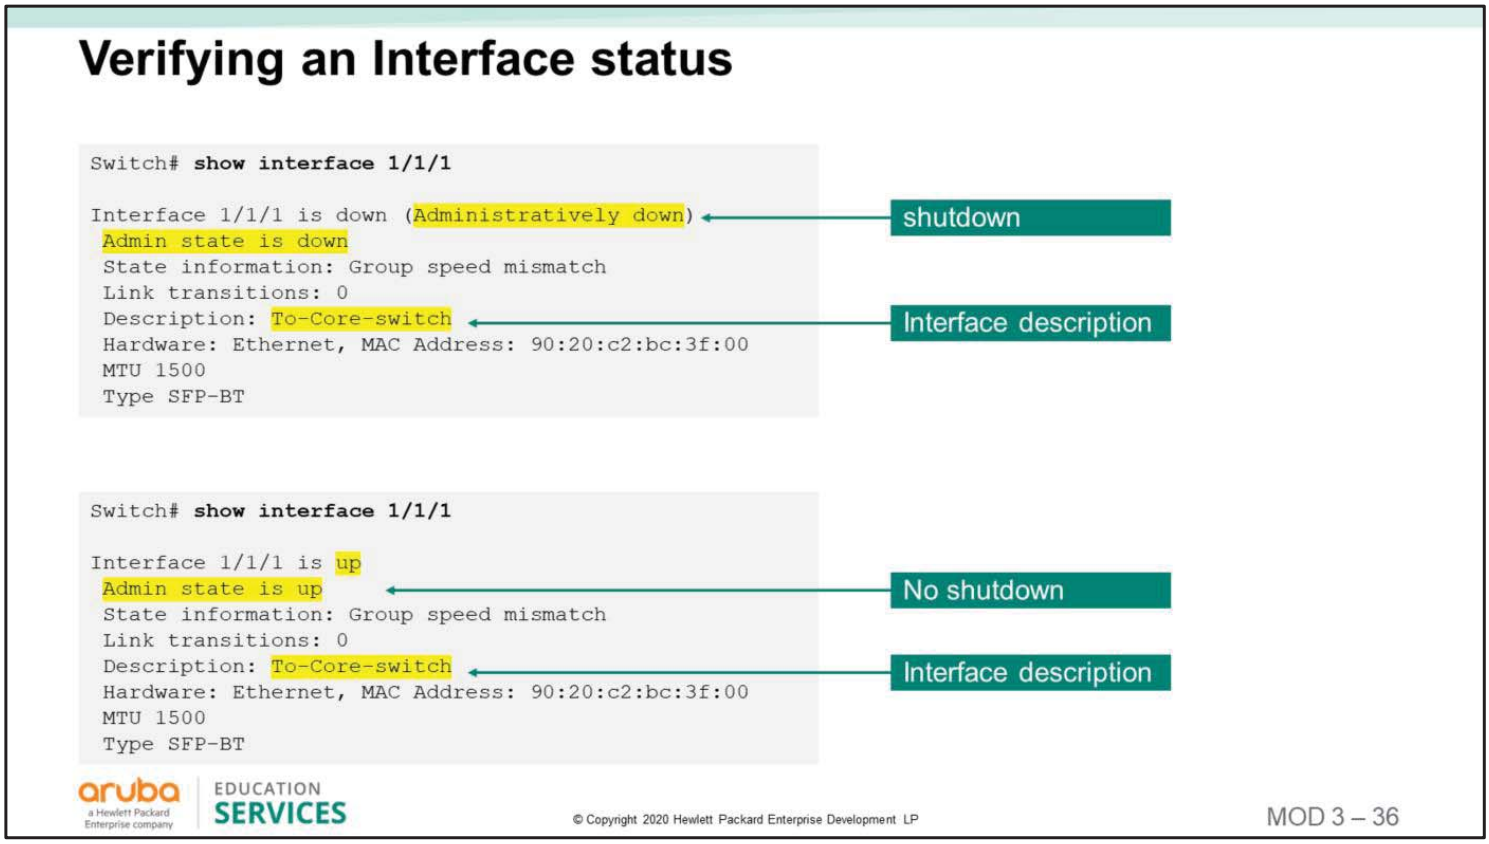

Show interface shows status and configuration for all interfaces on the switch.

|In figure's top example, "Administratively down" means that a shutdown command was configured on that interface.
|You can also see that an interface description was configured on that interface with the description command.

|In the bottom example, notice that the interface is now up. Somebody has entered the no shutdown command. The interface description remains unmodified.

Command context: Manager (#)
Syntax: show interface <Interface-ID>

Example: Interface Disabled

Discovering the network

LLDP

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Link Layer Discovery Protocol

Layer 2, IEEE standard protocol
- Network devices advertise identity/capabilities
- Helpful to discover and document

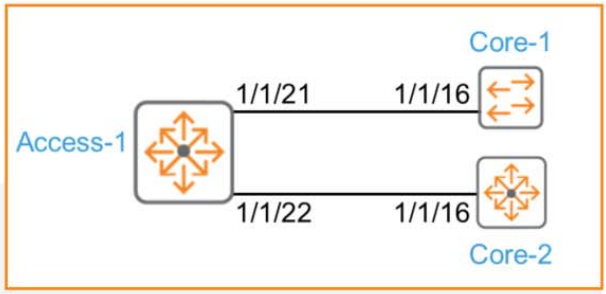


```
Access-1# show lldp neighbor-info

LLDP Neighbor Information
=========================
 <output omitted>

LOCAL-PORT  CHASSIS-ID         PORT-ID    PORT-DESC     TTL    SYS-NAME
-------------------------------------------------------------------------
1/1/21      90:20:c2:bc:a8:00  1/1/16     1/1/16        120    Core-1
1/1/22      90:20:c2:bb:8b:00  1/1/16     1/1/16        120    Core-2
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Imagine that you are remotely configuring an AOS-CX switch in a production environment. This switch is connected to multiple network devices, and you must configure only the interface that is facing an Aruba Controller. The onsite IT team is not available and somehow you must identify the correct interface so you can apply the configuration. This problem is easily addressed using LLDP.

LLDP is a vendor-neutral link layer protocol (Layer 2 protocol). It is used by network devices to advertise their identity and capabilities over a wired Ethernet connection. This protocol enables you to discover and document network device interconnections.

Media Endpoint Discovery is an enhancement of LLDP known as LLDP-MED. This provides for Auto-discovery of LAN policies such as VLAN-ID, Layer 2 priority and Differentiated services for QoS. This helps to enable plug and play networking.

Devices send LLDP information on a regular interval, encapsulated in Ethernet frames. Directly connected devices receive these frames and store the information in a table in local memory.
|You can view this information with commands like show lldp neighbor-info, as shown in the figure.

In the scenario, Access-1's port 1/1/21 connects to port 1/1/16 of Core-1. Access-1's port 1/1/22 connects to port 1/1/16 of Core-2. If someone has already taken the time to create a diagram like the one shown in the figure, you don't need to use LLDP for discovery – it has already been done for you. However, network documentation is often incorrect, out of date,

or was simply never created. You can verify existing documentation and create new documentation using LLDP information.

Note: Many folks do not believe that it is worth the time and money to create and maintain good, accurate network documentation – “The network is fine, why waste time making documents”. However, when the network fails, instead of having to discover the network (or rely on that one person who has the network memorized) everyone can look at the document and be far more effective troubleshooters. Documentation is an insurance policy against slow, ineffective troubleshooting. It can be the difference between a 4-hour network outage and a 15-minute outage.

In AOS-CX LLDP is enabled by default.

Show lldp neighbor-info. Displays information about neighboring devices for all interfaces or for a specific interface

Command context: Manager (#)
Authority: Operators or Administrators

To get more information about a specific entry you can append the interface at the end of the previous command. Show lldp neighbor-info 1/1/25, as shown below.

Access-1# show lldp neighbor-info 1/1/21

```
Port                        : 1/1/21
Neighbor Entries            : 1
Neighbor Entries Deleted    : 0
Neighbor Entries Dropped    : 0
Neighbor Entries Aged-Out   : 0
Neighbor Chassis-Name       : Core-1
Neighbor Chassis-Description  : Aruba JL635A  GL.10.04.0001
Neighbor Chassis-ID         : 90:20:c2:bc:a8:00
Neighbor Management-Address   : 10.250.1.3
Chassis Capabilities Available : Bridge, Router
Chassis Capabilities Enabled  : Bridge, Router
Neighbor Port-ID            : 1/1/16
Neighbor Port-Desc          : 1/1/16
Neighbor Port VLAN ID       : 1107
TTL                         : 120
```

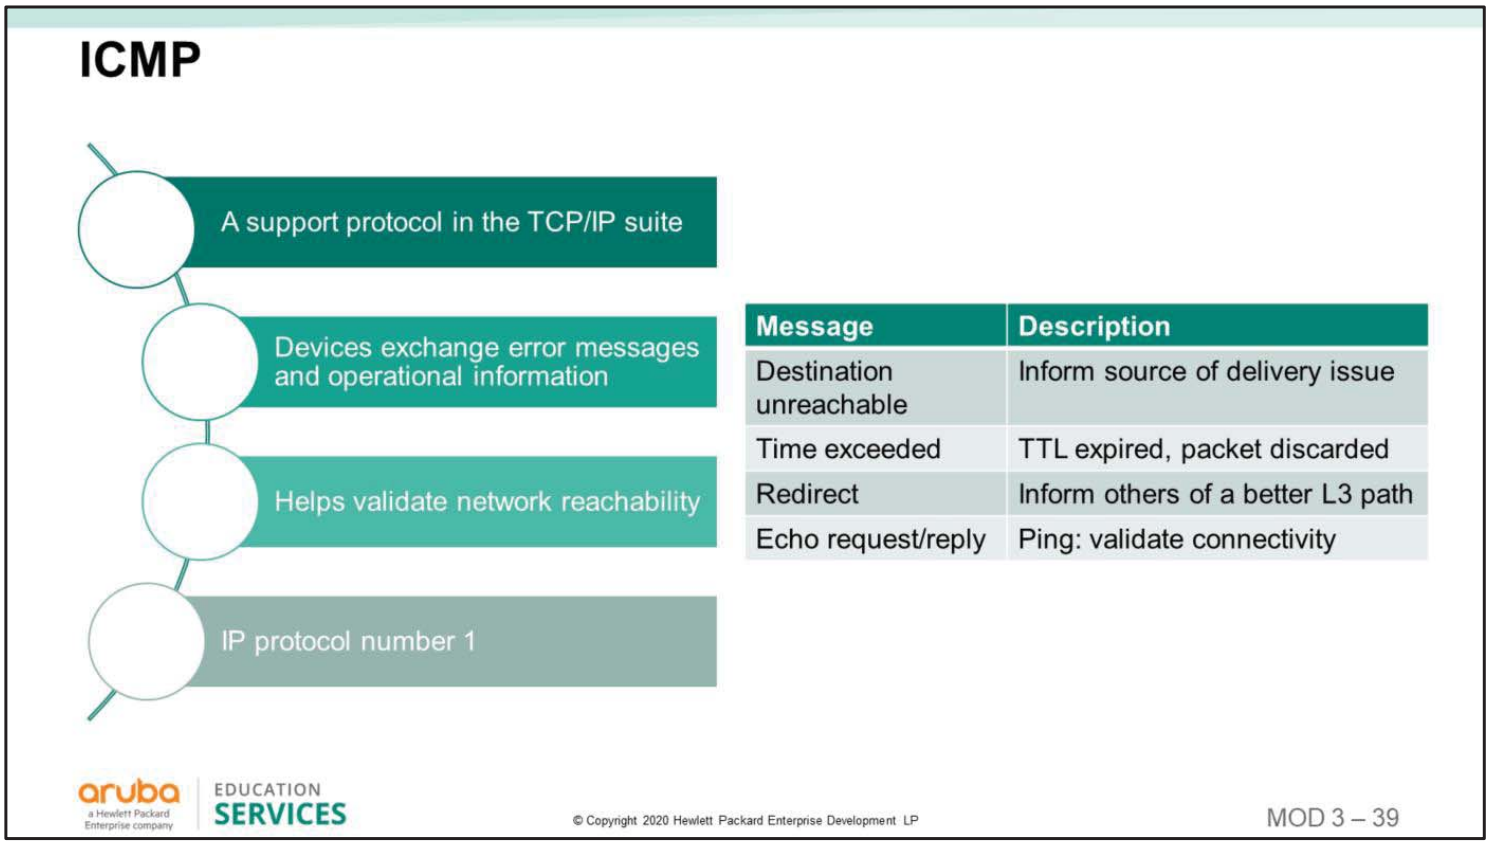

Suppose that a new printer is added into the network and a laptop wants to use this device. The first troubleshooting step is to verify connectivity between these two devices, to ensure that the network is properly configured. ICMP helps you to address this problem.

The Internet Control Message Protocol (ICMP) is a supporting protocol in the Internet protocol suite. Network devices use ICMP to send error messages and operational information, to indicate success or failure when communicating with another IP address. For example, an error is indicated when a requested service is not available or that a host or router could not be reached.

The table shows the common messages that are used in ICMP, starting with Destination unreachable. Other messages include:
Time exceeded
Redirect
Echo request / echo replay

Note. ICMP messages do not include any TCP or UDP header. ICMP messages are placed directly into an IP packet. Thus, ICMP is often considered to be an extension of the Network Layer.

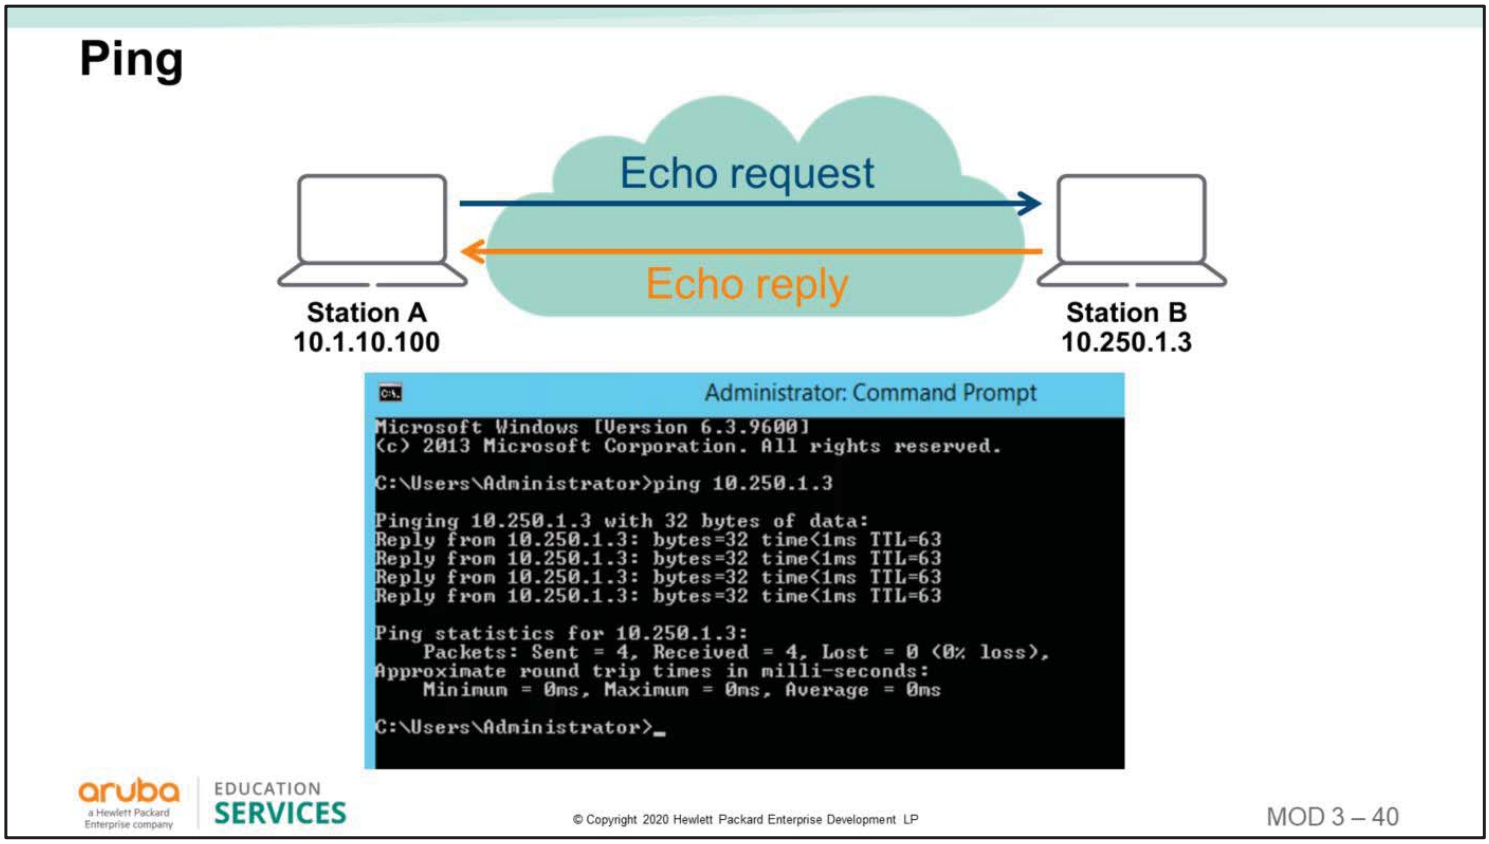

The ping command sends ICMP Echo request messages to the destination device. The remote device responds to each received Echo request with an Echo reply message. When the source device receives the messages back, it is sure that the network was properly configured – communication succeeded between the two devices.

In AOS-CX the ping command must specify the IPv4 or hostname.

Syntax: ping <IPv4-address or hostname>
Command context: Operator (>) Manager (#)
Authority: Operators or Administrators

|The ping command is supported in most of the Client Operating Systems. The example ping session shown in the figure is from the command prompt of a Windows Client.

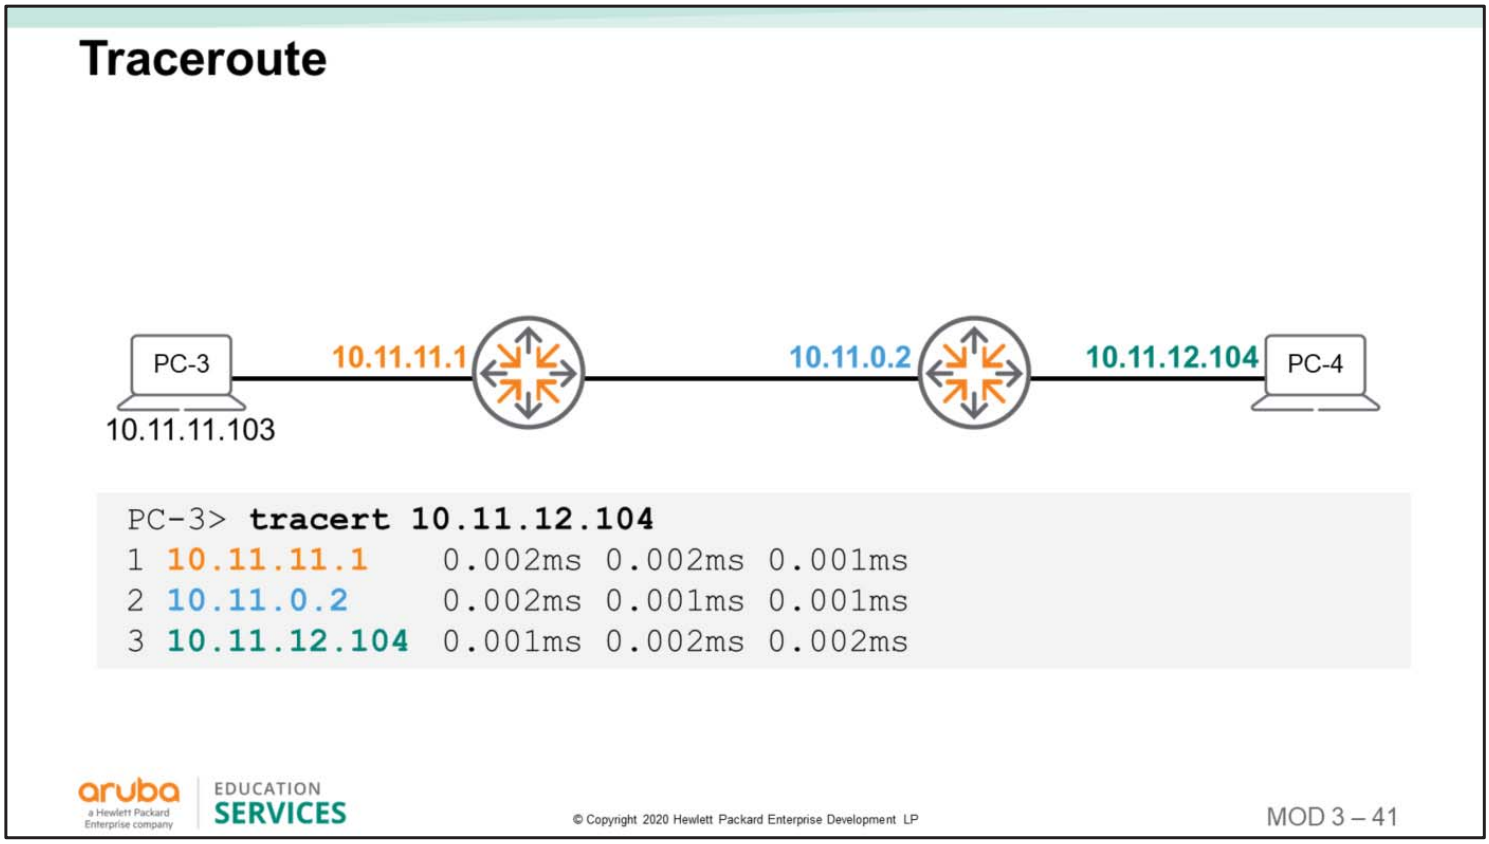

Traceroute is a variant of the ICMP protocol and it is a useful troubleshooting tool when you deploy a Layer 3 technology, such as a routing protocol. This command tracks the path that a packet takes to reach its destination listing the Layer 3 devices that are in the path.

In AOS-CX the traceroute command must specify an IPv4 address or hostname.

Syntax: traceroute <IPv4-address or hostname>
Command context: Operator (>) Manager (#)
Authority: Operators or Administrators

|In this example, the switch tries to establish connectivity to the destination 10.11.12.104 , via two intermediate routers (10.11.11.1 and 10.11.0.2). You will learn more about routing and Layer 3 in the next chapters.

Note: Many firewalls block ping and trace commands. If you are attempting to use ping and trace commands to test connectivity through a firewall, the tests may fail. This does not necessarily mean that anything is wrong – it just means that the firewall is dropping your ping and trace commands as a security measure.

Power over Ethernet

Power feature

- Grants power to endpoints using Ethernet cable
- Eliminates the need for external power sources

Name	IEEE Standard	Power to PD	Maximum power per port	Energized Pairs
PoE	802.3af	12.95 W	15.4 W	2-pairs
PoE+	802.3at	25.5 W	30 W	2-pairs
PoE++	802.3bt (Type 3)	51 W	60 W	4-pairs
PoE++	802.3bt (Type 4)	71.3 W	100 W	4-pairs

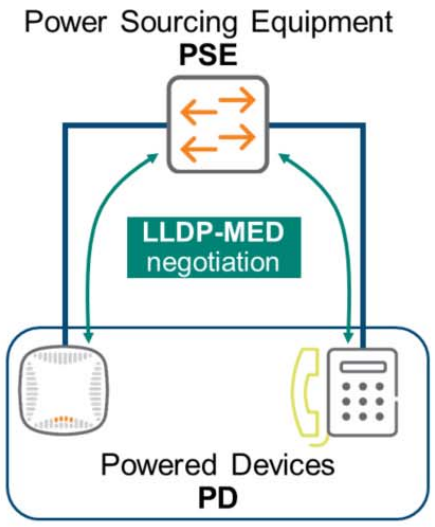

AOS-CX switches support all PoE standards

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 42

Power over Ethernet or PoE is a feature implemented on network devices such as Layer 2 and Multi-Layer switches, which provides power to endpoints using an Ethernet cable. This feature eliminates the need of external power sources saving expenses on materials and installation time.

PoE uses the twisted pairs of an UTP cable to send power to PoE-enabled devices. The first PoE standard uses two twisted pairs to transmit data while the other 2 are used for power transmission. With the new PoE standards power and data are sent over the four twisted pairs. If you are thinking how is that possible, the answer is that power and data use different frequencies, Electricity uses low frequency (60 Hz or less) and data uses high frequencies (10-100MHz).

The standard consider two types of devices, Powered devices (PDs) are devices that receive PoE power, while power sourcing equipment (PSE) provides the power. Typical PDs are Access Points, IP Phones, cameras and some IoT devices. On the other hand network switches are consider PSEs.

The table displays the 4 different standards that are available in the industry, AOS-CX switches support all standards.

PSEs and PDs can negotiate the power that PDs require more precisely using LLDP messages.

Question #5

Network discovery

Which of the options below accurately describe network discovery commands or techniques?

a. LLDP is a Layer 2 discovery protocol.

b. You typically use the ping command to help document your network.

c. If the ping or traceroute command fails, you know that you have a connectivity issue.

d. The ping command leverages ICMP echo requests and echo replies.

e. It is best to use LLDP after your network documentation is complete.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 43

Correct answers:
a.
d.

Lab Activity

Lab 3. Initial Setup

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Additional Content

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

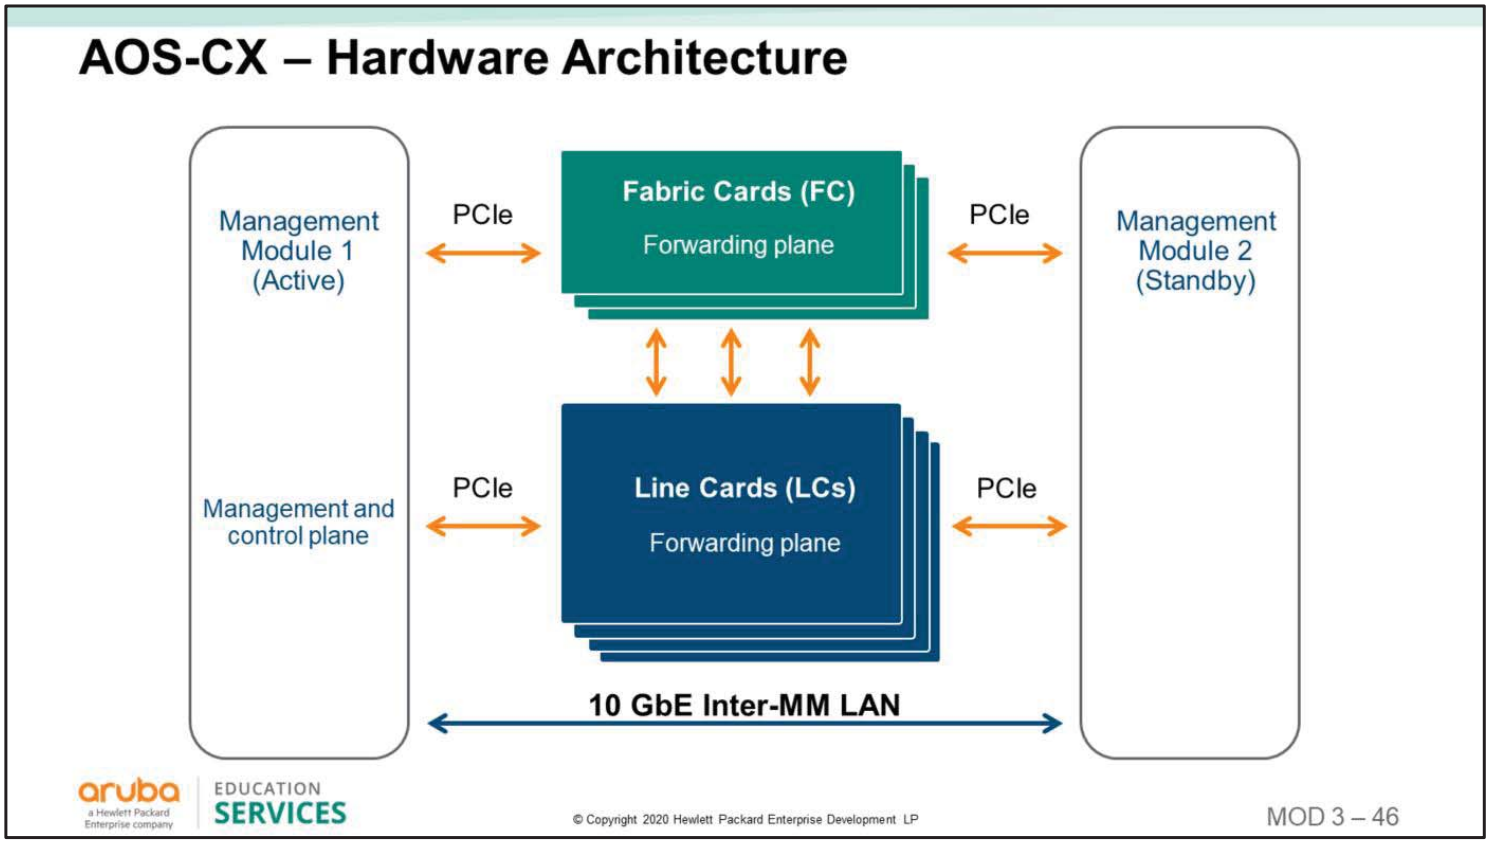

The AOS-CX modular switch series (Aruba 6400 and 8400) hardware architecture include three major components: Management Modules, Fabric Modules and Line Cards.

Management Modules. This component has to main purposes. First, it runs the management plane, for monitoring and configuration services. The management module also runs the control plane, which defines what to do with the incoming information by running protocols and algorithms.

Fabric Modules. This component helps to interconnect the multiple Line Cards that can be installed in the switch. A fabric card forwards data between the ingress Line card and the egress line card. This device makes decisions based on information derived from data packets, and so is considered part of the data or Forwarding plane.

Line Cards. This component works on the Forwarding plane. It decides where traffic should be sent. Data that must be forwarded to another port on the same line card uses and internal process within that line card. Traffic that must be sent to a different line card is sent via the Fabric Module, which selects the proper destination Fabric Card.

The diagram simplifies the components and the communication paths between them.

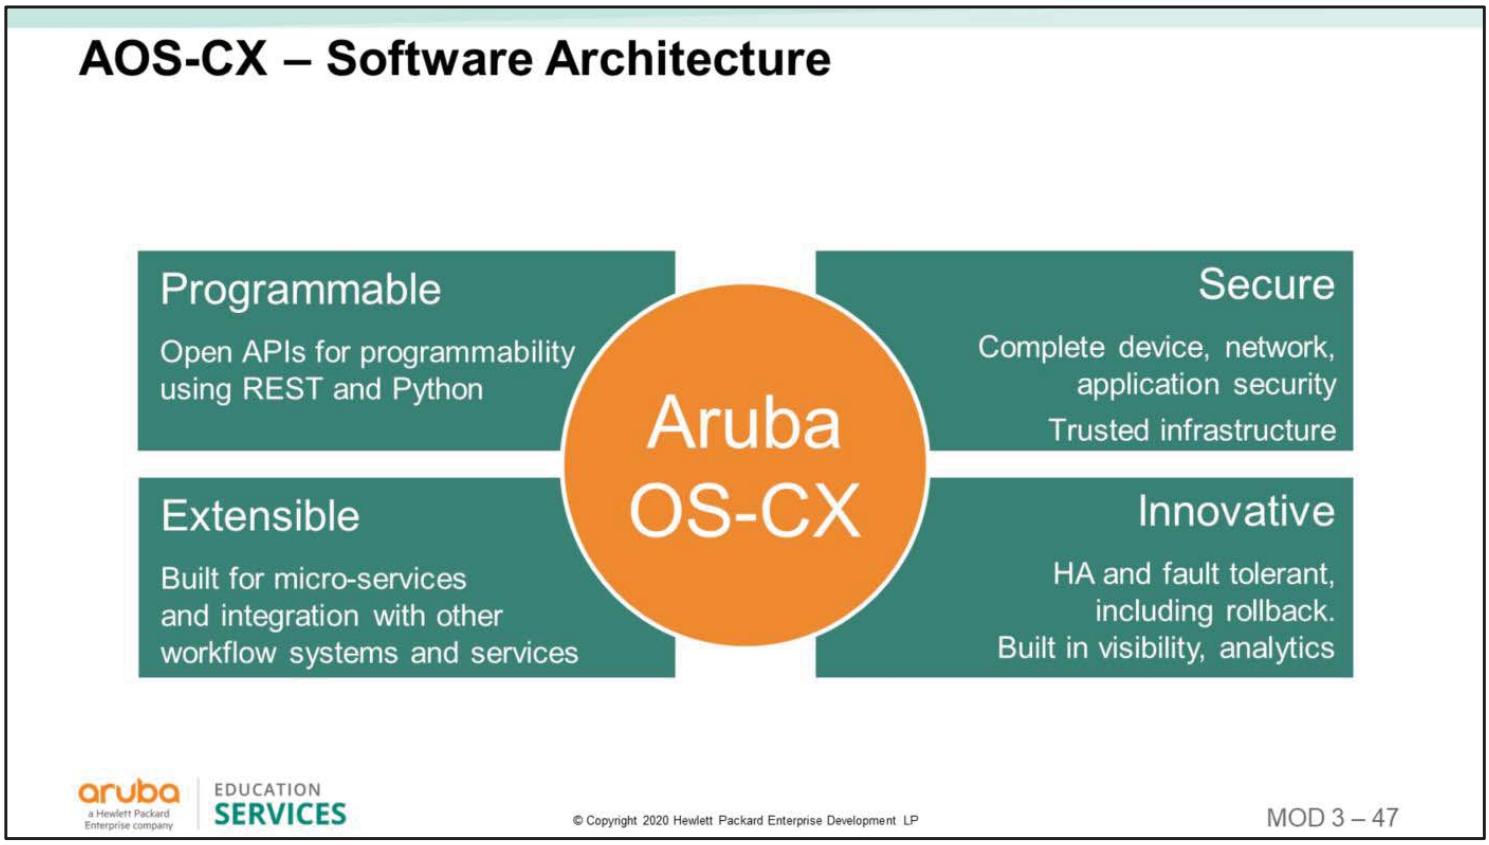

The AOS-CX software is a modern, database-driven operating system that automates and simplifies many critical and complex network tasks. A built-in time series database enables customers and developers to use software scripts for historical troubleshooting and to analyze past trends. This helps to predict and avoid future problems due to scale, security, and performance.

This network operating system is built on a modular Linux architecture with a stateful database which helps to offer the following unique capabilities:

Easy access to all network state information allowing unique visibility and analytics.
REST APIs and Python scripting for fine-grained programmability of network tasks.
A micro-services architecture that enables full integration with other workflow systems and services.
Continual state synchronization that provides superior fault tolerance and high availability.
Security best practices were applied to completely create a trusted platform.

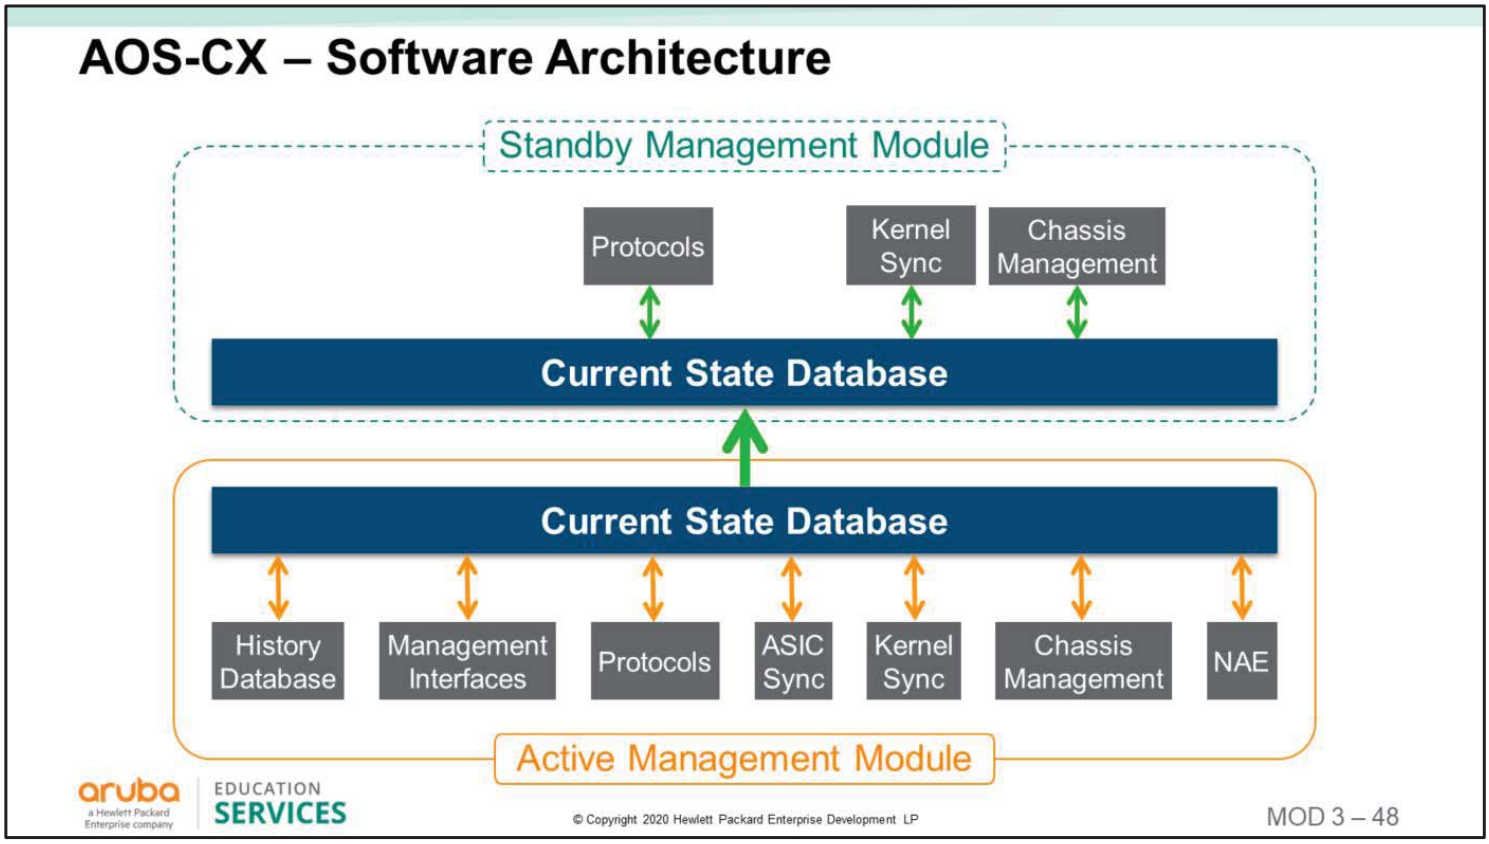

The Current State Database is the most important aspect of the AOS-CX software architecture. All software processes communicate with the data base rather than each other. This model ensures near real-time state and resiliency. Using the Current State Database, you can upgrade software modules independently.

The figure shows how processes like the History Database or Protocols interact directly with the database and not between each other. This streamlined approach allows all processes to only use one language to talk to the database. Without this model, direct inter-process communication would be less efficient, wasting CPU resources.

The database also maintains the current configuration, status of all features and statistics. The unified database ensures that all information is visible in a single place. Thus, interaction is accomplished through a single, open API.

The diagram also shows how the Active Management Module (MM) can synchronize information on the standby MM. This ensures a fault tolerant system that reduces downtime. Network Protocols do not have to wait and re-converge.

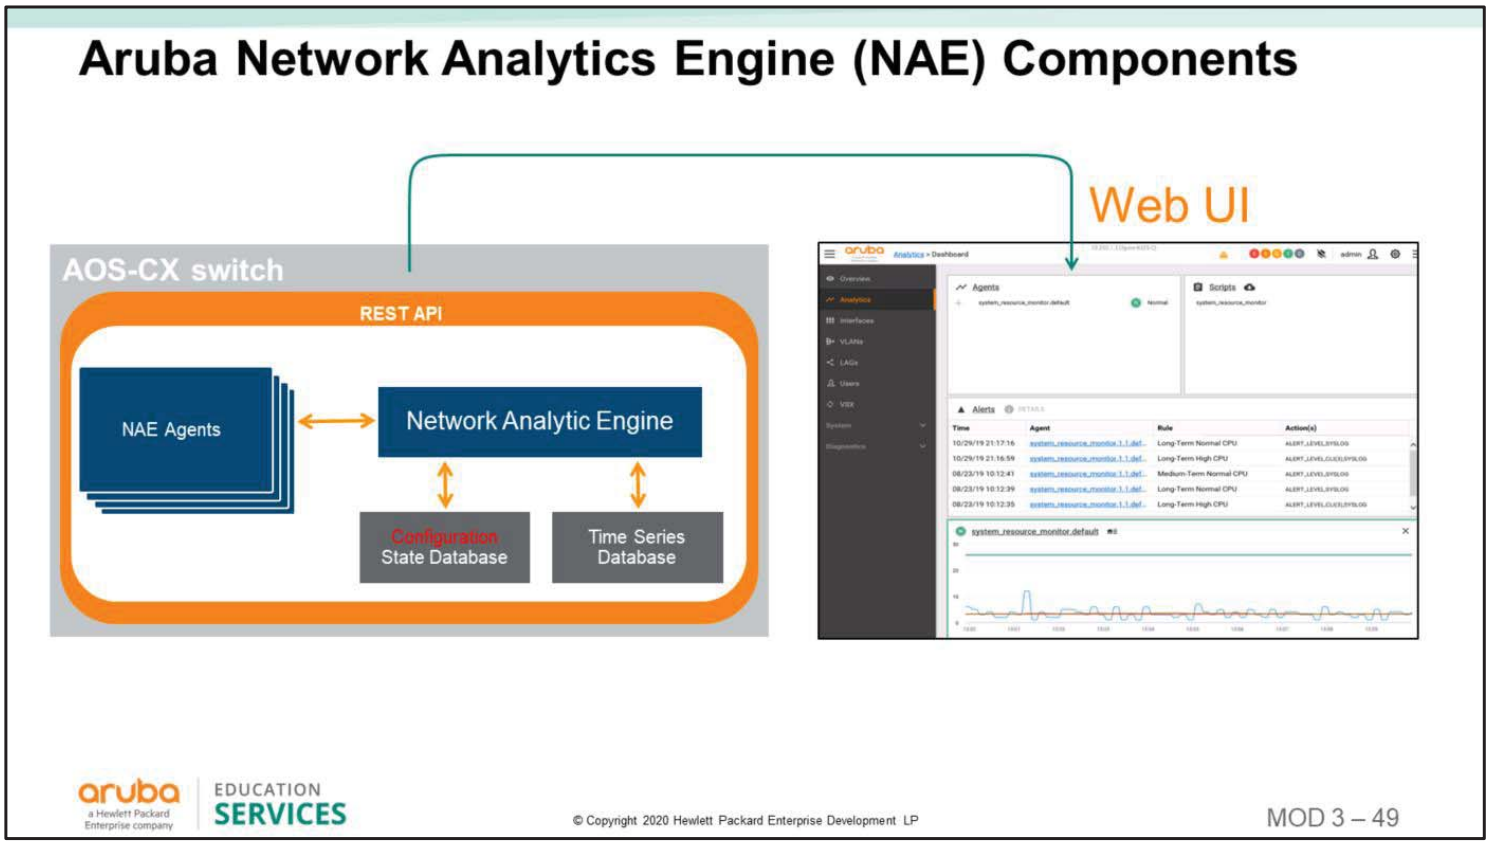

NAE is made up of agents, rules, databases, APIs, and a Web UI, as shown in the figure and described below.

NAE Agents. The built-in NAE makes use of agents to collect context. Agents are scripts, triggered in the device when a specific event occurs. Then it collects additional interesting and relevant network information.
NAE Rules. Agents are triggered by user-defined rules. For example, you can create a rule to collect information when CPU utilization exceeds a certain threshold for a specified period.
Configuration and state database. This database enables NAE's direct access to the current configuration and switch operational states. Data retrieved from this database can be used to analyze trends and predict future capacity requirements.
Time series database. This database gives the users the ability to rewind and playback the network context surrounding a network event. Under normal use, storage is estimated at 400 days.
REST APIs. This communication method enables integration with external systems, such as SIEM tools and log analytic engines. Also, operators can use the APIs to request information from other devices in the network. This helps to create a complete picture of the network state when a specific event occurs.
Web UI. Lets you access, view, and configure NAEe agents, scripts and alerts. Automatically generated graphs provide additional context, useful for troubleshooting networks.

AOS-CX – Top memory

This command displays memory utilization information

```
Switch# top memory
top - 20:41:22 up 17 days, 23:46, 1 user, load average: 0.31, 0.26, 0.24
Tasks: 238 total, 2 running, 193 sleeping, 0 stopped, 0 zonbie
%Cpu(s): 2.5 us, 1.6 sy, 0.0 ni, 95.9 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 7868164 total, 3013768 free, 1398760 used, 3455636 buff/cache
KiB Swap: 2097148 total, 2095868 free,     1280 used. 6080024 avail Mem

    PID USER      PR    NI      VIRT      RES      SHR S    %CPU %MEM        TIME+ COMMAND
   2366 root      20     0   2013036   129684    45568 S     0.0 1.6     739:59.05 /usr/bin/switchd_agent -i 1 -s
   3047 root      20     0    791184    84416    32232 S     0.0 1.1     189:43.37 /usr/bin/prometheus --
    701 root      20     0   1212616    75652    19316 S     0.0 1.0      18:02.19 /usr/bin/hpe-restd
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 50

Top cpu. This command shows CPU utilization information.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

Top memory. This command shows memory utilization information.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

AOS-CX – Show interfaces transceiver detail

Lists the transceivers installed in the switch.

```
Switch# show interface transceiver detail
Transceiver in 1/1/1
 Interface Name      : 1/1/1
 Type                : SFP28DAC3
 Connector Type      : No separable connector
 Transfer Distance   : 0.00km (SMF), 0m (OM1), 0m (OM2), 0m (OM3)
 Diagnostic Support  : None
 Product Number      : JL488A
 Serial Number       : CN97KBG0BJ
 Part Number         : 8121-1743
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 3 – 51

Show interfaces transceiver detail. This command lists the transceivers installed in the switch.

Command context: Operator (>) or Manager (#)
Authority: Operators or administrators

You began this module by exploring the networking devices used to build a typical network. This knowledge can help you to design, implement, and troubleshoot networks. You learned about important network services that are vital for proper network function and end-user interactions, along with exploring 2-tier and 3-tier network design models. You reviewed Aruba's portfolio of switches, to help ensure that you can select the proper device for a given scenario.

Then you learned how to use Aruba switches – how to connect to them and use the CLI to glean basic network health and status information, perform basic configuration, and discover how an unknown network is connected. This ability to document a network will serve you well when you must diagnose network outages.

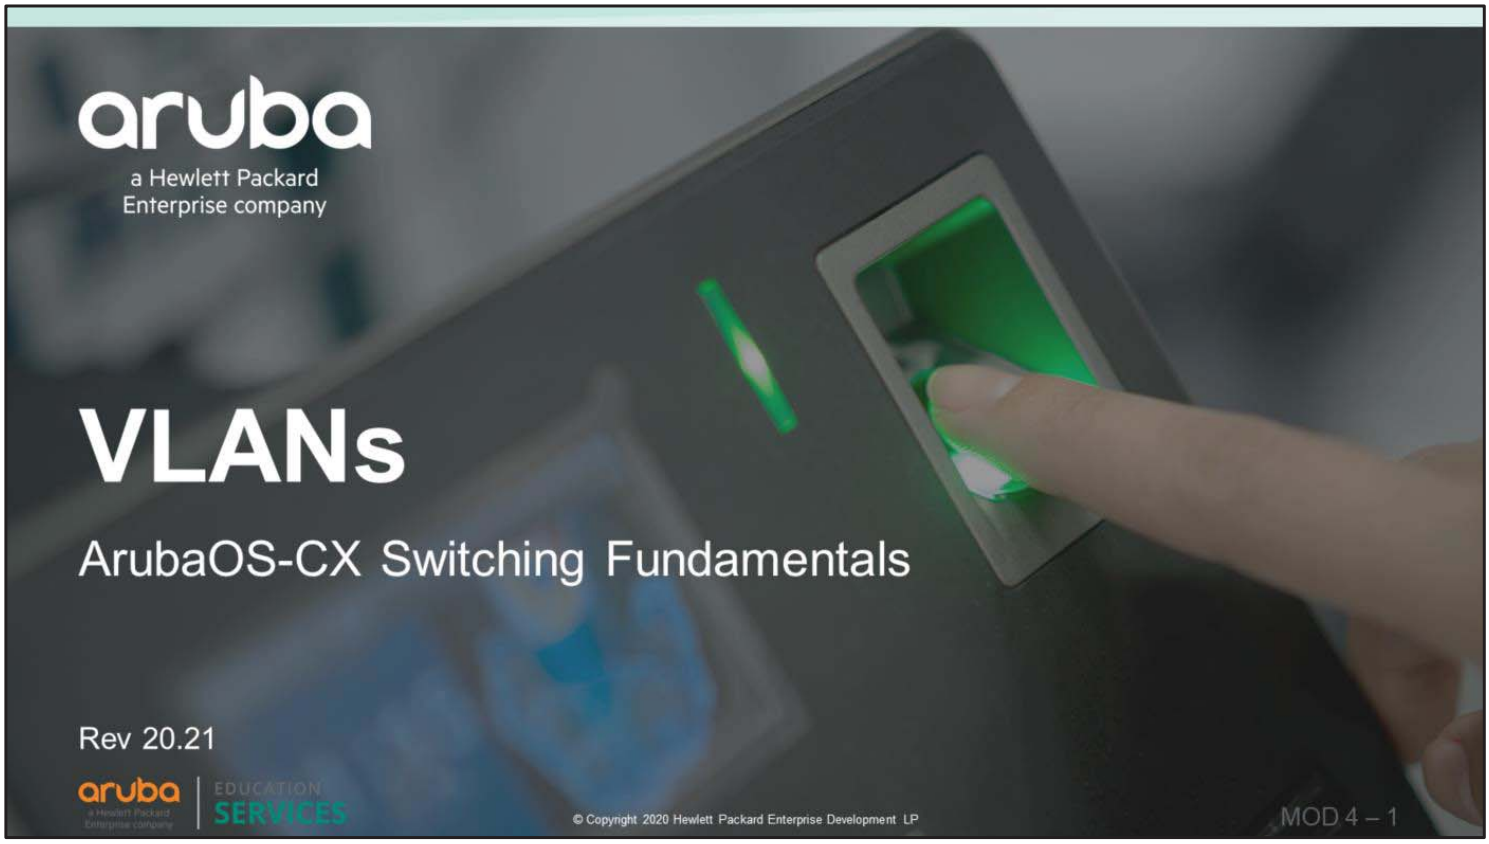
aruba
a Hewlett Packard Enterprise company
VLANs
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 4 – 1

Objectives

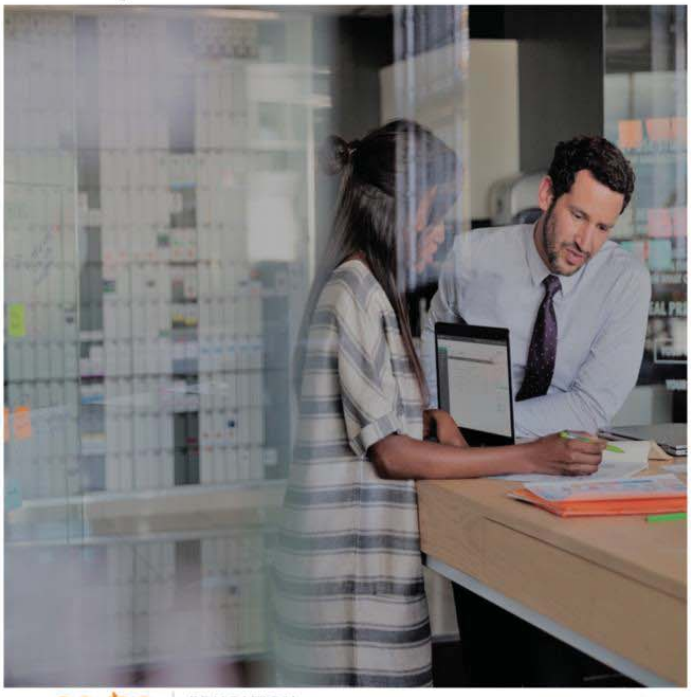

–Learning Objectives

- Compare and describe collision and broadcast domains
- Describe and configure VLANs
- Explain 802.1Q tagging and
- Configure trunk ports
- Describe MAC address and ARP forwarding tables
- Describe the frame delivery process

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 2

After completing this module, you will have a very good understanding of vital Layer 2 concepts and protocols, with a particular focus on Virtual LANs (VLAN).

| You will be able to compare collision domains and broadcast domains, and how you can leverage this knowledge to improve network performance.

| You will be able to describe and configure VLANs – perhaps the most important foundational aspects of deploying resilient, secure, and high-performance networks.

| You will be able to explain 802.Q tagging, and how this is used to extend the VLAN concept across multiple switches.

| You will be able to leverage this knowledge about tagging to configure trunk ports.

| You will be able to describe MAC address tables and forwarding tables, and how they are used to improve network connectivity and efficiency.

| And finally, you will be able to describe the frame delivery process, to more fully understand network communications.

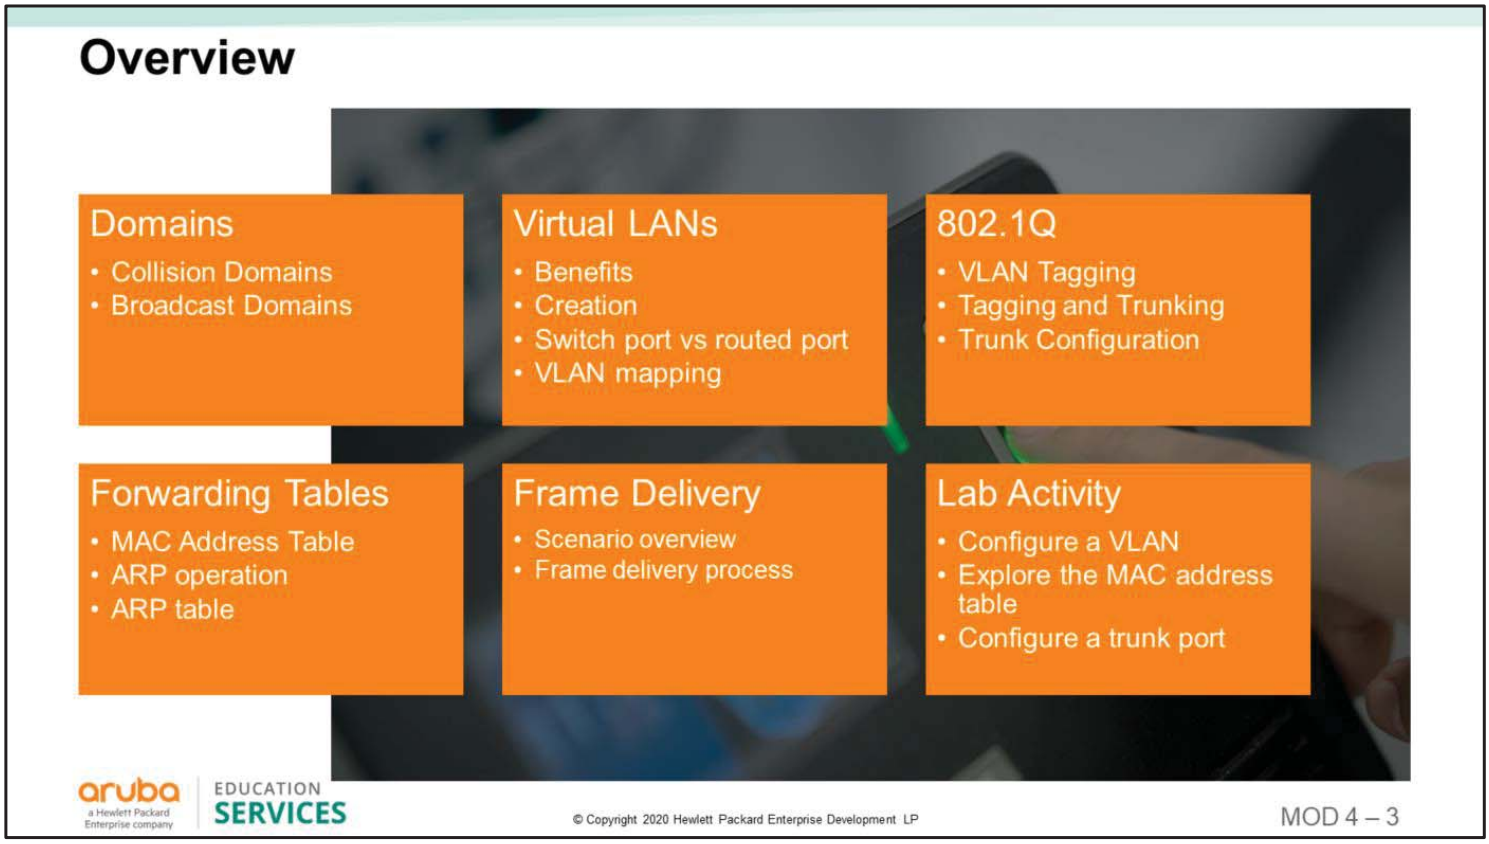

This module is intended to improve your understanding of vital Layer 2 switch concepts, which will speed your journey toward mastery of Layer 2 network design, deployment, configuration, and troubleshooting.

First you will learn about Layer 2 collision domains and broadcast domains – important for both your theoretical knowledge, and your practical ability to improve network performance. Then you will learn about VLANs, and how vital they are to creating scalable, flexible, and secure networks. Related concepts include switch ports, Switch Virtual Interfaces (SVI), and physical routed ports. You will turn this theoretical knowledge to practical use by learning how to configure all of these entities.

You will learn about some potential limitations of VLANs, and how they are eliminated with the 802.1Q trunking protocol. With this knowledge and configuration ability, you will be able to extend the VLAN concept across multiple physical switches.

With such an ability to create larger, more scalable networks comes a need to understand forwarding tables, including the MAC address table that is built by Layer 2 switches, and the ARP cache, which is maintained by both end systems, Layer 2 switches, and Layer 3 routers.

Armed with all of this prerequisite knowledge, you will be ready to explore a typical scenario of two devices communicating over a network. This will unite the earlier information from this module into a real-world scenario. You will see the interaction of Layer 2 frame addressing, VLANs, 802.1Q, the MAC address table, and the ARP cache.

In the Lab activity you will create a couple of VLANs and you will configure a trunk port to

expand the VLAN concept across different switches. Finally you will explore the MAC address table.

Domains

Collision and Broadcast

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

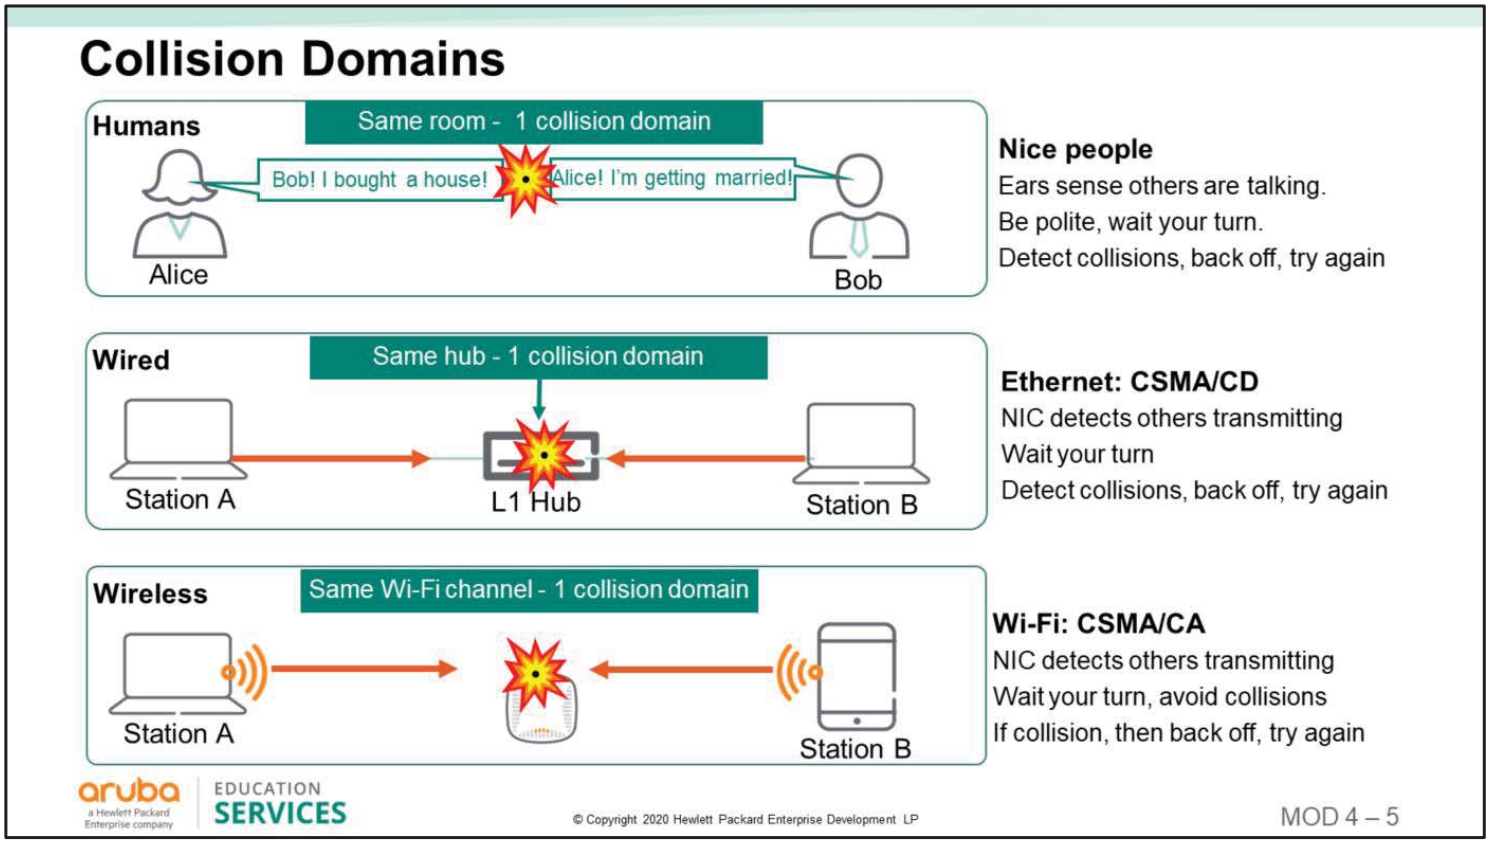

Suppose that Alice and Bob are having a conversation. Alice has something to say, but her ears detect that Bob is speaking, so she politely waits her turn.
|If Alice and Bob happen to speak at the same time, the sound of their voices collide, making it difficult to understand. Alice and Bob realize what is happening and stop. Bob says, “I’m sorry, please go ahead and speak”. Alice says, “No, please you go ahead”. They both back off, and then try again.

|Ethernet is similar - only one Ethernet host on a particular segment may transmit at a time. This is all controlled by an algorithm running on all Ethernet NICs called Carrier Sense Multiple Access / Collision Detection (CSMA/CD). This works very much like humans, as described above.

Multiple humans can be in the same small room, and all of them can access the airwaves and speak any time they want - it’s a Multi-access system, like Ethernet. Before speaking however, a polite human first listens to sense if others are currently speaking – this is like Ethernet’s Carrier Sense mechanism. Ethernet cards sense the state of a “carrier signal” that indicates a currently transmitting NIC.

|If two devices transmit at the same time, the electrical signals collide and become corrupted. This is called a collision. This condition is detected, and each station backs off and tries again. This is like when Alice and Bob both start talking at the same time, realize what is happening, and stop. Bob says, “I’m sorry, please go ahead and speak”. Alice says, “No, please you go ahead”. They both back off, and then try again.

Ethernet collisions occur when a hub is deployed in networks. A hub is a Layer 1 device that receives digital 1’s and 0’s, and repeats them, near instantaneously, out all other ports.

Therefore, a hub is also known as a multi-port repeater. All 16 devices connected to a 16-port hub are in the same collision domain - if one device transmits, 15 other devices must wait. If two hosts transmit at the same time, a collision occurs, everyone backs off, waits a few milliseconds, and tries again.

Thankfully, hubs are outdated devices that are no longer used. These unintelligent, layer 1 repeaters have been replaced by intelligent Layer 2 switches. For a properly deployed and configured switch under normal circumstances, collisions do not occur on Layer 2 switches.

|Similarly, Wi-Fi uses an algorithm called Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA). This works very much like Ethernet's CSMA/CD, only it adds additional mechanisms to preemptively avoid collisions.
|Everyone attached to the same wireless Access Point (AP) channel is in the same collision domain. Therefore, only one of them may transmit at a time.

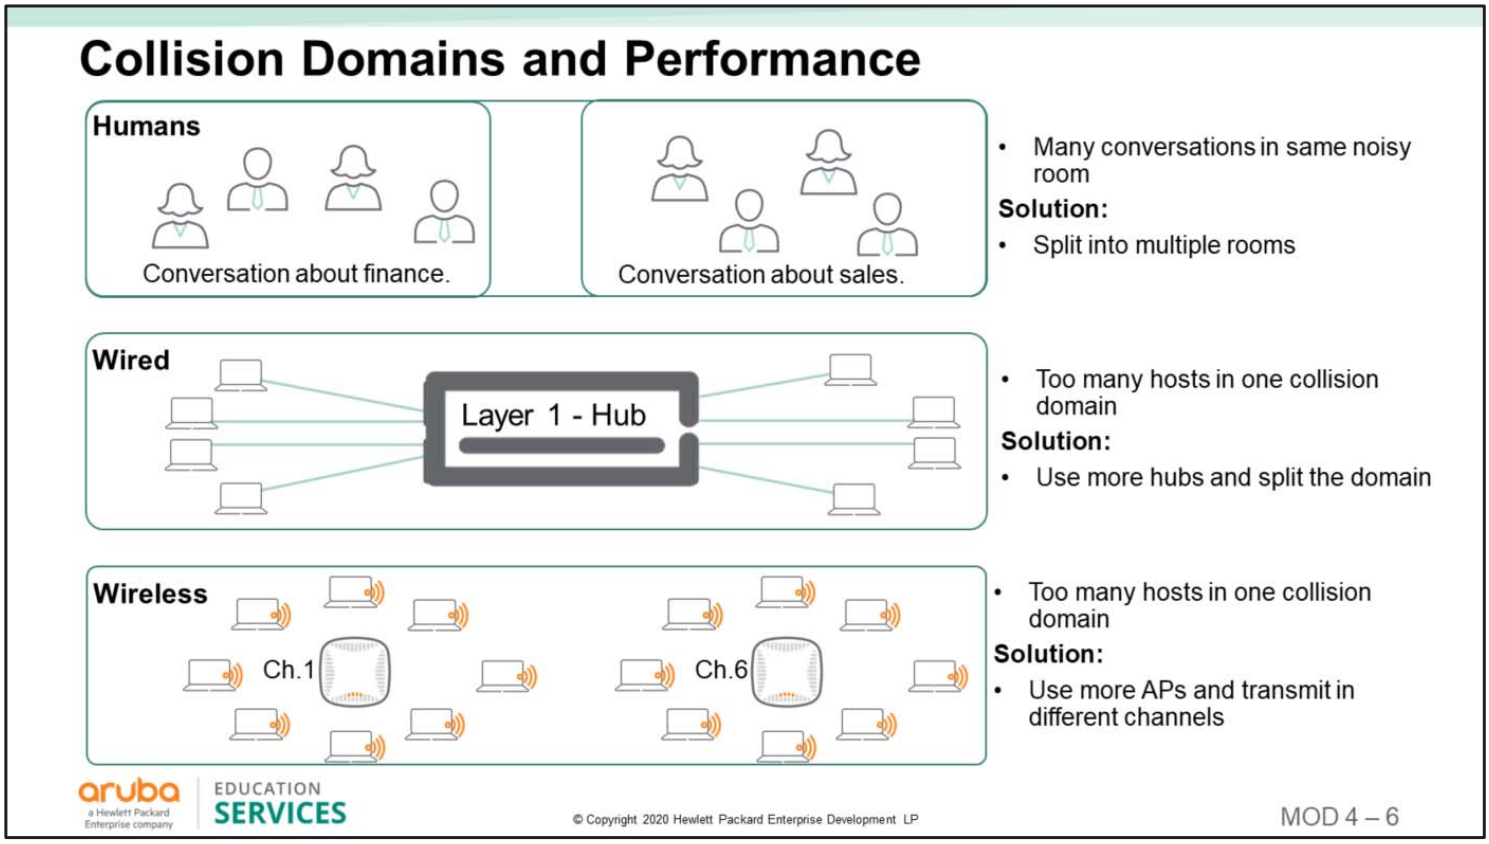

Your understanding of collision domains enables you to optimize network performance.

Consider the humans. With only two people in the room, conversation is quick, easy, and efficient. Then there are 8...10...50 people in the room.
|Multiple conversations occur, causing distractions and “collisions” (people talking at the same time).
|You may want to split people off into different rooms to improve communications and prevent people from having to wait too long.

|Suppose that you have an Ethernet hub with 8 connected devices. While one host transmits, seven hosts must wait their turn. If you had a 48-port hub, then 47 hosts must wait. Imagine the performance impact!
|You would want to split this network up into multiple collision domains.
Thankfully, hubs are relatively outdated devise that are no longer used. These unintelligent, layer 1 repeaters have been replaced by more intelligent Layer 2 switches. For a properly deployed and configured switch under normal circumstances, collisions do not occur on Layer 2 switches.

|Similarly, consider 1 AP, with transmit power set to maximum. This increases its coverage area, such that it can service all 75 people in the room. This is nice, because you save money – you provided coverage for everyone with a single AP. However, while 1 host transmits, 74 people must wait. People complain that wireless is slow. Nobody is happy.
|To resolve the issue, you purchase more APs, set their transmit power much lower, and set each to a unique channel. Now about 25 people connected to each AP – each is a different collision domain. Now three people can transmit at a time, and performance is greatly improved. Most good wireless engineers add more APs and reduce their power, such that

only 10-25 people connect. This can drastically increase performance.

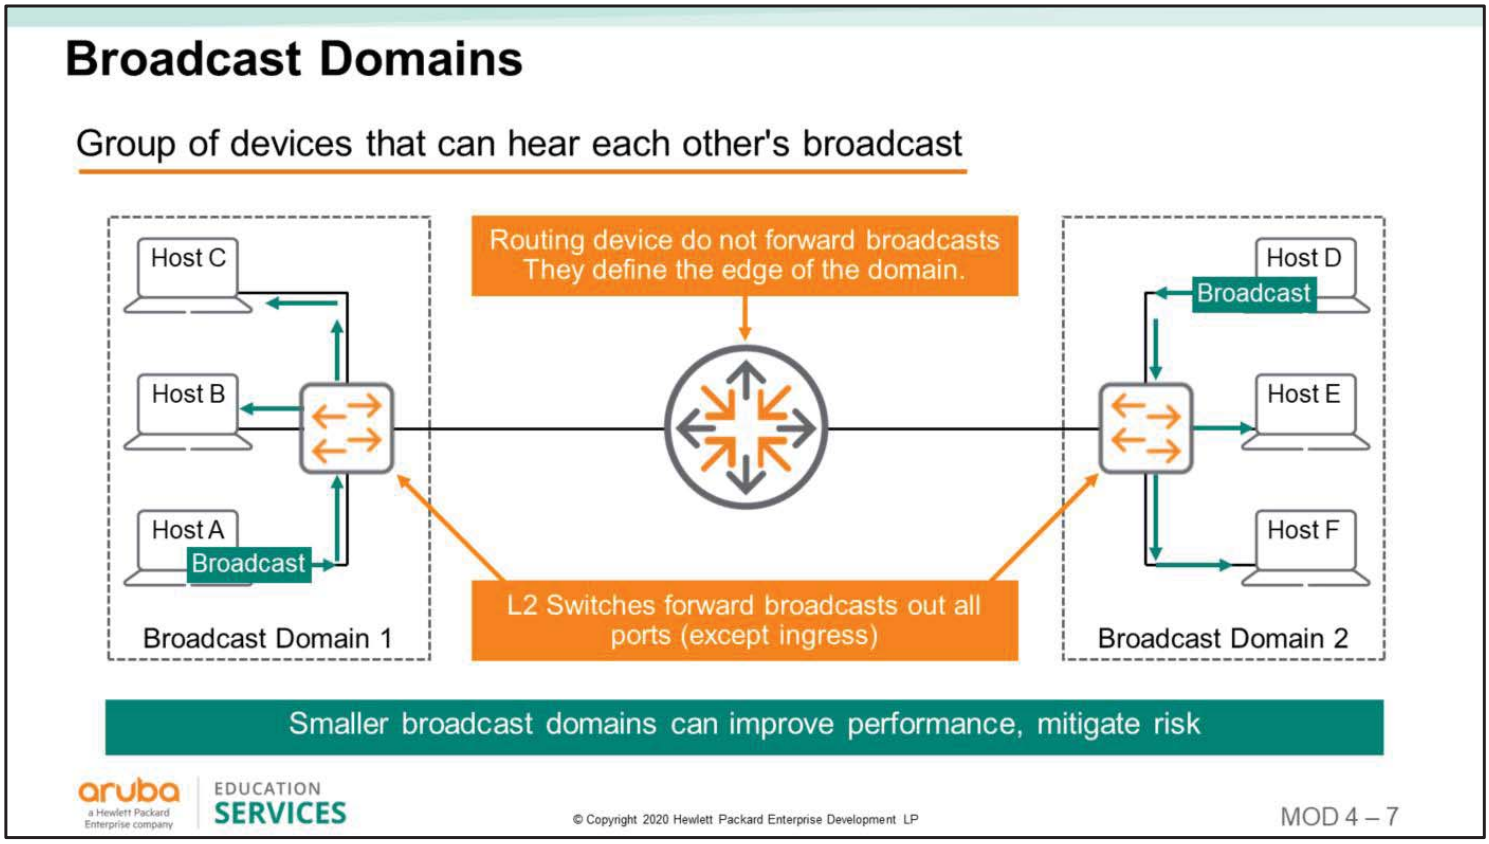

A broadcast domain is simply a group of devices that are on the same network, capable of receiving and responding to a broadcast frame from any device.

Recall that Module 1 introduced a broadcast as a type of communication where a single device contacts all other devices. A broadcast message uses a special Layer 2 MAC address: FF:FF:FF:FF:FF:FF and a special Layer 3 IPv4 address: 255.255.255.255.

|Layer 2 switches always forward any received broadcast frame out all ports, except for the port on which the broadcast was received (the ingress port).
|In the figure, Host A transmits a broadcast address. Perhaps it is saying, “Hey everyone! Who has IP address 10.1.1.1?” Because the switch forwards this out all other ports, hosts B and C receive the message. Of course, nobody in broadcast domain 2 receives the message – they are not even connected!

|Similarly, when Host D sends a broadcast, all hosts in domain 2 receive it, and nobody in domain 1 receives it.

But now suppose that you want to connect these two networks into an internetwork. Do you recall what type of device does this?
|A router. Now the two networks are connected. However, unlike Layer 2 switches, layer 3 routers do not forward broadcasts. In other words, they define the edge of the broadcast domain. Thus, broadcast traffic travels exactly as before – all domain 1 hosts receive broadcasts from domain 1, but not from domain 2, and vice versa. Of course, now that they are connected, all devices can communicate with unicast or multicast traffic.

But why not simplify this network? Eliminate the router, and simply connect all hosts together on a single network, and in a single broadcast domain. You save money – no need to purchase a router.

|As with collision domains, large broadcast domains cause performance issues. For unicast traffic, from Host A to Host F, only those two stations must fully process packets. However, with broadcast frames, all stations must process the traffic. Switches must forward, or flood broadcast frames out all ports (except the ingress port). This can increase utilization on switches, and increase bandwidth utilization – every switch link must carry every broadcast. Hackers might even write programs to generate millions of broadcast packets and flood the network, leaving no resources for valid traffic. This is called a Denial of Service (DoS) attack. Smaller broadcast domains means better performance.

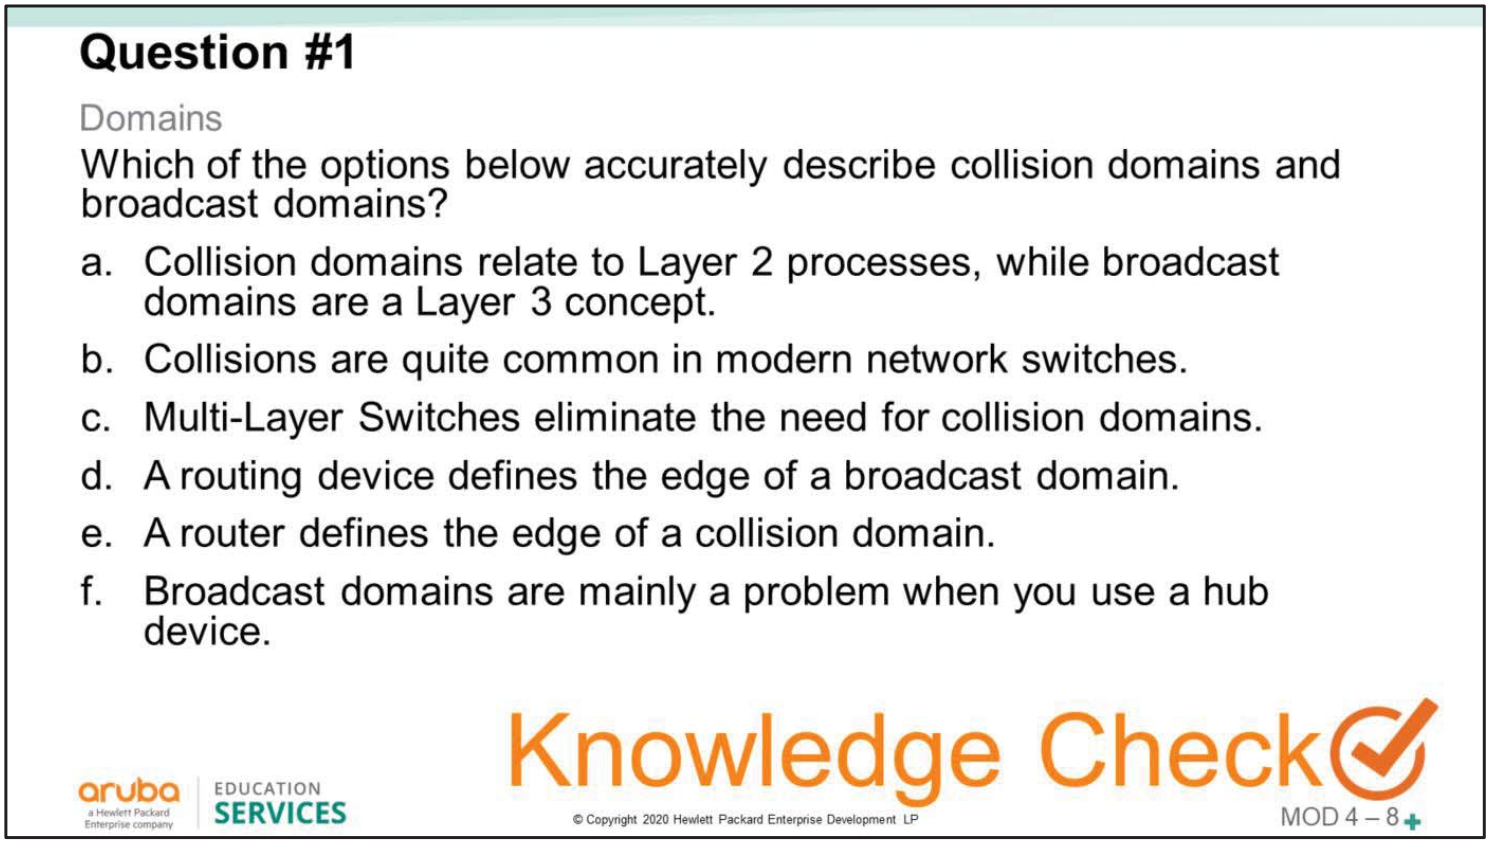

Correct answer:
d.

Virtual LAN

Benefits, creation and mapping

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

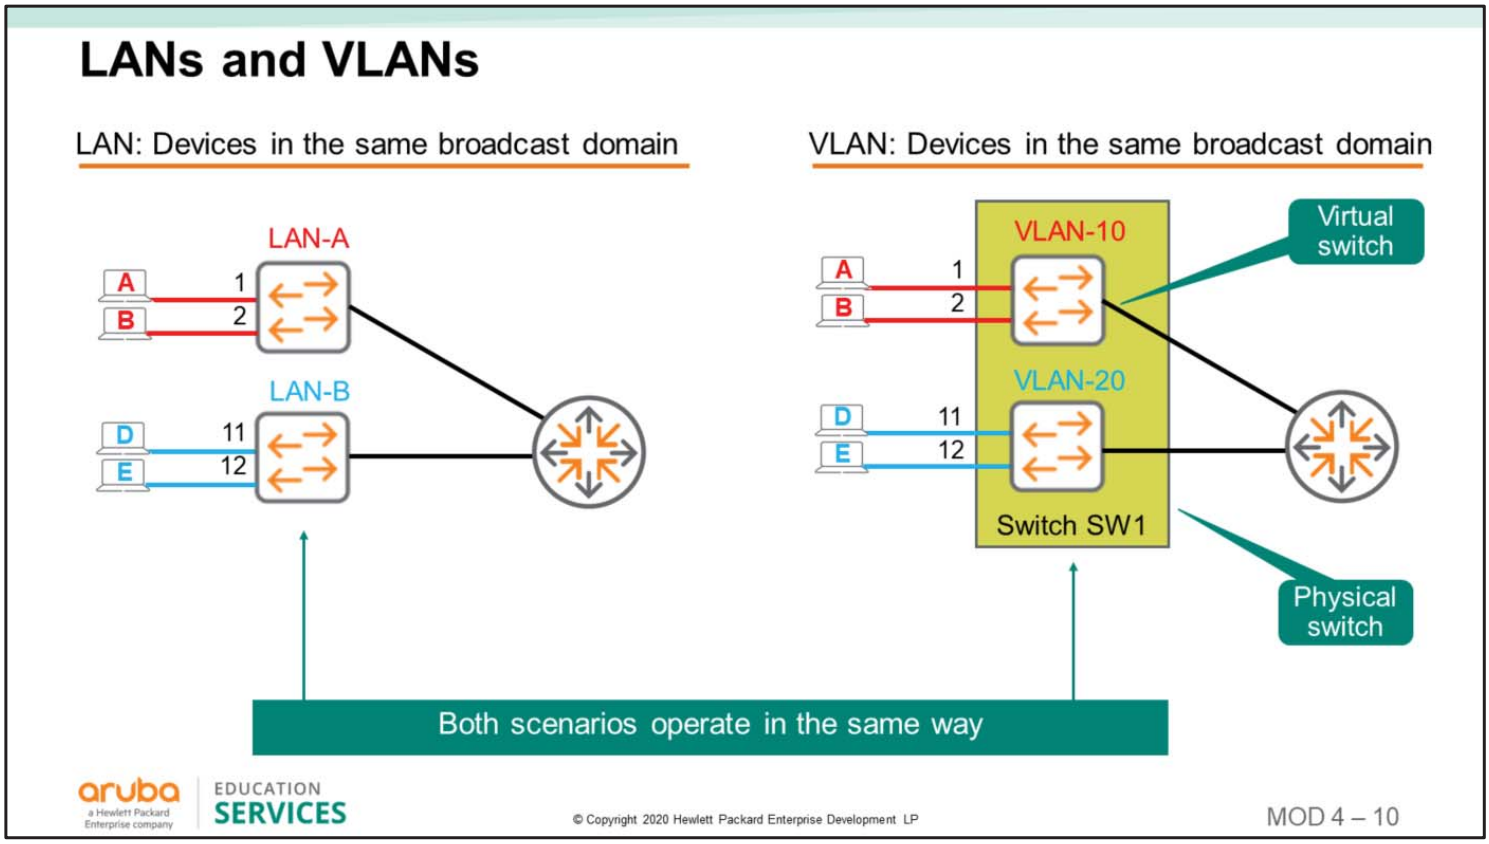

A Local Area Network (LAN) is a group of devices in the same broadcast domain. The figure shows two LANs, labeled LAN-A and LAN-B. LAN-A switch has two connected devices. Host A connects to Port 1 and Host B connects to port 2. The other switch has similar connections.

Recall that if a switch receives a broadcast on any port, it floods it out all other ports, except the ingress port. Therefore, if Host A sends a broadcast, the LAN-A switch forwards it out all other ports, and so host B receives the frame. Of course, hosts D and E do not receive this frame – there is no connection between LAN-A and LAN-B. They are physically separate LANs, connected to physically separate switches.

What if you wanted to connect these three Local Area Networks into an Internetwork? |Simply add a router, which can router unicast and multicast traffic between the LANs.

|Now consider a Virtual LAN (VLAN), which is, just like a physical LAN, a group of devices in the same broadcast domain. |Suppose that you have a single, physical Aruba switch named SW1.

|Hosts A – E connect to ports 1, 2, 11, and 12 on this switch. By default, all these devices are in the same broadcast domain. If Host A sends a broadcast, all other hosts receive the frame. Now suppose that you learned some new switch syntax, and created a Virtual LAN named “VLAN10”.

|It is as if you have created a small virtual switch, inside the physical switch. This virtual switch exists, but it is not connected to any of the physical switch ports.

|And so you define or “map” physical ports 1 and 2 as being members of the red VLAN 10.

|Similarly, you create VLAN20, and assign ports 11 and 12 as members of the blue VLAN 20.

You have now created two separate broadcast domains on a single physical switch. When host A sends a broadcast, the switch knows to only forward that frame out all ports that are in the same VLAN, except the ingress port. Only host B receives the broadcast. If host D broadcasts, only E receives it.

|You have effectively recreated the scenario on the left, using a single physical switch. Just like the physical scenario on the left, there is NO connectivity between the separate VLANs. No unicast, multicast, nor broadcast traffic shall pass between the VLANs.
|Of course, you could always connect a router to create an internetwork, just like you did with the physical switches.

But if the scenarios operate in the same way, why bother with VLANs ? Let’s talk about the advantages of VLANs and learn some new syntax.

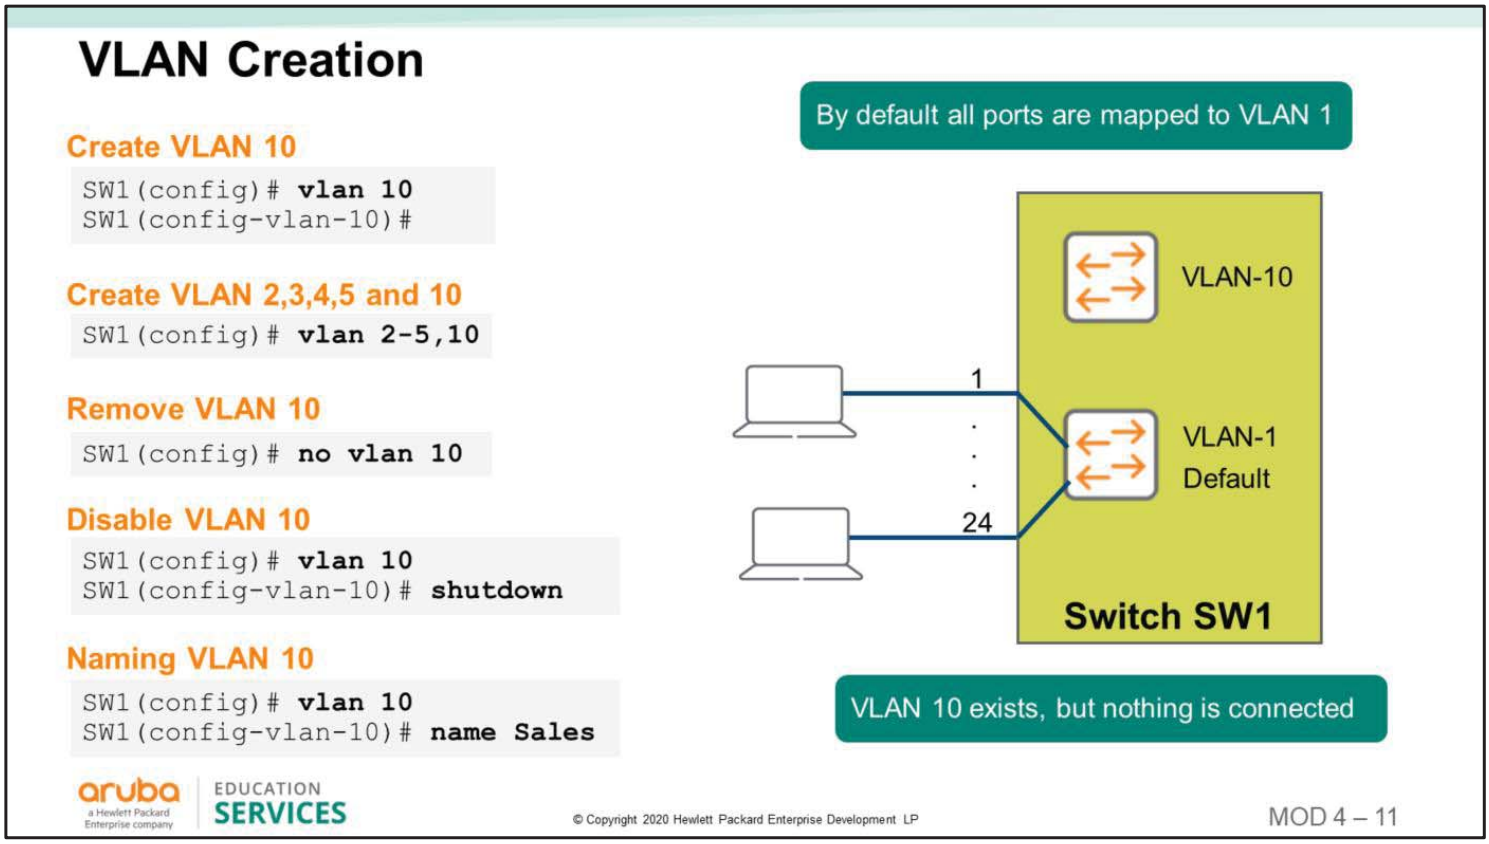

ArubaOS-CX refers to VLANs by their VLAN-ID - a number between 1 and 4094.
|In ArubaOS-CX VLAN 1 is created by default and cannot be removed. By default, all ports are members of this VLAN. This is a common default for many switches.

The vlan command creates a VLAN, which is enabled by default.
|In the figure's top example, you create VLAN 10.
|You can also create a series of VLANs with a single command, as shown in the next example.

|To remove the VLAN, use the command no vlan 10.

| Or, instead of deleting it, you could use the shutdown command to disable the VLAN.

|It is also a good idea to name the VLAN, as shown in the figure, with the name command.

Remember, when you use the command vlan 10 to create a VLAN, it is as if you have created a virtual switch inside of the actual, physical switch.
|Although it exists, no ports have yet been defined as members of this VLAN. No devices are yet attached to this virtual switch, as shown in the figure.

The next thing you probably want to do is associate or "map" physical interfaces as members of the VLAN

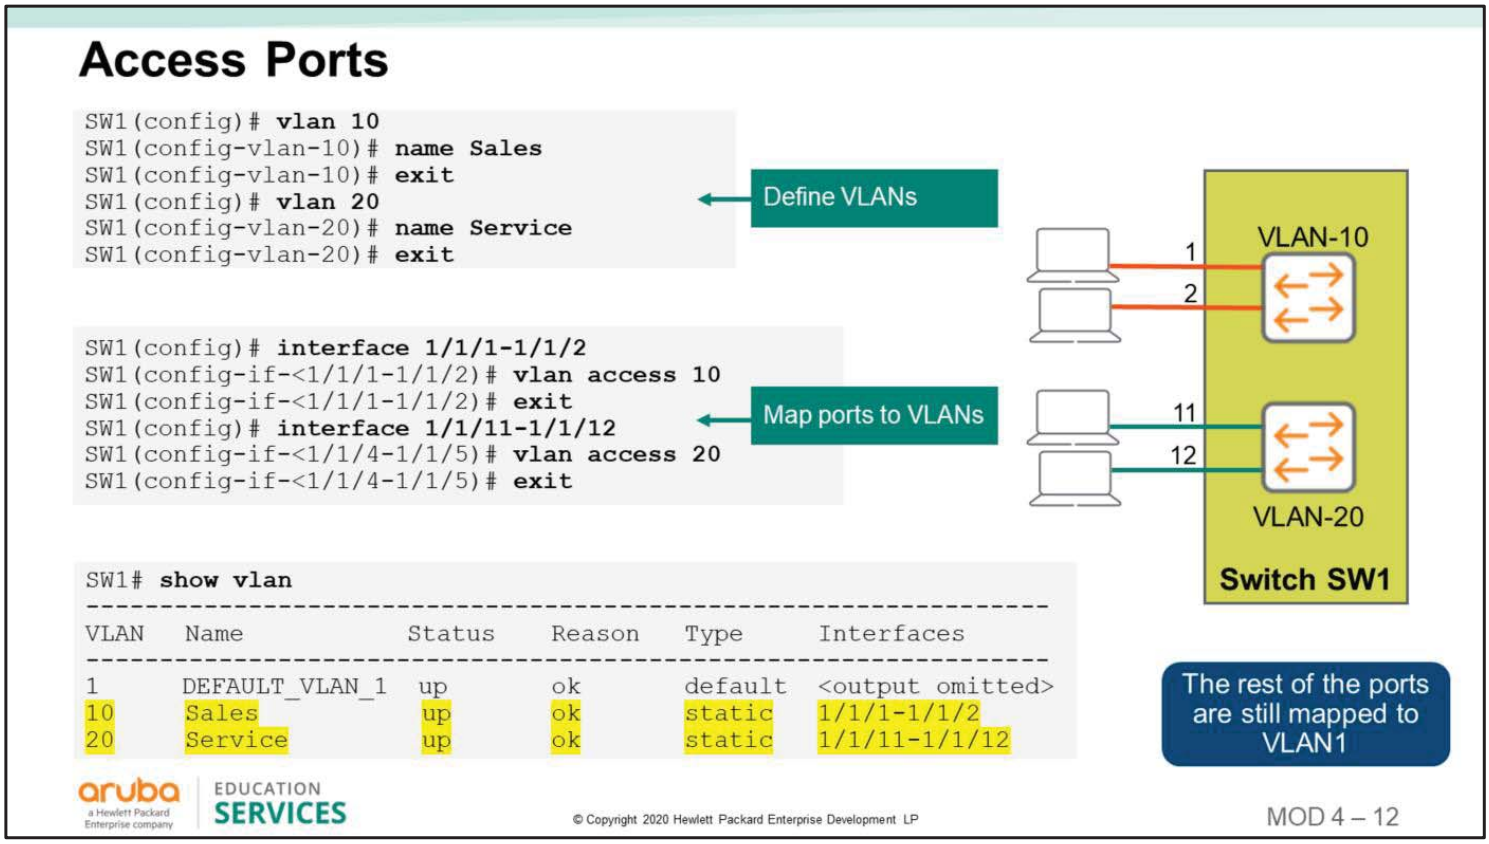

The figure shows SW1, with four connected devices, all in VLAN 1 by default.
|You define VLAN 10, name Sales.
|You define VLAN 20, name Service.

Now you need to assign or “map” ports to these VLANs.
|From the global configuration context, you choose to configure a range of interfaces at the same time – ports 1/1/1 through 1/1/2. Then you assign these ports to VLAN 10 as shown.
|You repeat these steps for ports 1/1/11 – 12, as shown.

Understand that only one VLAN ID can be assigned to an interface – interface 1/1/1 can’t be a member of both VLAN 10 and VLAN 20 at the same time. This would be like you trying to be in the sales room for a meeting and in the engineering room for another meeting at the same time.

|Use the show vlan command to verify VLAN creation and mapping.

ArubaOS-CX 6300 series interfaces are considered Layer 2 interfaces by default. In the rest of the family switches of ArubaOS-CX, interfaces are by default Layer 3. To convert these interfaces to Layer 2 mode, use the command “no routing”.

Question #2

VLANs

What are the benefits of creating VLANs?

a. Simplify the network deployment
b. Minimize the need for routing
c. Smaller broadcast domains can improve performance
d. MAC address tables are smaller
e. Separate VLANs mitigate risk

Knowledge Check

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 13

Correct answers:
c.
d.

802.1Q

Tagging, Trunk ports

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

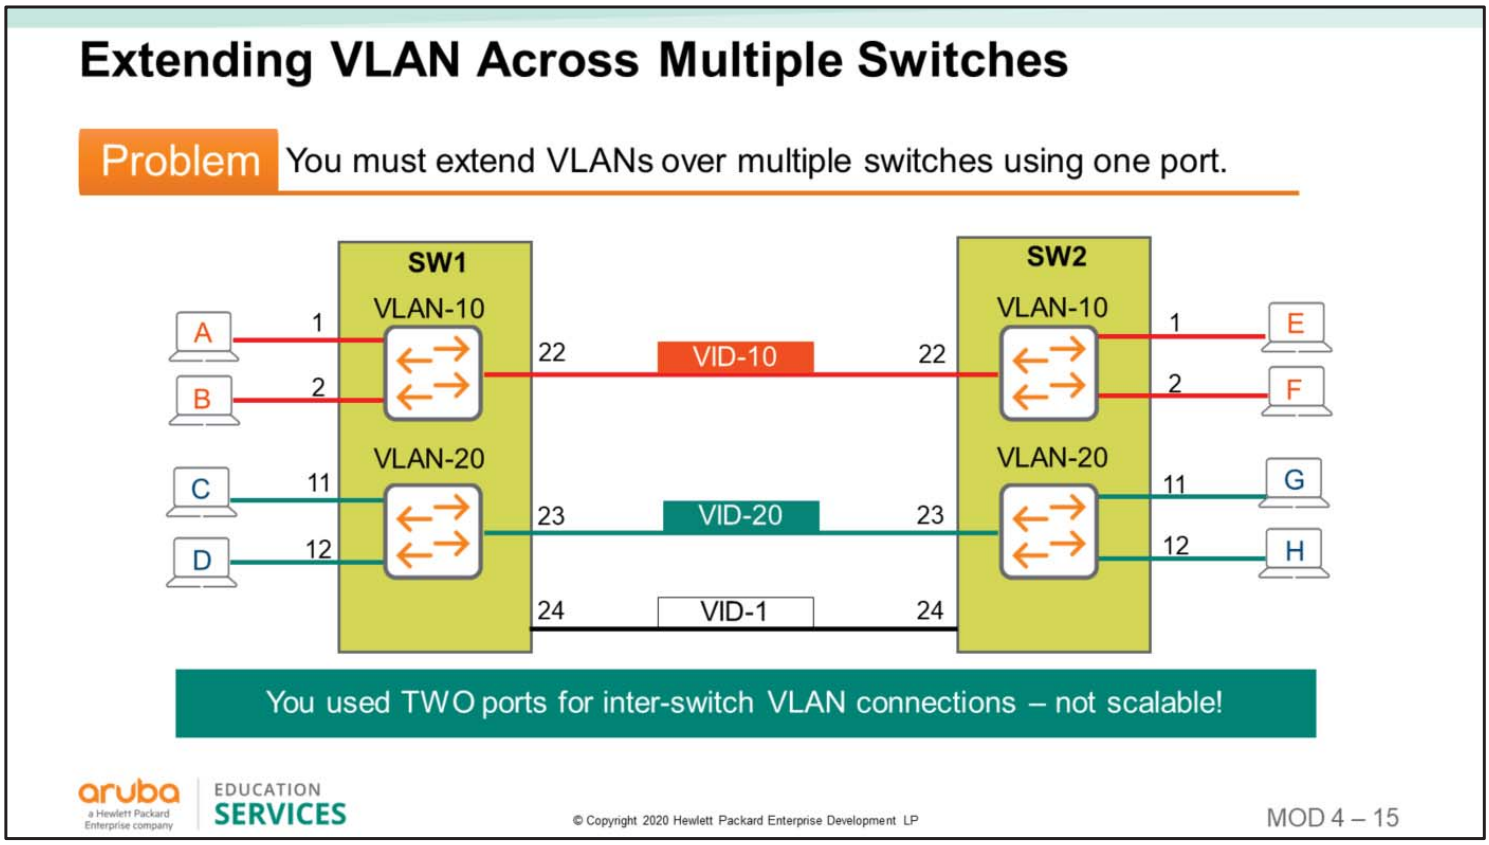

The figure shows two switches, each with ports in VLANs 10, and 20.
|Now you want to connect these two switches, so you connect port 24 of SW1 to port 24 of SW2. What happens? Will all devices be able to communicate?

Recall that by default ports are members of VLAN1, including port 24 on both switches. This means that port 24 is not connected to VLANs 10 and 20, and so cannot carry that traffic. Effectively, you still do not have connectivity between the switches for VLANs 10 and 20.How can you fix this?

One way is to connect two more pairs of ports.
|Map port 22 on both switches to VLAN 10, and port 23 as to VLAN 20. Now all VLANs can communicate. A broadcast from host A is flooded out all ports in VLAN 10 (except ingress port), and so all members of VLAN 10 receive the broadcast. This is good, but there is a problem.
|You used up 2 physical ports on each switch. What if you had 100 VLANs? We'd need 100 physical switch ports just for interconnects – that is not feasible.

We need a way to use one single physical port to connect multiple VLANs.

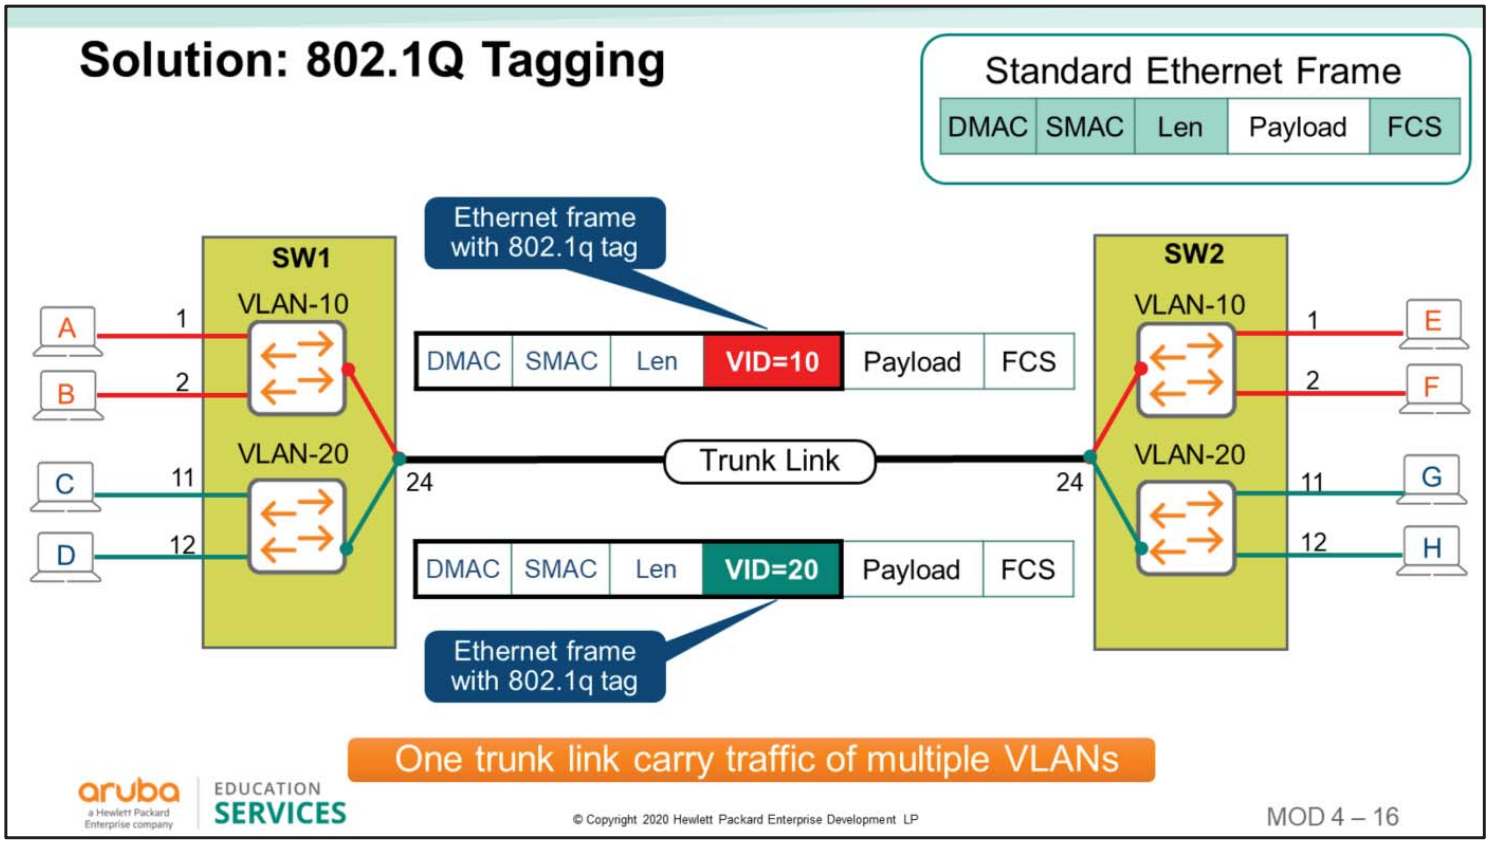

Look at the standard Ethernet frame shown in the diagram. Suppose that this was a broadcast frame from host A, flooded out all ports in VLAN 10, including port 24. There is no field in this frame to inform receiving switches about the intended destination VLAN. If SW2 could talk, it might ask, “What do I do with this? Is this for VLAN 10? VLAN 20? VLAN1?

We need a tagging mechanism, as defined in the IEEE 802.1Q standard. This standard adds an additional field between the length and payload fields – the 802.1Q Tag field. |The most important part of this field is the VLAN tag or VLAN-ID or simply VID. Before SW1 floods this broadcast frame across this specially-defined “trunk port”, it adds this tag to the frame.

Thus, SW2 receives the frame, looks at the tag and knows, “Ahhh. This frame is for VLAN 10”. It strips off the special tag, and forwards a standard Ethernet frame out all ports in VLAN 10 – to hosts E and F.

|Next host G sends broadcast. SW2 floods a standard Ethernet frame out its local ports. Before flooding the frame out port 24 , it’s a tag, “This frame is for VLAN 20”. SW1 receives this frame, sees that it is for VLAN 20, strips off the tag, and forwards a standard Ethernet frame out all ports in VLAN 20 (11 and 12).

Thus, you can now extend the VLAN concept across multiple switches, using a single physical port.

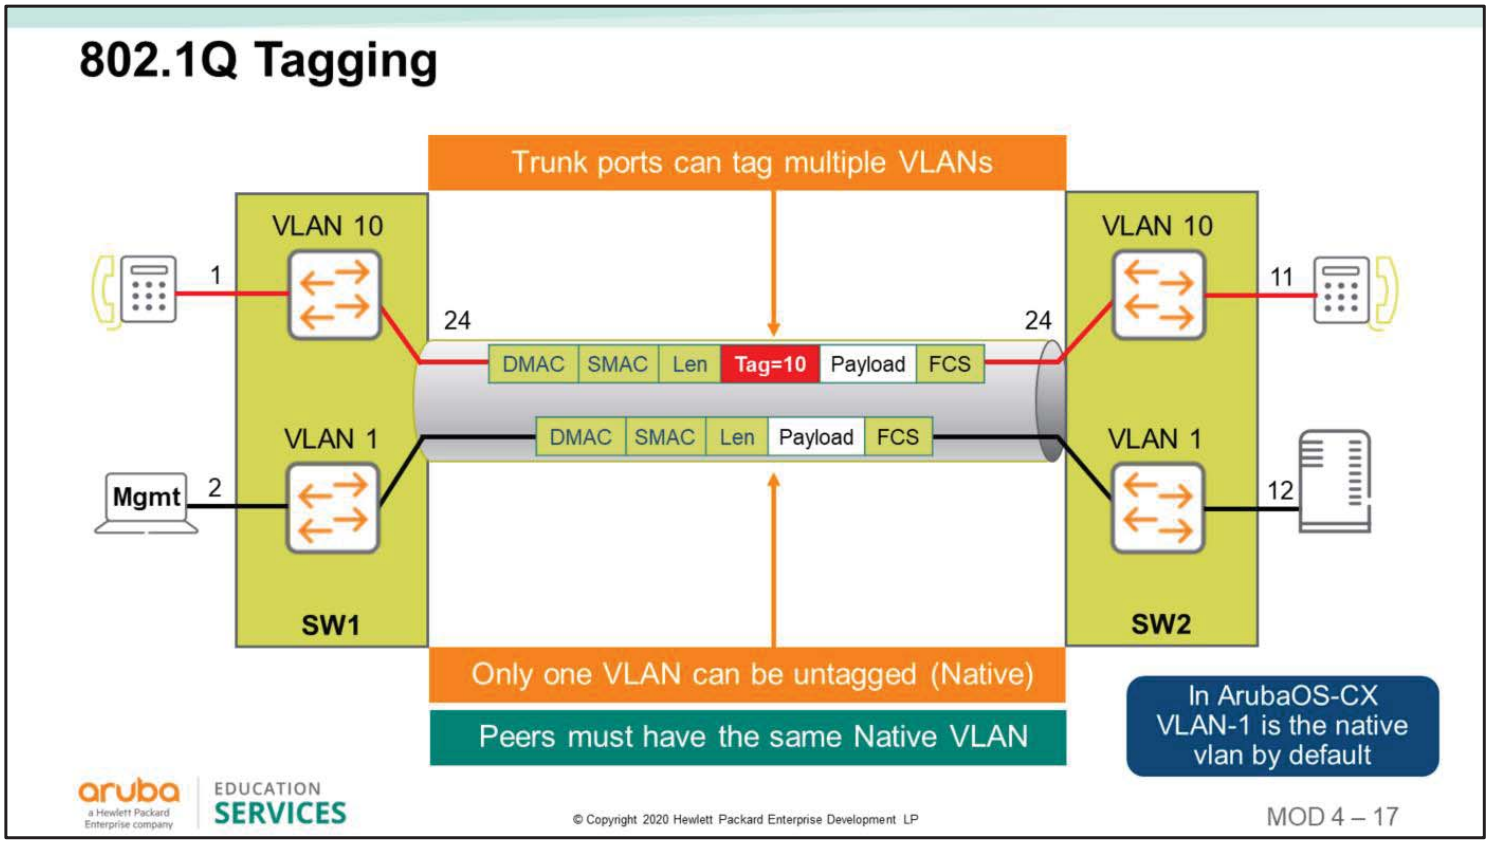

|Note. The VLAN-ID field is 12-bit long, this length allows a switch to carry traffic up to 4094 VLANs.

|The 802.1q standard allows administrators to select a single VLAN on which no tag or VID is included in the Layer 2 header. This means that a standard, untagged Ethernet frame is send out of the port. This VLAN is known as the Native VLAN or Untagged VLAN. By default in ArubaOS-CX VLAN 1 is considered the native VLAN. You can modify this by using vlan trunk native <vlan-id> command.

|To avoid inconsistencies Aruba recommends you use the same Native VLAN in both peers.

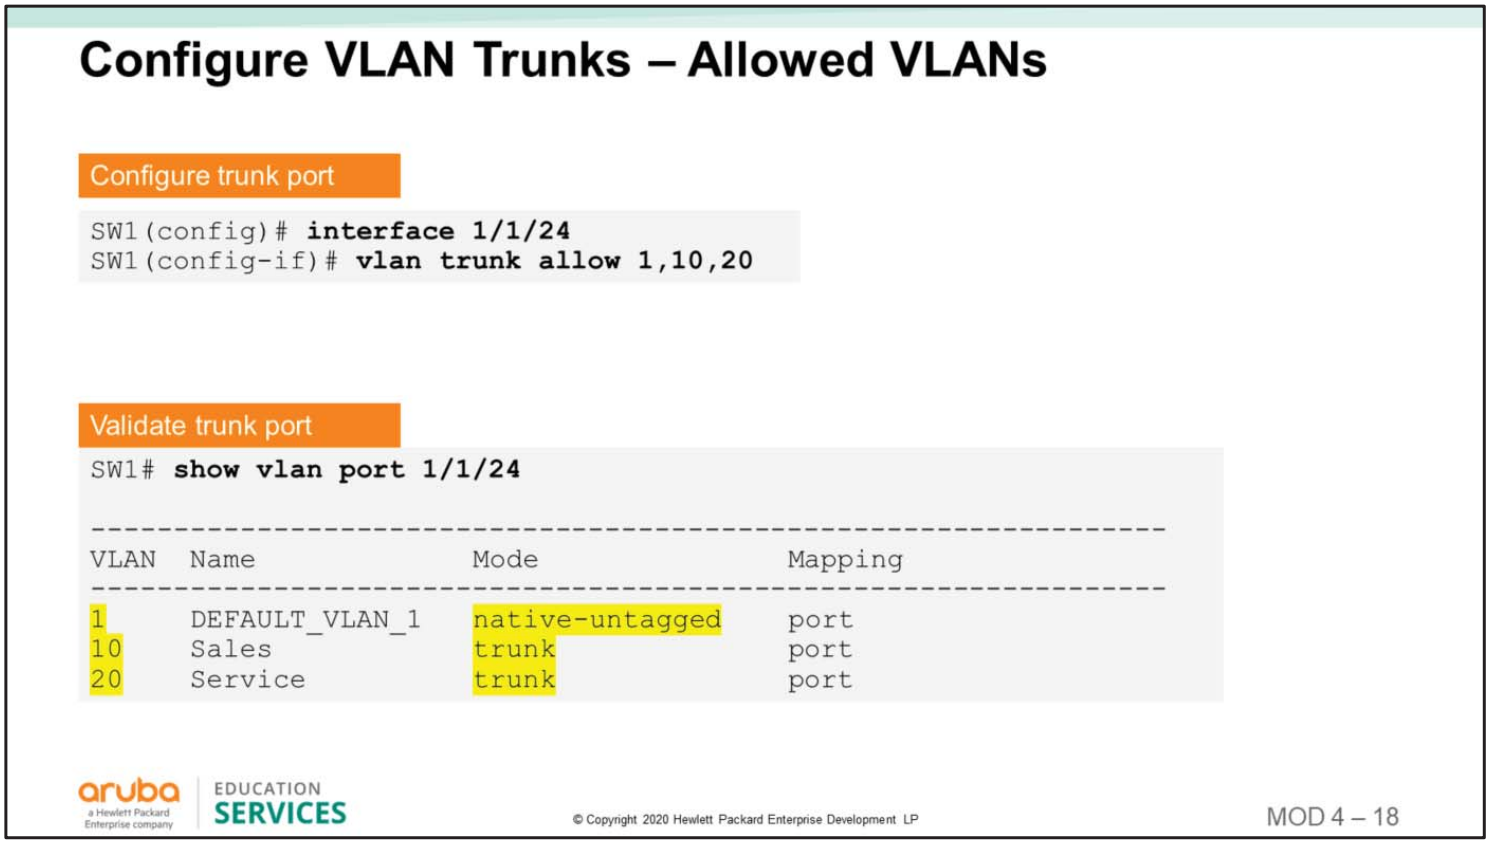

In ArubaOS-CX, use the command vlan trunk allowed to allow a VLAN-ID to traverse a trunk interface. Multiple VLAN-IDs can be assigned to a trunk interface. In the figure, interface 1/1/25 is configured as a trunk link, and only allows VLANs 1, 10, and 20. Other VLANs shall not traverse this link.

Note. This example is focused on Trunk port configuration. It is assumed that VLANs are already configured, and access ports are mapped to them.

|To verify the Trunk interface you can use the show vlan port command

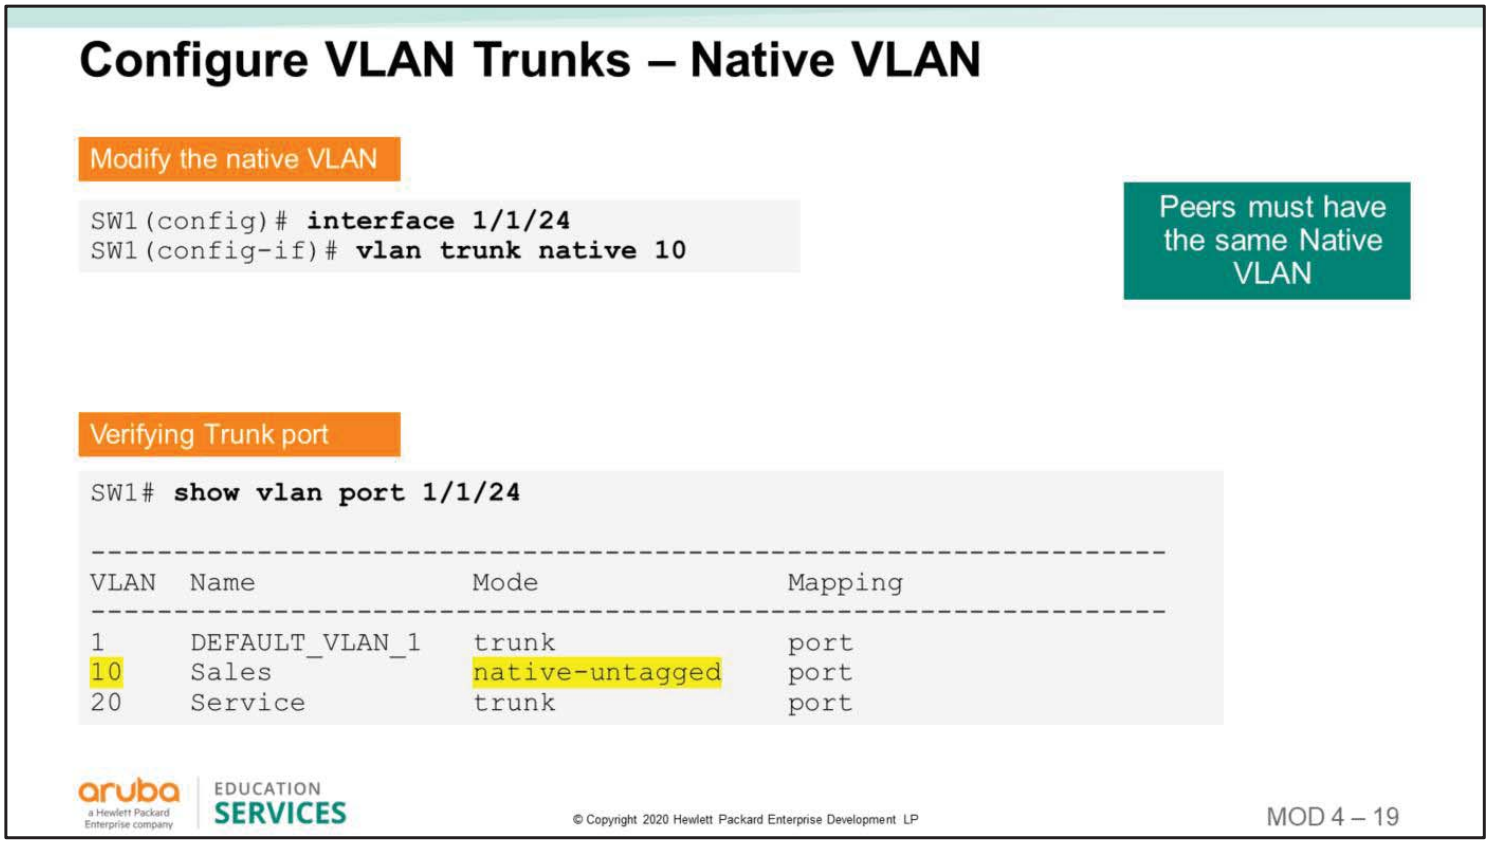

Recall that by default, all VLAN traffic is tagged to traverse a trunk link except the native VLAN, which defaults to VLAN 1. However, you can change the native VLAN as desired. As in the example, use the vlan trunk native command to change the untagged VLAN - to 10 in this case.

This configuration must match on both switches, so be sure to apply this command to the appropriate port of the attached switch.

Note: This is a good opportunity to remember earlier lessons. You can use the “show lldp neighbor-info” command to see which switch is connected to port 1/1/24 of the switch you are currently configuring.

|Verify your efforts with the command show vlan port.

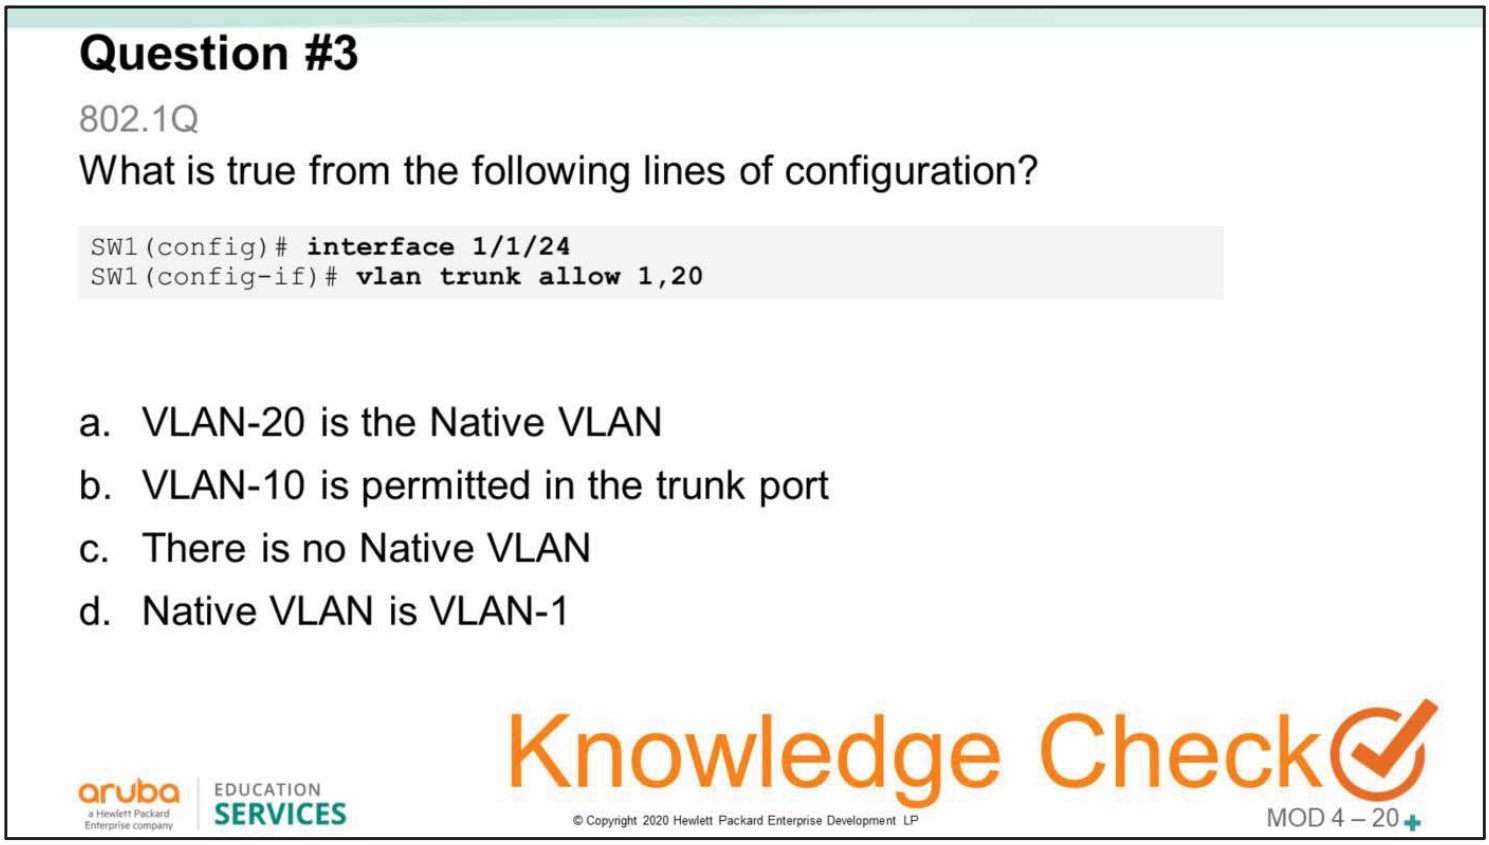

From the configuration of the port you can conclude the following:

VLAN 1 and 20 are the only ones permitted, this means that any other VLAN-ID cannot use the port 1/1/24.
The default Native VLAN is VLAN-1 in ArubaOS-CX, there is no need to explicitly specify this information.

Correct answer: D

Forwarding tables

MAC and ARP tables

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

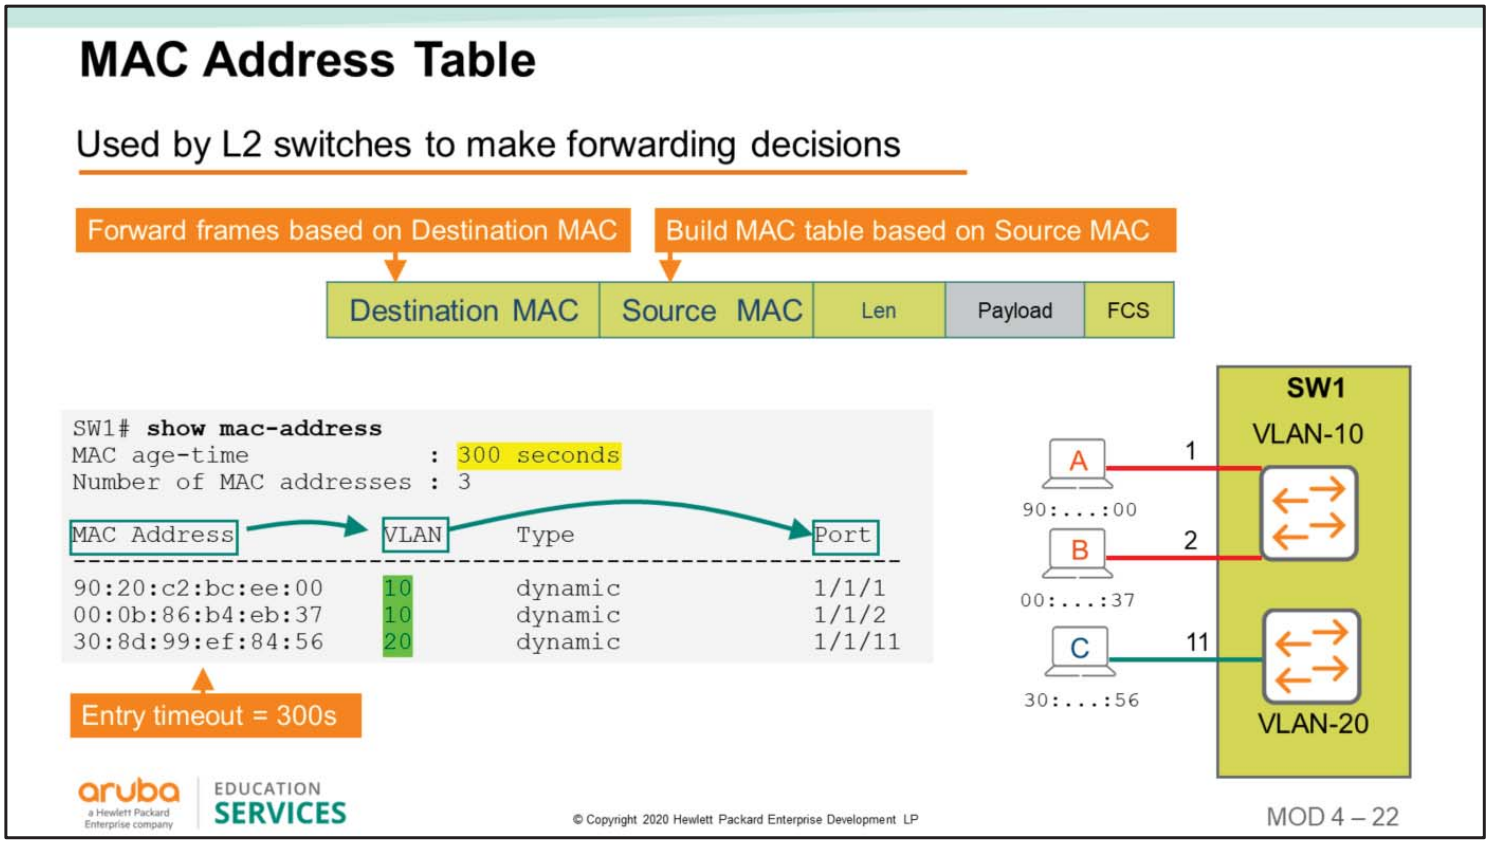

Layer 2 switches use the MAC address table to make forwarding decisions. The switch builds this table automatically, based on the source MAC address of the frames that it receives frames from connected devices. How?

Suppose that you power up switch SW1. Its MAC address table is blank. You see this here in the show mac-address output.
Then Host A transmits a frame to Host B. Switch 1 receives this frame with source MAC address = 90:…:00 on Port 1, defined as being a member of VLAN 10.
Thus, it adds the top entry to the MAC address table, as shown in the figure.
It maps the MAC address to its VLAN and port.

Suppose that at this point there are no other entries in the table. The switch does not yet know where to forward frames to destination 00:…:37, so it floods this frame out all ports in VLAN 10. All VLAN 10 hosts receive this frame and see the destination MAC address – “This is not for me”, and they discard the frame. Except for Host B, “This is my MAC address, I will respond”.
And so Host B responds to Host A with source MAC = 00:…:37, and destination MAC = 90:…:00.

Now the switch adds the second entry you see in the table MAC 00:0b:86:b4:eb:37 is in VLAN 10, connected to port 2.

A similar process will happen with host C on port 11.

|So, switches automatically build the MAC table based on source MAC addresses, and forward frames based on destination MAC addresses.

|By default, table entries are maintained for 300 seconds (5 minutes). You can verify the MAC address table in ArubaOS-CX using the show mac-address command.

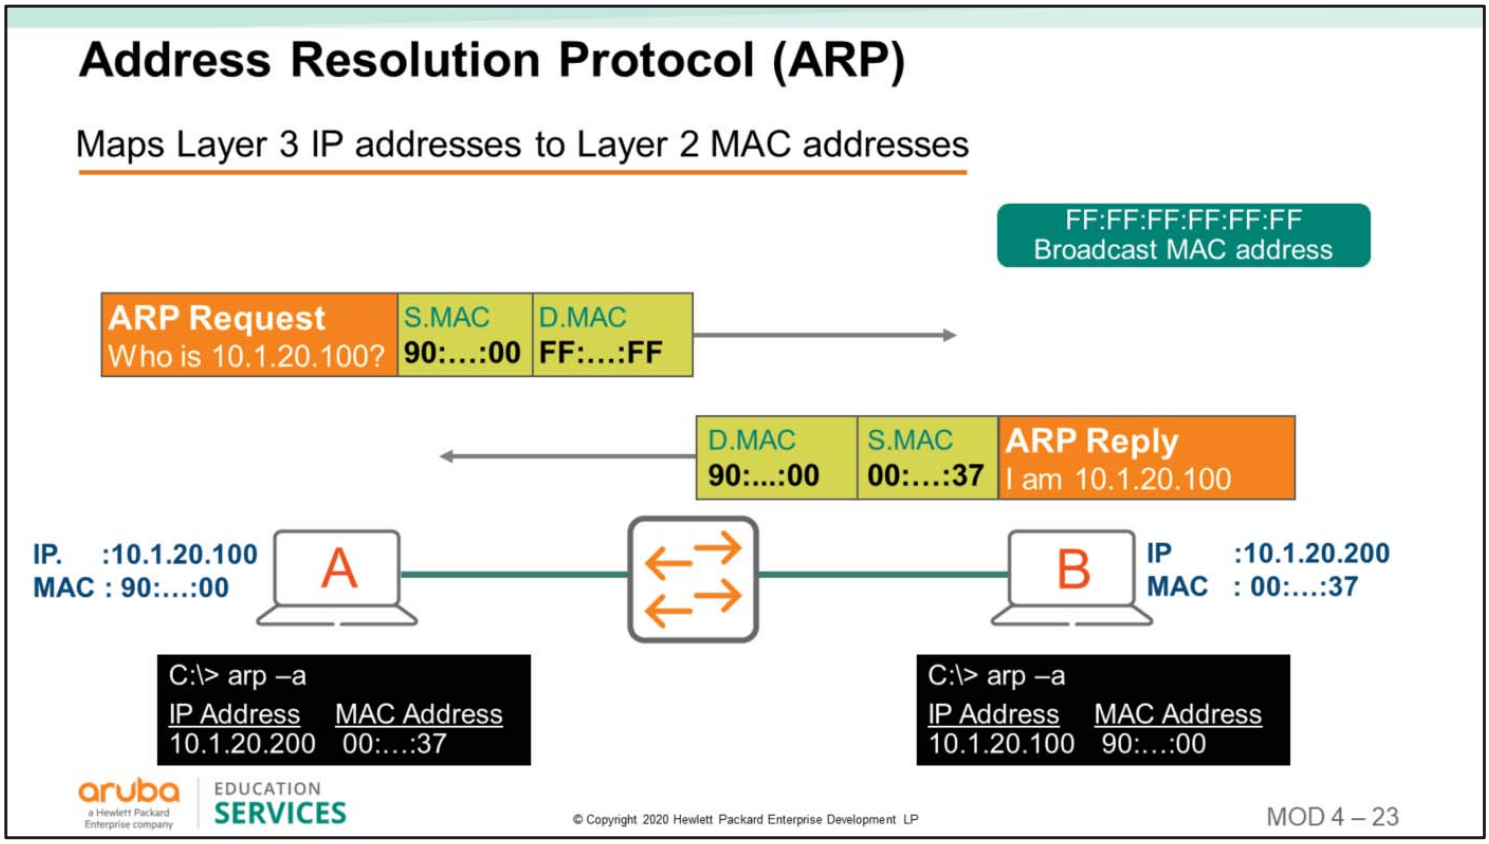

Imagine that you are in a conference room, leading a discussion with people that you have just met. You don't yet know everyone's name. Someone enters the room and hands you a note that says, "Important message for Alvin Rogers".

You do not know which person is Alvin, so you say, "Excuse me everyone. Who is Alvin Rogers?". Everyone hears your request, but only Alvin responds, "I'm Alvin". You thus can associate the name Alvin to an actual person, seated at a particular chair in the room.

This is very much like the Address Resolution Protocol (ARP), which maps Layer 3 IP addresses to Layer 2 MAC addresses.

In the figure, Host A knows that it must communicate with the host at IP address 10.1.20.200. To do this it must build an Ethernet frame, which requires its own source MAC address (90::00), and Host B's destination MAC address, which it does not yet know.
|To learn it, Host A broadcasts an ARP request - destination MAC address ff:ff:ff:ff:ff:ff. (In the IP header, the destination IP address is 255.255.255.255). You know that switches forward broadcasts out all ports in the same VLAN (except for ingress port). Thus all hosts in the VLAN receive this frame and ignore it, "I'm not 10.1.20.200".
|Except for Host B, which respond with an ARP reply, unicast to Host A's MAC address, "I am 10.1.20.100, and my MAC address is 00::37".

|Host A receives this reply and creates an entry in its ARP table, sometimes called an ARP cache. It maps 10.1.20.100 to 00:0b:86:b4:eb:37. Thus, the next time Host A must communicate with Host B, it need not use ARP.

|Also know that when Host B received the ARP request from host A, it learned about Host

A‘s MAC address and IP address. Thus, Host B creates an entry in its ARP table, mapping 10.1.20.100 to MAC address 90:20:c2:bc:ee:00.

On a Windows PC’s command prompt, use the command arp –a to see which IP addresses have been resolved to a MAC address.

Question #4

MAC address and ARP tables

Which options below accurately describe MAC address and ARP tables?

a. Switches automatically build the MAC table based on destination MAC addresses.

b. MAC address tables are build based on source IP addresses.

c. The ARP reply packet uses a broadcasting mechanism.

d. The ARP table maps IP addresses to MAC addresses.

e. Switches use the MAC address table to properly forward frames.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

Correct Answers:
d.
e.

Frame Delivery

Layer 2 switch communication

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

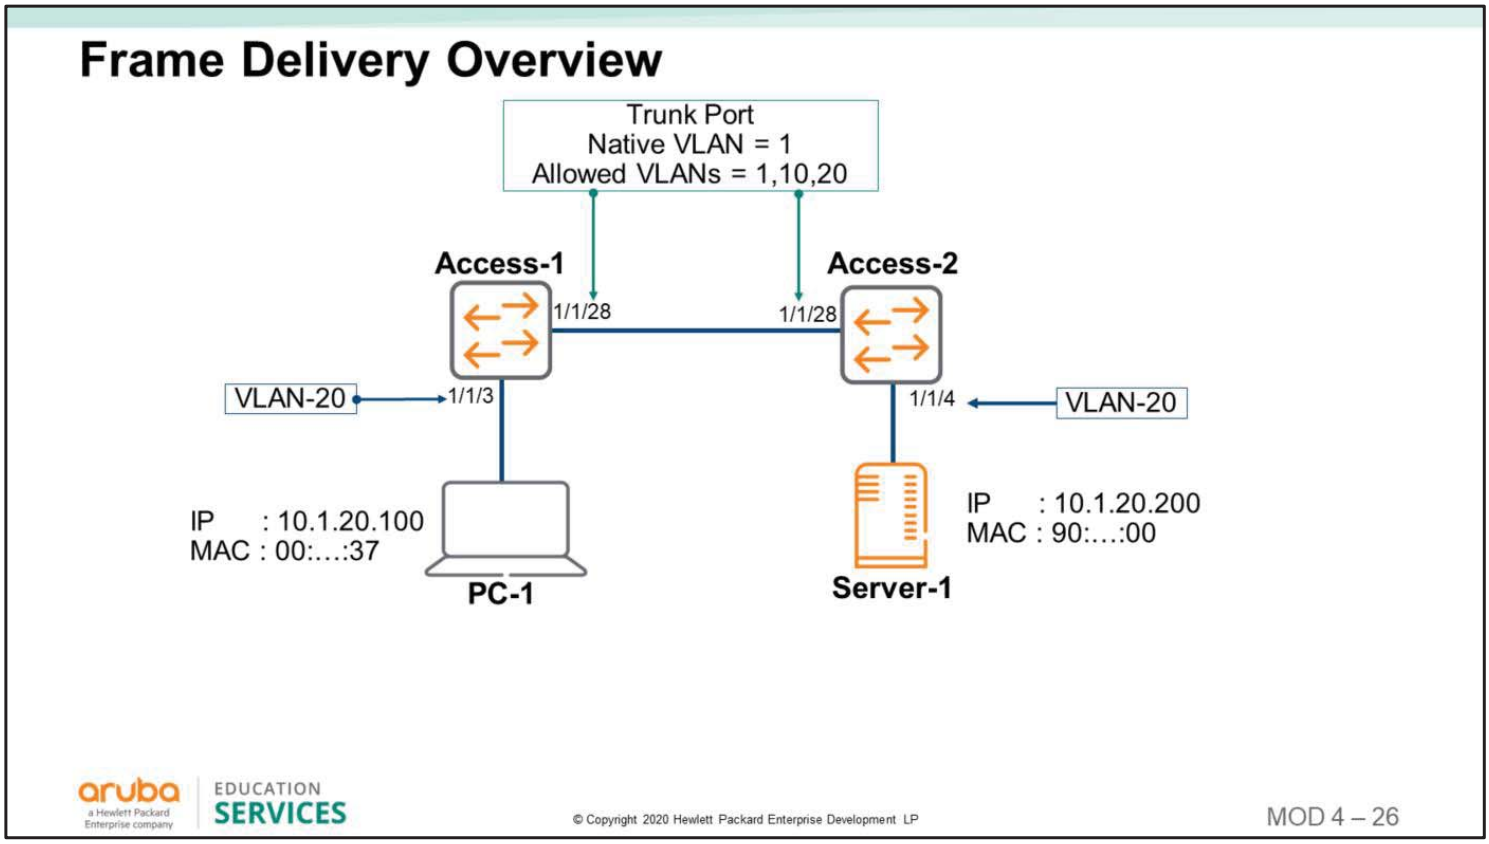

In this section you will learn how two devices in the same VLAN communicate across multiple switches. The scenario is as follows:

PC-1 and Server1 are correctly configured with an IP address.
Both switches are configured with 3 VLANs: 1, 10 and 20.
The switch-to-switch links are configured as trunk ports allowing VLANs 1, 10 and 20, and the Native VLAN = 1.
Port 1/1/3 on Access-1 switch and Port 1/1/4 on Access-2 switch are mapped to VLAN=20
PC-1 has not initiated any previous communication to Server-1.

In this scenario, the user on PC-1 wants to establish an FTP session to download a file from PC-2.

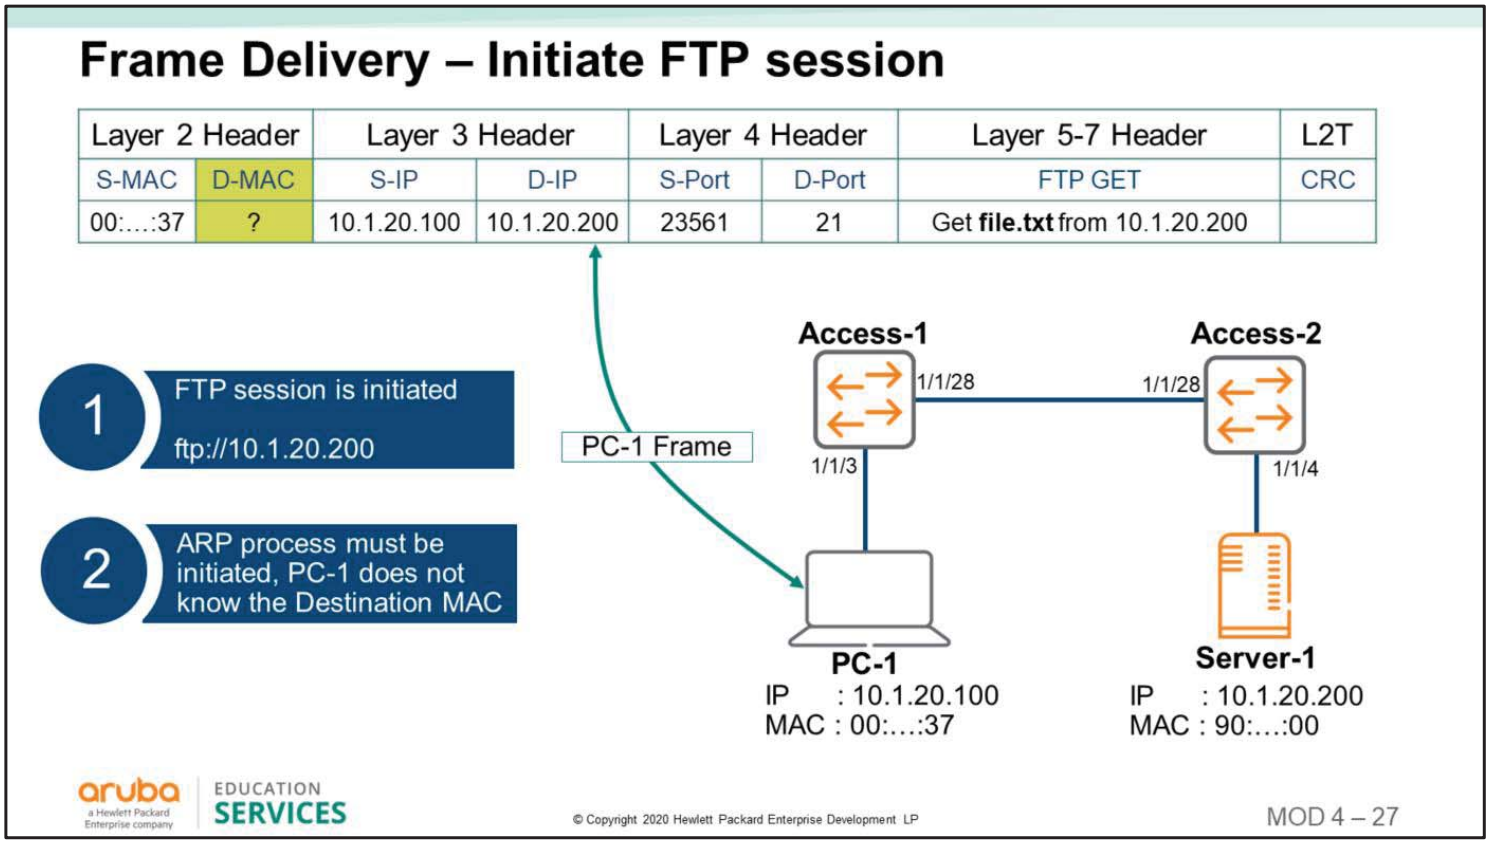

1. The user on PC-1 opens a browser and types the IP address of Server1 - ftp://10.1.20.200

2. Notice that even if PC-1 knows the Destination IP address it does not know the MAC address associated to this IP. Therefore an ARP process must be initiated.
Also notice that PC-1 has the rest of the information to build the Layer 2 frame including Layer 3 to Layer 7.

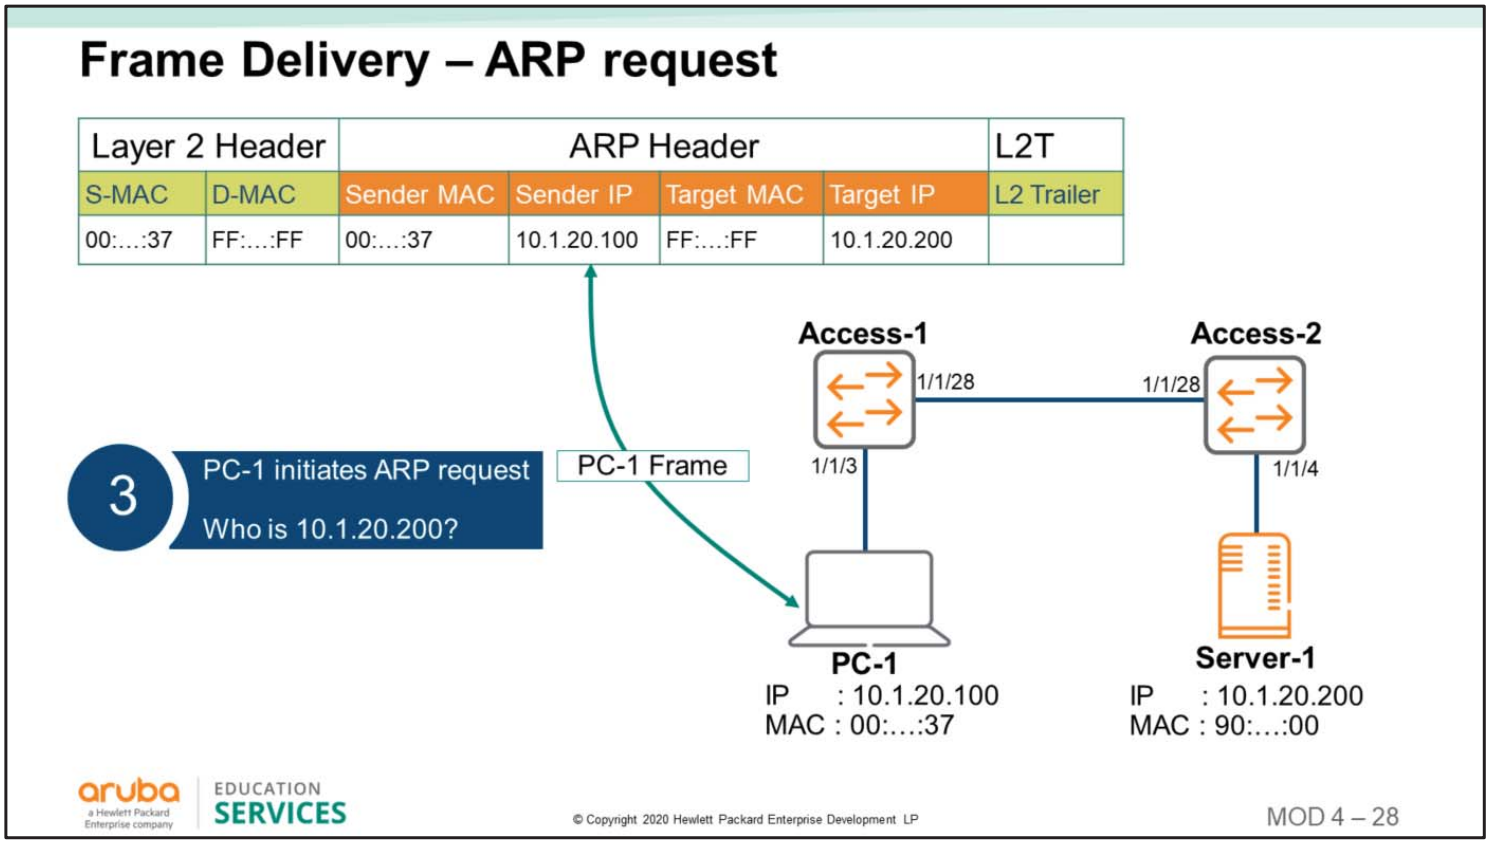

3. PC-1 generates an ARP request. Notice that the ARP request has a Layer 2 broadcast destination MAC address (FF:FF:FF:FF:FF:FF)

Frame Delivery – Flood ARP Broadcast

Layer 2 Header		ARP Header				L2T
S-MAC	D-MAC	Sender MAC	Sender IP	Target MAC	Target IP	L2 Trailer
00:…:37	FF:…:FF	00:…:37	10.1.20.100	FF:…:FF	10.1.20.200	

4 Access-1 adds entry to the MAC address table

5 Destination is broadcast, Access-1 floods the frame on all VLAN 20 ports

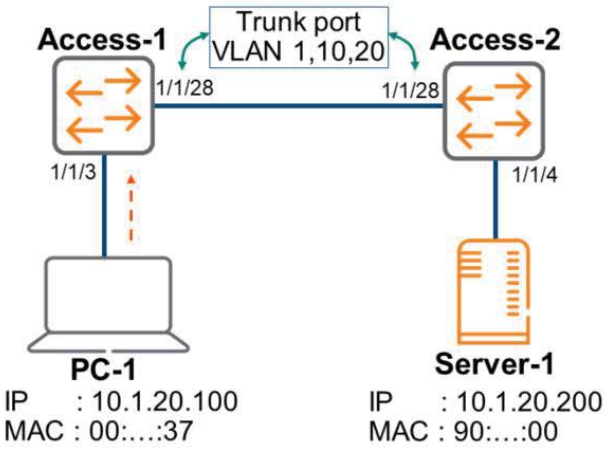

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 29

4. The frame is received by Access-1, which creates an entry in the MAC address table with PC-1 MAC on port 1/1/3.
5. Because the ARP request is a broadcast, the switch floods the frame out all the ports in VLAN 20. In this example, the frame is sent to Access-2 switch via port 1/1/28.

Frame Delivery – dot1q Tagging

Layer 2 Header			ARP Header				L2T
S-MAC	D-MAC	VLAN-ID	Sender MAC	Sender IP	Target MAC	Target IP	L2 Trailer
00:…:37	FF:…:FF	VLAN-ID = 20	00:…:37	10.1.20.100	FF:…:FF	10.1.20.200	

6 Access-1 adds VLAN-ID and sends the frame to Access-2

7 Access-2 creates an entry in the MAC address table

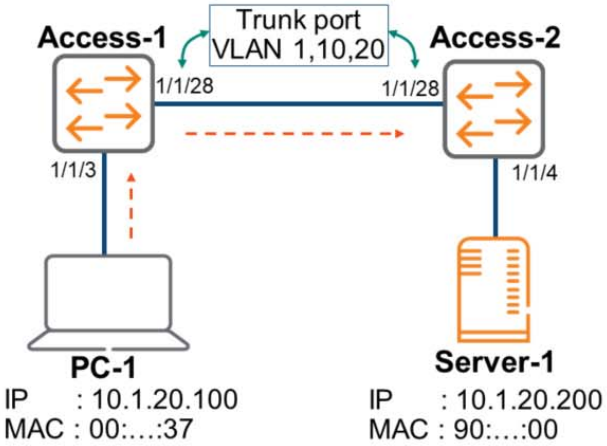

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 30

6. Port 1/1/28 is configured as a Trunk port, so the frame uses a dot1q tag of VLAN-ID = 20.
7. The frame is received by Access-1. This switch will create an entry in the MAC table.

Frame Delivery – Forward Untagged Frame

Layer 2 Header		ARP Header				L2T
S-MAC	D-MAC	Sender MAC	Sender IP	Target MAC	Target IP	L2 Trailer
00:…:37	FF:…:FF	00:…:37	10.1.20.100	FF:…:FF	10.1.20.200	

8 Access-2 forwards the frame to Server-1 using port 1/1/4

Frame to Server-1 will not include VLAN-ID

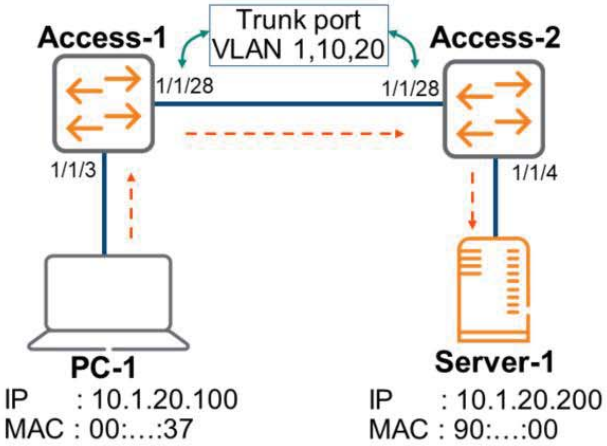

aruba a Hewlett Packard Enterprise company EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 31

8. Access-2 forwards the frame to all the ports mapped to VLAN-20. In this case the frame is sent out interface 1/1/4. This interface is not a trunk port which implies that the frame will not include an 802.1q tag with its VLAN-ID field.

Frame Delivery – ARP Reply

Layer 2 Header		ARP Header				L2T
SMAC	DMAC	Sender MAC	Sender IP	Target MAC	Target IP	L2 Trailer
90:…:00	00:…:37	00:…:37	10.1.20.100	90:…:00	10.1.20.200	

9 Server-1 responds with an ARP reply

This frame is unicast

10 Access- 2 learns the Server-1 MAC address

11 Access- 2 sends the frame to Access-1. This frame includes the VLAN-ID

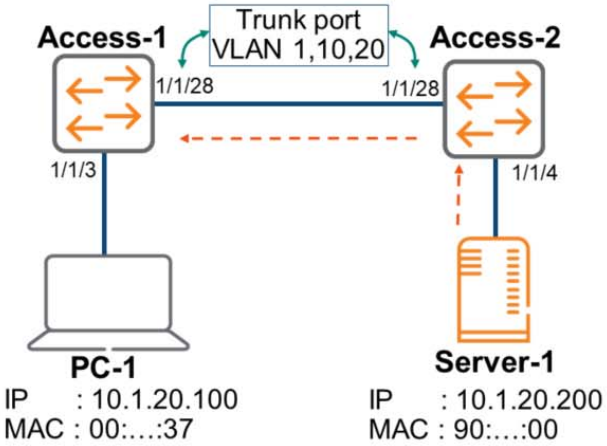

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 32

PC-2 receives the ARP request and generates an ARP reply. PC-2 knows the source MAC, so this is a unicast frame.

Access-2 processes the ARP reply message. The switch learns the Server-1 MAC address and adds the entry to the MAC address table.

Access-2 sends the ARP reply to Access-1. This frame includes the VLAN-ID and the frame is send out from interface 1/1/28.

Frame Delivery – PC Receives ARP Reply

Layer 2 Header		ARP Header				L2T
SMAC	DMAC	Sender MAC	Sender IP	Target MAC	Target IP	L2 Trailer
90:…:00	00:…:37	00:…:37	10.1.20.100	90:…:00	10.1.20.200	

12 Access-1 forwards the ARP reply to PC-1, frame is untagged

13 Access-1 adds an entry of Server-1 MAC address

14 PC-1 knows the Destination MAC
PC-1 can initiate the FTP session

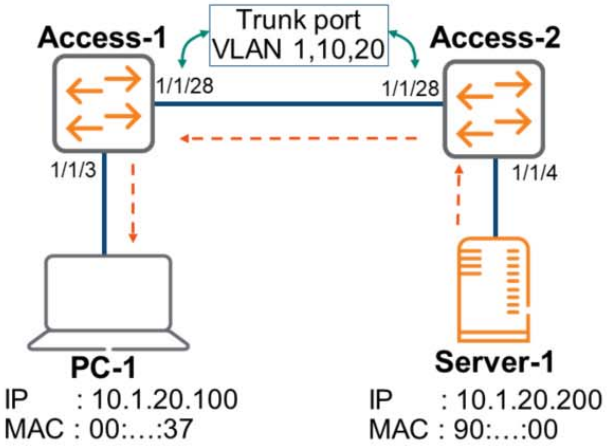

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 33

Access-1 forwards the ARP reply to PC-1 via port 1/1/3. This port is not trunk, so no VLAN tag is used.

The switch creates entries for PC-2 in the MAC and ARP tables.

Now that PC-1 knows PC-2's MAC address it can build the Ethernet frame and initiate the FTP communication to get the file.

Question #5

Frame Delivery

When does a switch add a VLAN tag to a frame?

a. When it forwards the frame to a server
b. When it forwards the frame to a PC or client
c. When it forwards the frame to another switch
d. When it accepts a frame from another switch
e. When it accepts a frame from a server

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 4 – 34

Correct answer:
c.

Lab Activity

4.1 Configure a VLAN
4.2 Add a second switch to the topology
4.3 Add a core switch to the topology

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

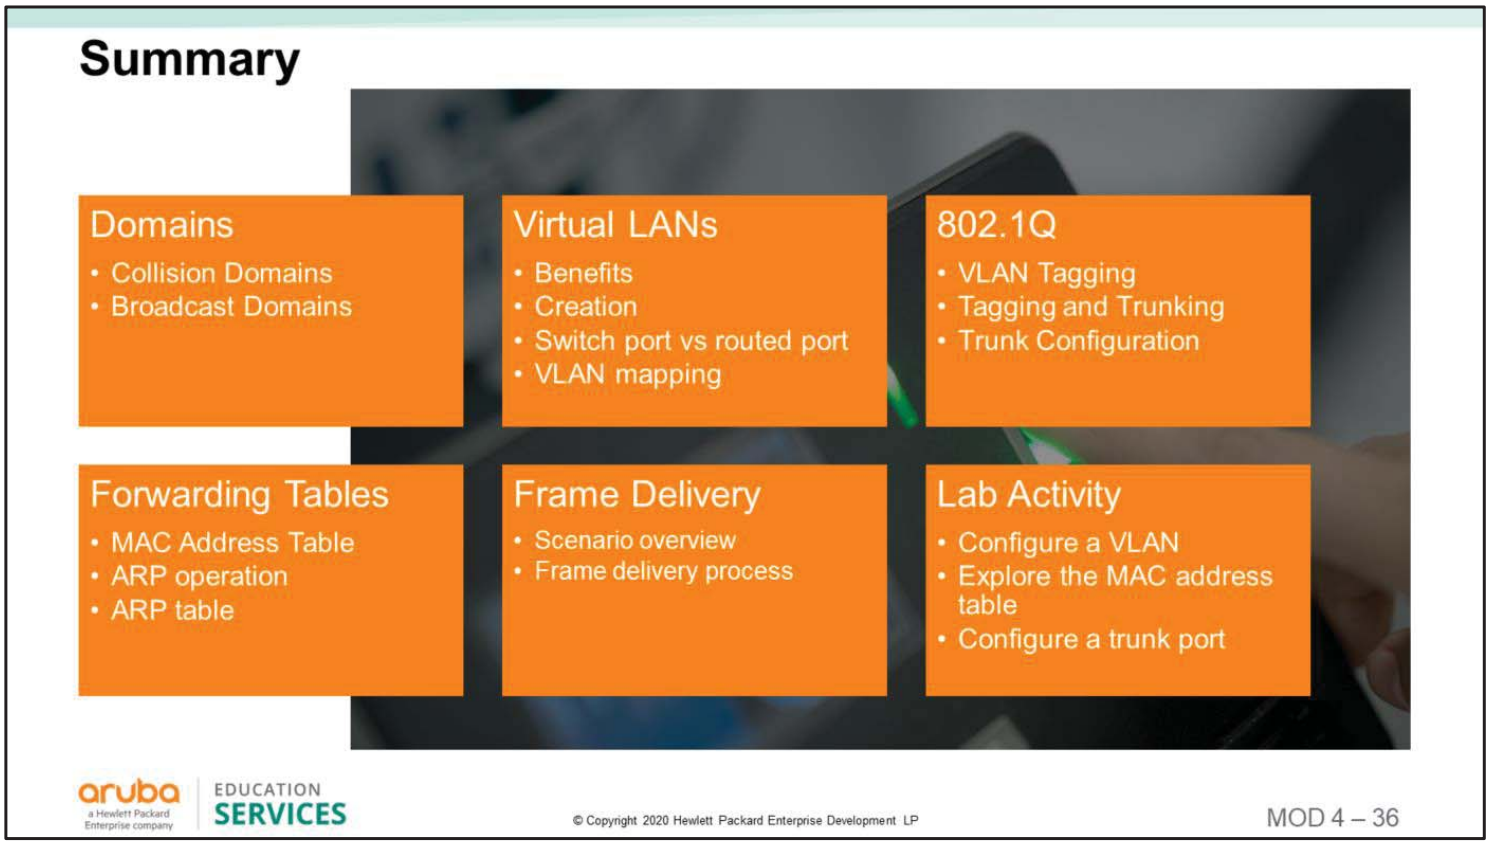

Congratulations! Now you have the essential knowledge and skills to understand and configure vital Layer 2 switch concepts, which will speed your journey toward mastery of Layer 2 network design, deployment, configuration, and troubleshooting.

You learned about Layer 2 collision domains and broadcast domains – important for both your theoretical knowledge, and your practical ability to improve network performance. Then you learned about VLANs, and how vital they are to creating scalable, flexible, and secure networks. Related concepts included switch ports, Switch Virtual Interfaces (SVI), and physical routed ports. You turned this theoretical knowledge to practical use by learning how to configure all of these entities.

You learned about some potential limitations of VLANs, and how they are eliminated with the 802.1Q trunking protocol. With this knowledge and configuration ability, you are able to extend the VLAN concept across multiple physical switches.

With such an ability to create larger, more scalable networks comes a need to understand forwarding tables, including the MAC address table that is built by Layer 2 switches, and the ARP cache, which is maintained by both end systems, Layer 2 switches, and Layer 3 routers.

Armed with all of this prerequisite knowledge, you explored a typical scenario of two devices communicating over a network. This united the earlier information from this module into a real-world scenario. You saw the interaction of Layer 2 frame addressing, VLANs, 802.1Q, the MAC address table, and the ARP cache.

aruba
a Hewlett Packard Enterprise company
Spanning Tree Protocol
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba | EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP

Objectives

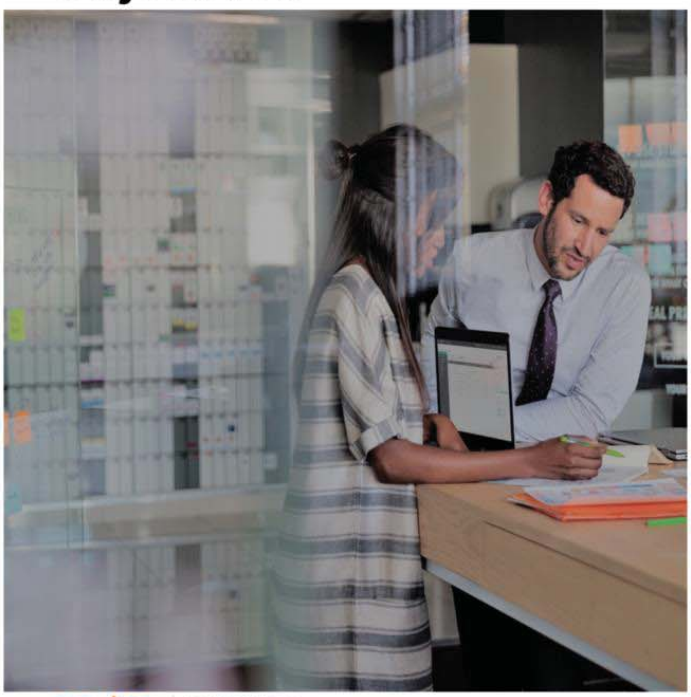

- Learning Objectives
  - Describe redundant network links
  - Describe Spanning Tree Protocols
  - Explain and Configure Rapid STP
    - Elements and operation
    - Ports and links
    - Proposals and Agreements
  - Describe and Configure MST

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 2

After completing this module, you will be familiar with the concepts and configuration related to various Spanning Tree Protocols (STP).

You will be able to describe how the original IEEE 802.1d protocol can resolve issues related to redundant network links, along with its performance limitations.

You be able to explain how the Rapid Spanning-Tree Protocol (RSTP) improves this performance in some detail, including RSTP elements and operation, ports and links, and proposals and agreements.

You will be able to describe and configure the Multiple Spanning-Tree protocol to gain additional improvements over RSTP.

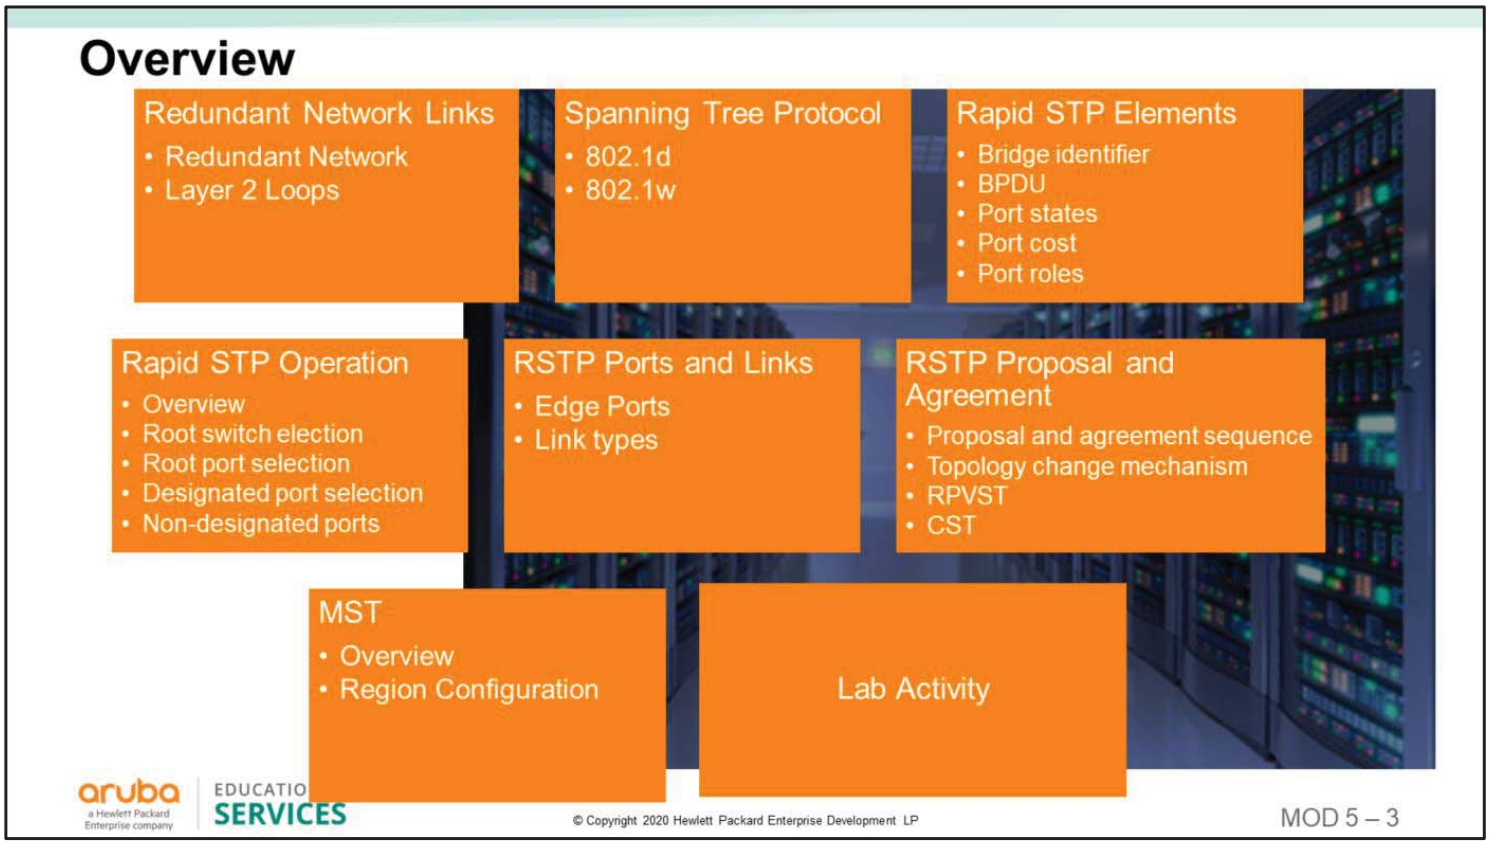

You are about to explore the most vital aspect of designing Layer 2 switched networks, especially as relates to resiliency, path optimization, and network efficiency – the Spanning-Tree protocols.

You begin by learning how single points of failures can be mitigated by connecting redundant switches and redundant links. However, this causes the network-breaking problem of Layer 2 loops – broadcast storms and MAC table instability.
You learn to solve these problems with the 802.1d Spanning-Tree Protocol (STP). But this is too slow for modern networks, so you upgrade to the newer IEEE 802.1w standard – the Rapid Spanning-Tree Protocol (RSTP). You learn about RSTP elements, and how they work together to create a functional system that provides layer 2 resiliency, while avoiding loops.

You dive deeper into RSTP operation. You explore how edge ports and link types, as well as the RSTP proposal and agreement process can be leveraged to further increase the efficiency and uptime of your systems.
As always, this new RSTP solution brings new problems and challenges. In large networks, this solution can waste bandwidth and CPU resources, potentially preventing your network from reaching its true potential.

You learn how to resolve this issue by deploying the IEEE 802.1s standard Multiple Spanning Tree (MST) solution.

Finally, you engage in some hands-on lab activities to solidify your knowledge.

Redundant Network

Redundant Networks and Layer 2 Loops

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

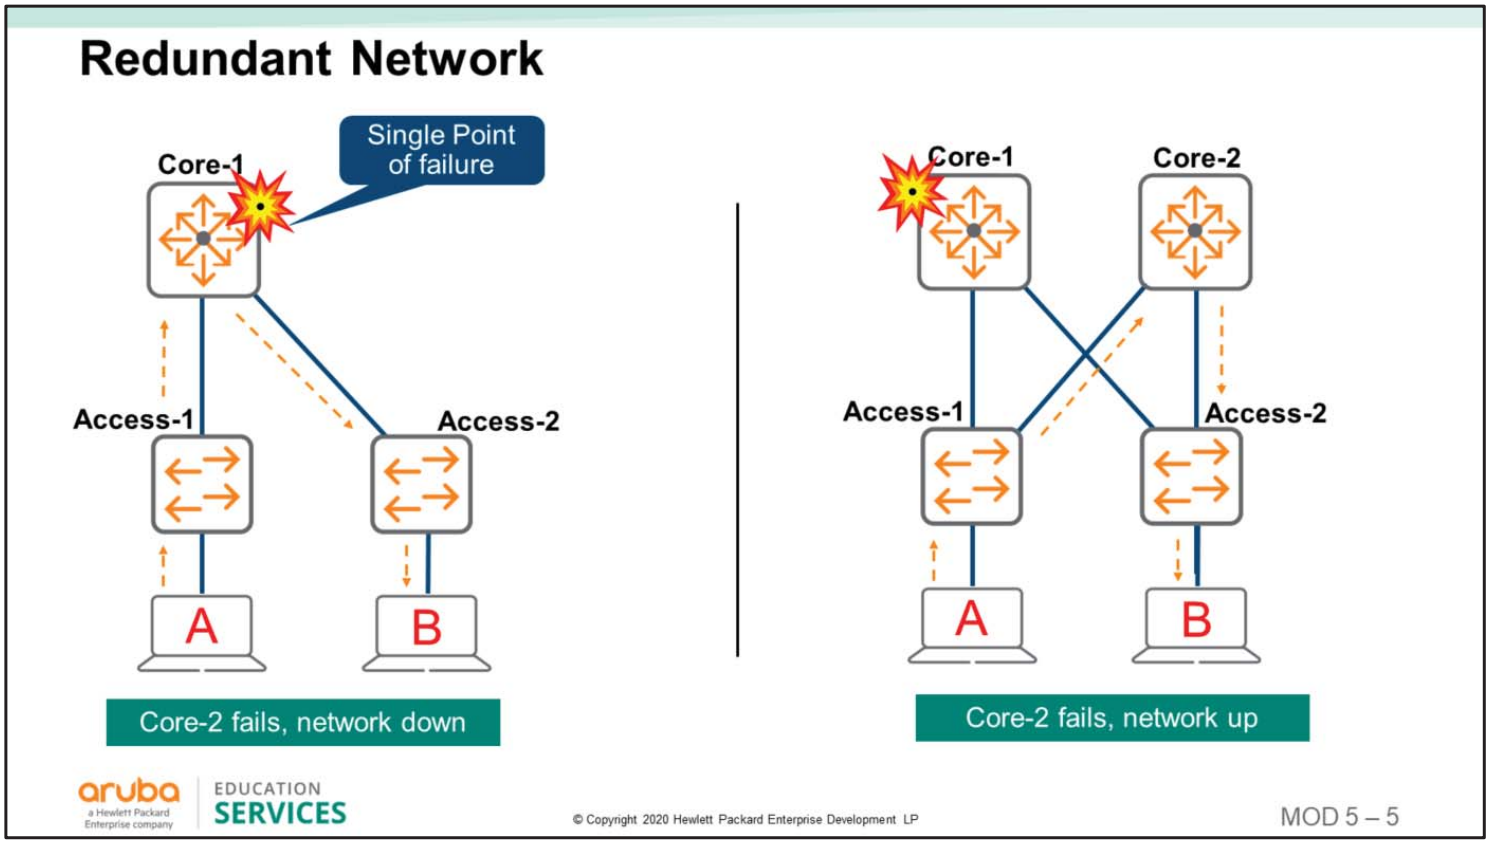

Network communications is often mission critical. Layer switch or link failures cannot be tolerated. For example, the left-hand example shows a single path to connect two Access switches using Core-1. This design has single points of failure. If Core-1 fails, the network is down. Host A and B cannot communicate.

One common way to mitigate this is by adding a redundant Core switch, as in the figure's right-hand example. In this scenario, if the Core-1 fails, the network remains viable. Hosts A and B can still communicate over the redundant link using Core-2

While this redundant link mitigates a single point of failure issue, it creates a new challenge.

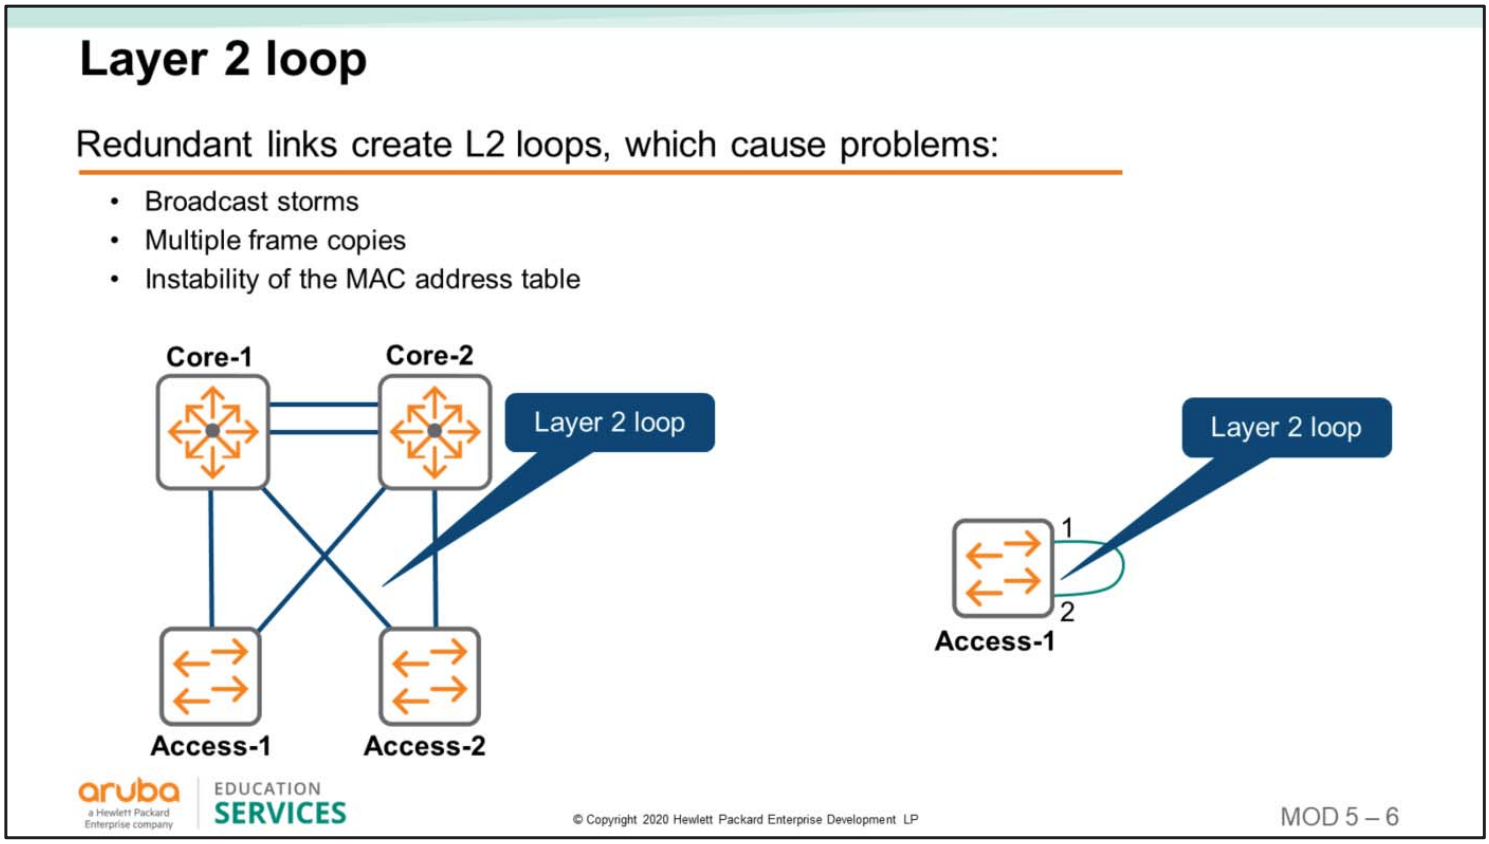

Connecting Layer 2 switches with redundant links creates Layer 2 loops. A loop is even created if you connect an Ethernet cable from one port on a switch to another port on that same switch. The slide shows three different ways to create Layer 2 loops.

If not properly handled, these loops cause serious problems that can effectively disable a network.
These loops can cause Broadcast storms, and Multiple frame copies, which can in turn cause instability of the MAC address table.

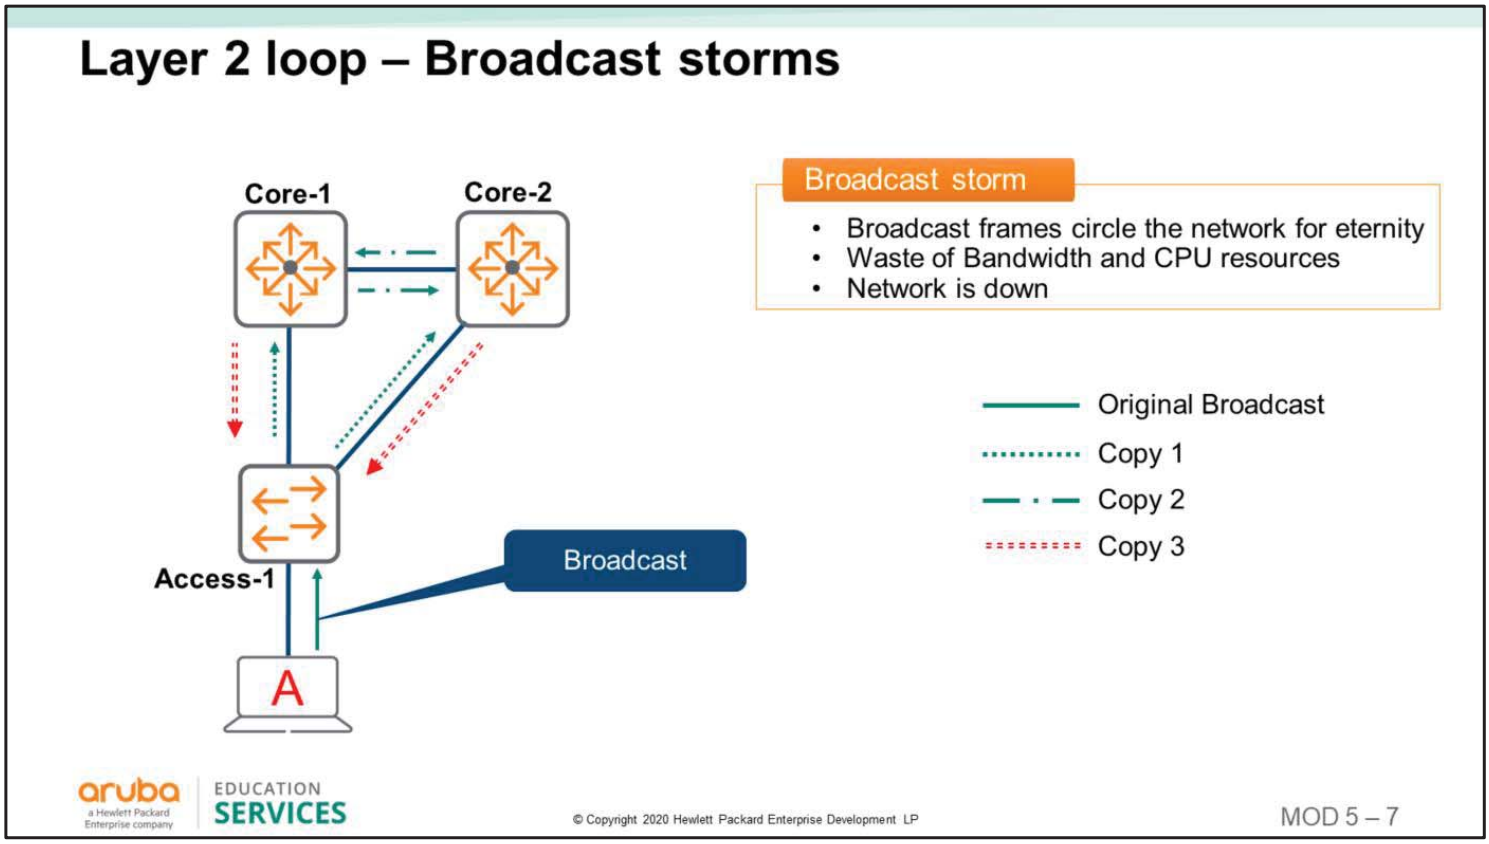

You learned in module 3 that switches flood broadcasts out all interfaces in the same VLAN, except the ingress port. This can cause problems on a redundant, looped network like the one in the figure.

|Suppose that Host A sends a broadcast frame.
|Switch Access-1 receives this frame, and floods it out (copy 1)
|Core-1 and Core-2 receive this and floods it out their connections to each other (copy 2)
|Core-1 and Core-2 and Access-2 receive the second copy, and flood them out their other ports (copy 3)
|Access-1 receives the third copy, and the cycle continues.
|This frame circles around the network for all of eternity.

And remember that Access-1 flooded this broadcast to Access-2, which forward it to Core-2, etc. So a copy circles the network in the other direction, forever.

As a natural part of network and endpoint operation, nearly all devices send broadcast frames, often many times per minute. Every broadcast from every device circulates around the network forever.
|Soon, all available bandwidth and CPU cycles are used up processing broadcasts. No resources are available to process normal data communication frames. The network is effectively unusable.

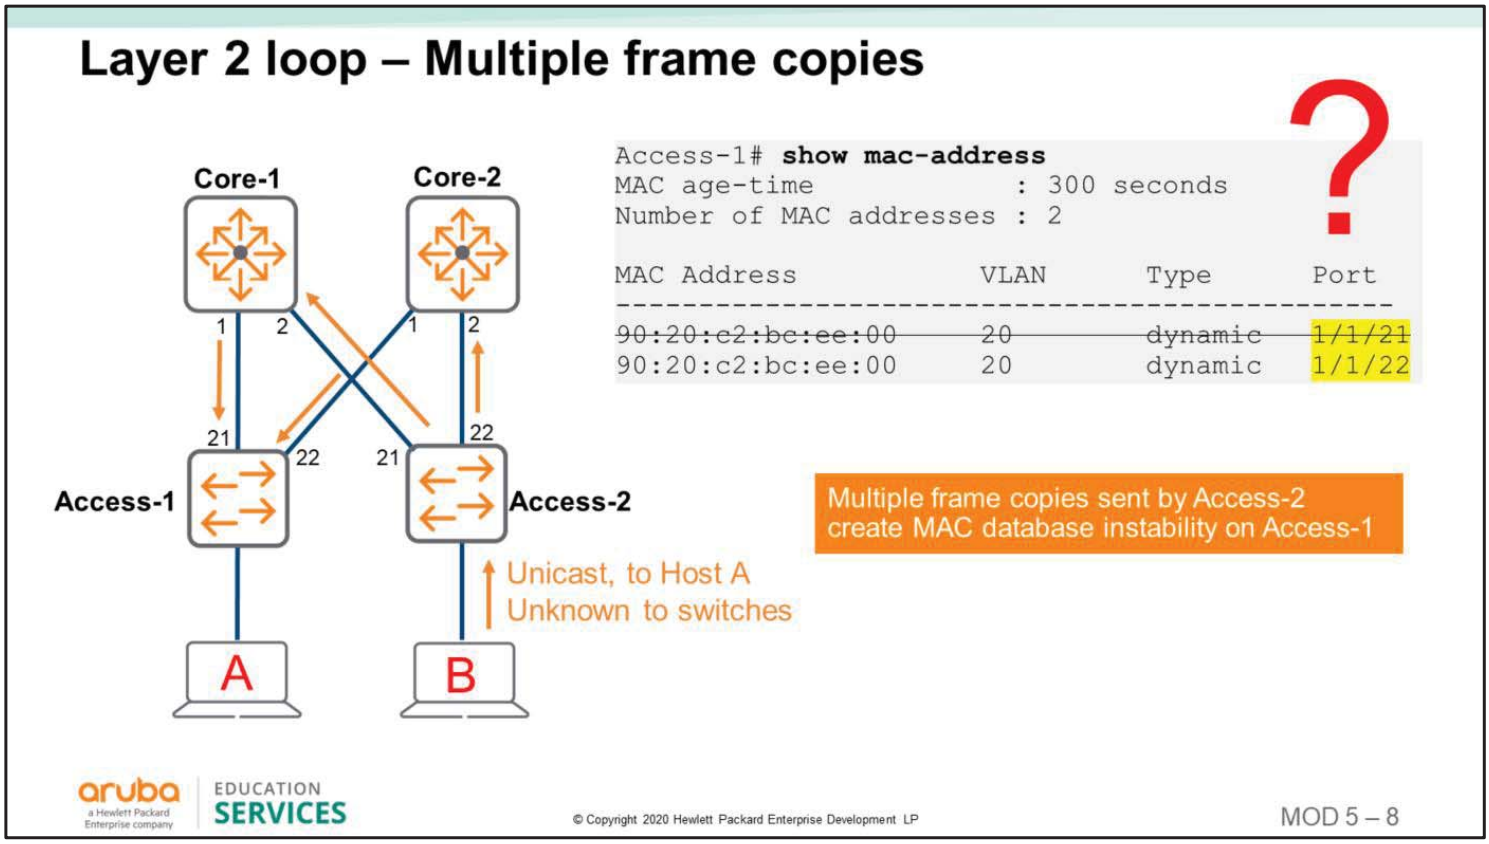

You may recall that switches not only flood broadcasts out all ports in the VLAN (except the ingress port), but they do the same thing with unicasts to an unknown destination. If a switch receives a unicast packet to some destination MAC address, and that address has not yet been learned in the MAC address table, then the switch floods the frame out all ports in the VLAN (except ingress port). These can also circulate and waste bandwidth and CPU cycles. They also cause other issues.

Suppose that Host B sends a frame to Host A.
Access-2 has not yet learned Host A's MAC address, and so it floods this frame out ports 21 and 22. Thus, Core-1 and Core-2 both switches receive a frame with source MAC 90:…:00, inbound on port 21.
Both cores send a copy to Access-1 switch.
Access-1 now believes that it can reach Host B via port 21 and 22.

Access-1 is thus confused – Host B can only exist in one place, but it appears to exist in two places.
The multiple frame copies generated by Access-2 creates MAC database instability on Access-1.

As you can see, these are serious problems. You need redundant links for reliability, but this redundancy can bring your network to its knees. It is time to learn about the solution to these challenges – the Spanning Tree Protocol. Invented in 1984 by the brilliant Miss Radio Perlman. Legend has it that she created the algorithm for this protocol in a few hours. She was so pleased with her creation that she then wrote a poem about it, and then when home for the day.

Spanning Tree Protocol

802.1d and 802.1w

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

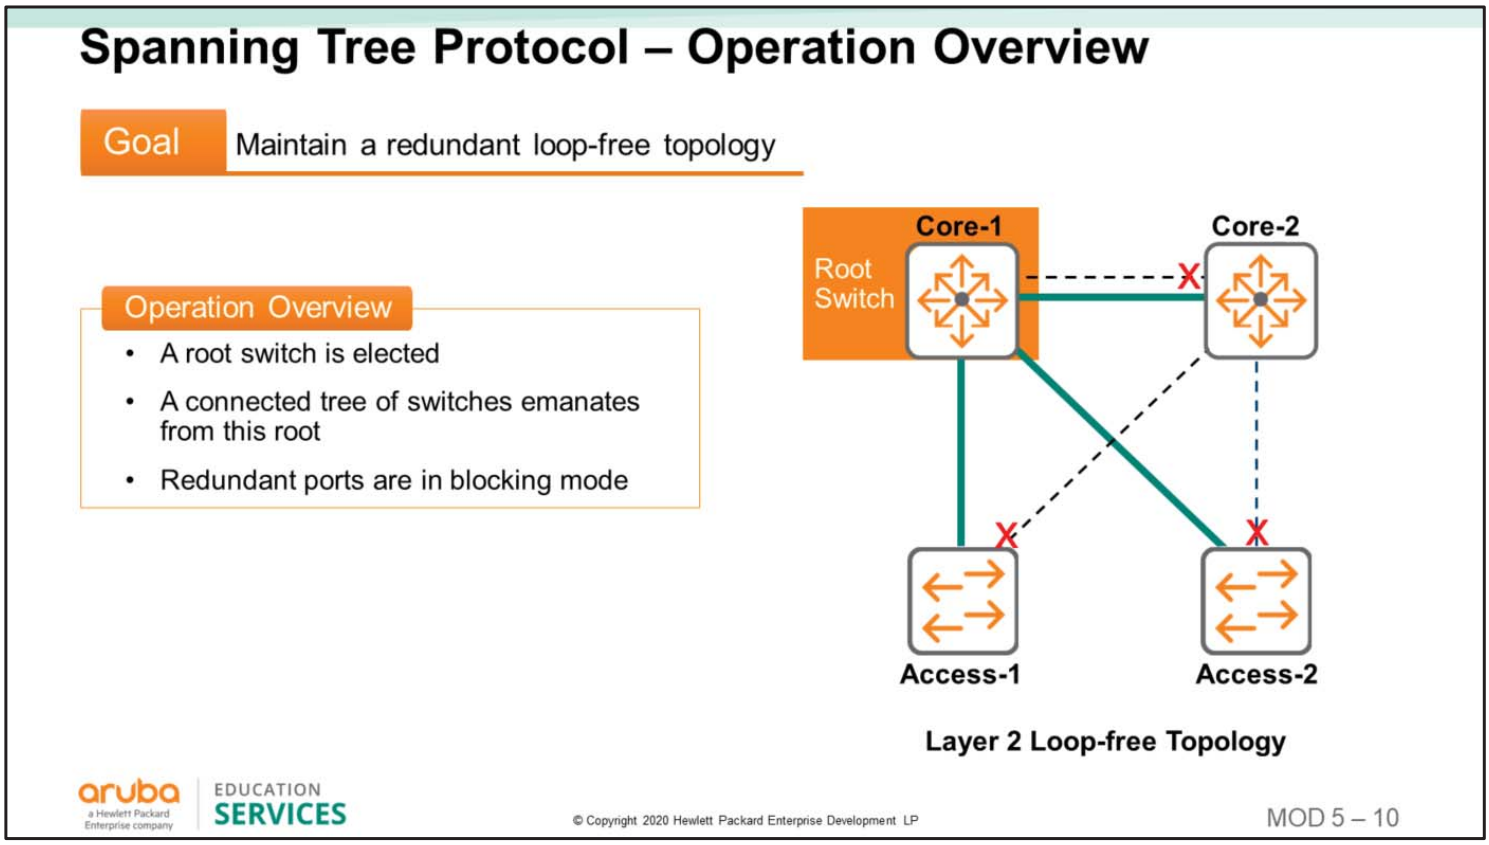

The IEEE 802.1d standard version of the Spanning Tree Protocol (STP) was developed to build and maintain redundant, yet loop-free networks. With STP, you eliminate single points of failure, while avoiding loops and MAC table instability. STP creates a loop-free topology by automatically disabling redundant links.

The figure shows a highly redundant network, for resiliency and fault-tolerance. These redundant links could cause loops and their associated issues.

However, once STP is engaged, certain redundant links are automatically disabled – there are no loops, so there are no problems with broadcast storms or MAC address table instability.

However, these disabled links remain available to provide redundancy as needed.

To accomplish its goal, One switch in the Spanning Tree domain is elected as a root bridge or root switch. Many trees grow out of the ground, from its roots, as a single trunk, and then branch out from there. Likewise, the root bridge is the reference from which the spanning tree grows. All other switches are non-root bridges (switches), sometimes knows as designated switches. A loop-free path “grows” from this root switch out to all non-designated (non-root) switches.

Note: What is now called a “switch” is a very fast, multiple-port version of what used to be called a “bridge” (many years ago). When discussing STP, you will often see the term bridge. In this context, “bridge” and “switch” can be thought of as the same thing. When you see “bridge’, know that the discussion is actually about a switch.

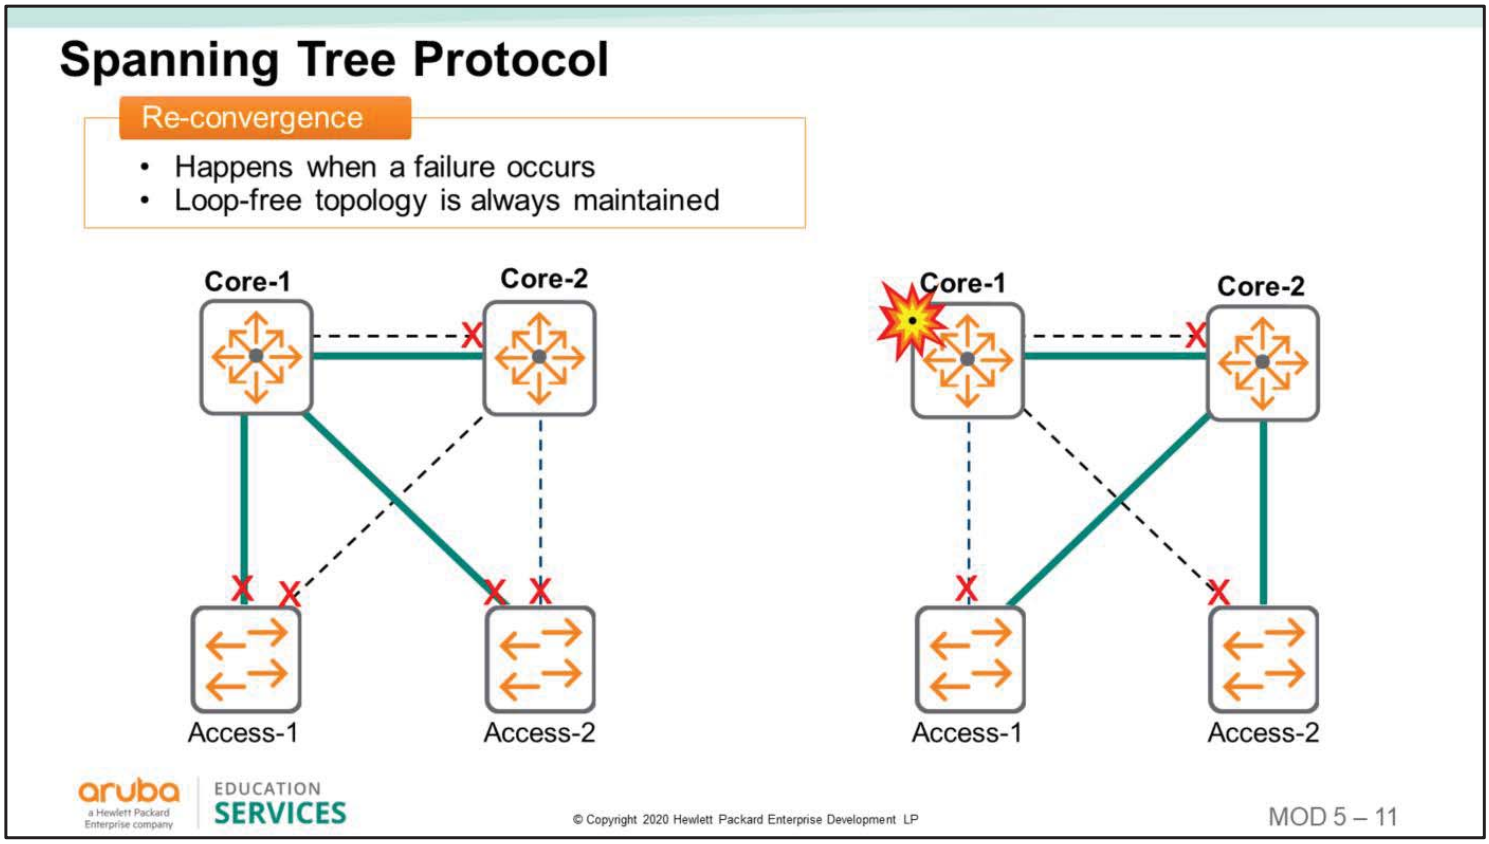

The STP algorithm runs on all switches, and redundant links have been disabled to avoid loops.
However, if a failure occurs, STP converges on a new topology of active links, with are used to forward frames. A loop-free topology is always maintained.

Note: The word “converge” can mean something like “to meet at some point”. Convergence, in the context of networking, means that the devices all come to an agreement about a new network topology. All devices converge on a new set of active paths to be used for frame forwarding.

Overview of Spanning Tree Protocol

STP	RSTP	MSTP
802.1D (pre-2004)	**802.1w** merged into **802.1D 2004**	**802.1s** merged into **802.1Q 2005**
• Original protocol • Slow convergence	• Faster convergence • Costs that fit with modern port speeds	RSTP behavior in multiple instances
• Obsolete • Not recommended	Default *operation* between AOS-CX switches (when no MSTP region is set up)	Default *mode* for **AOS-CX** switches (but no MSTP region settings)

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 12

In the original standard the failure detection is based on timers. With the IEEE 802.1d default, the root switch (and only the root switch) originates a “hello” packet every 2 seconds. These hello packets are forwarded on out to the rest of the switches in the domain. If some switch downstream (farther away from the root switch) ceases to receive hello packets for 20 seconds (the default Max Age timer), an outage is assumed, and all switches begin to converge on a new topology. The duration of this process is defined by the Forward Delay timer, which is 15 seconds by default. This protocol is considered obsolete and its use is no longer recommended in modern networks.

Rapid Spanning Tree Protocol (RSTP) was developed in 1998 to speed convergence. Instead of only the root switch originating hello packets, ALL switches originate them. This means that RSTP now has a true keep-alive mechanism that can respond in seconds (or less). The need for old and slow Max Age and Forward Delay timers is eliminated. You will soon learn the details of this new convergence process.

Note. RSTP is backward compatible with the original standard 802.1d. However, to maintain this capability, some benefits of RSTP are lost.

802.1s or Multiple Spanning Tree was developed to improve the performance of the protocol for implementing multiple Loop-free Topologies or Instances to load balance the traffic across all links. This protocol helps to create optimal paths.

MSTP runs on AOS-CX switches by default. However, if no setting is configured then this protocol behaves like RSTP. You will understand why this happens later in this training.

Note:

AOS-CX 6300 Spanning-Tree is globally enabled by default
AOS-CX 8325 Spanning-Tree is globally disabled by default

Rapid Spanning Tree Protocol Elements

Bridge ID, BPDU, Port state, Port cost, Port roles

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Rapid Spanning Tree Elements Overview

Key Components

- Bridge Identifier (BID)
- Bridge Protocol Data Unit (BPDU)
- Path Cost
- Port States
- Port Roles

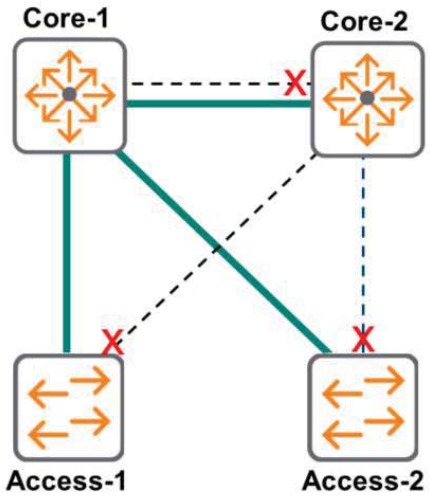

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 14

In this section you learn about key elements that the Rapid Spanning Tree Algorithm uses to decide which ports to enable for traffic forwarding, and which ones will not.
This includes the following

- Bridge Identifier
- Bridge Protocol Data Unit
- Rapid Spanning Tree Port States
- Rapid Spanning Tree Port Cost
- Rapid Spanning Tree Port Roles

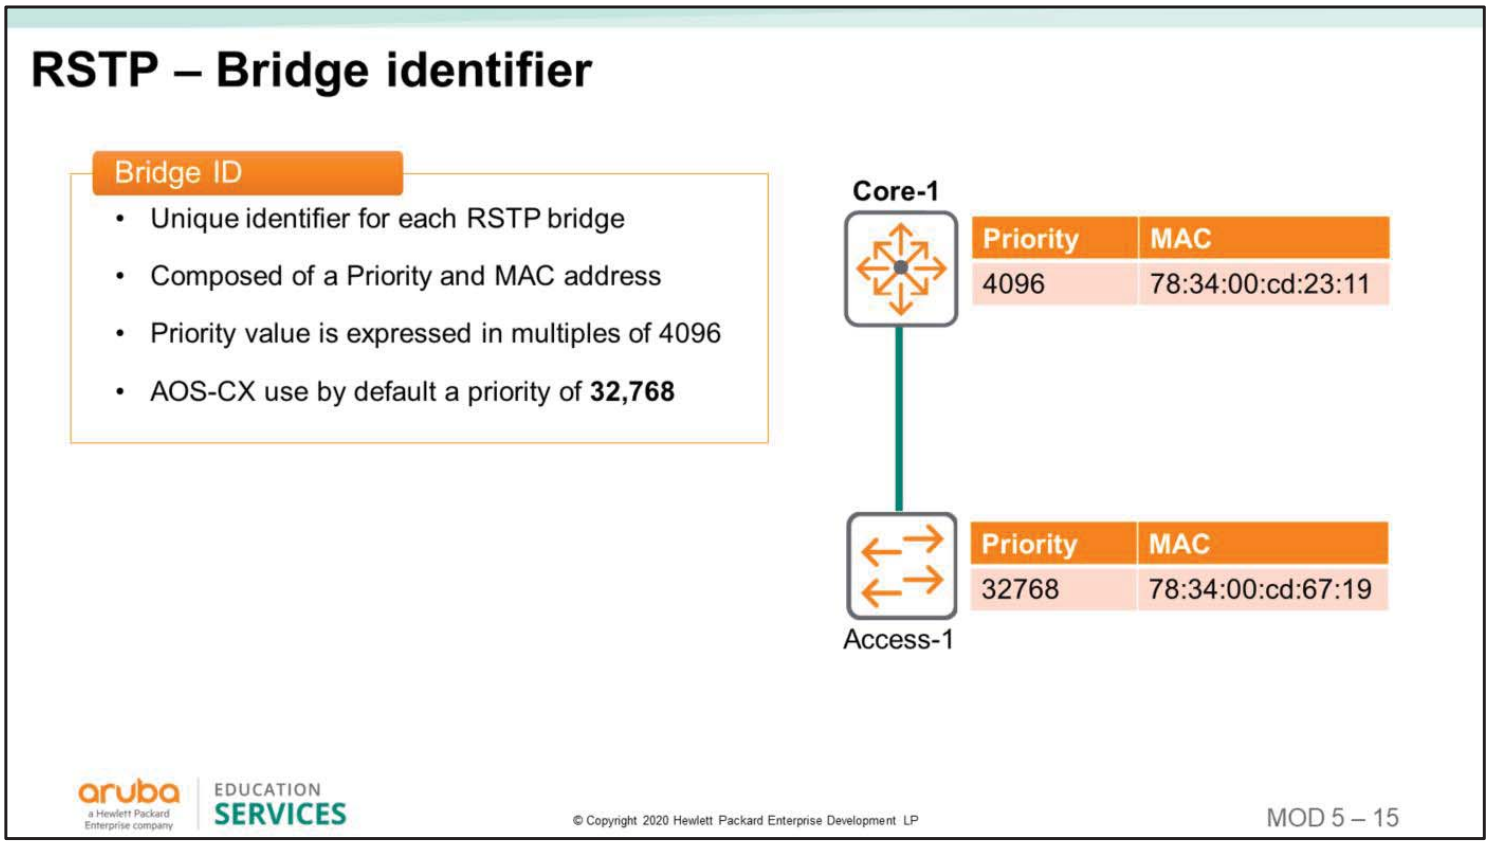

Bridge Identifier

Spanning Tree assigns each switch a unique identifier called Bridge ID. This identifier is composed of a 2-byte priority and 6-byte MAC address. The priority defaults to 32768, so by default, all switches have the same priority. However, each switch has a unique MAC address, and so each Bridge ID (BID) will be unique.

In AOS-CX switches have by default a priority of 32,768.

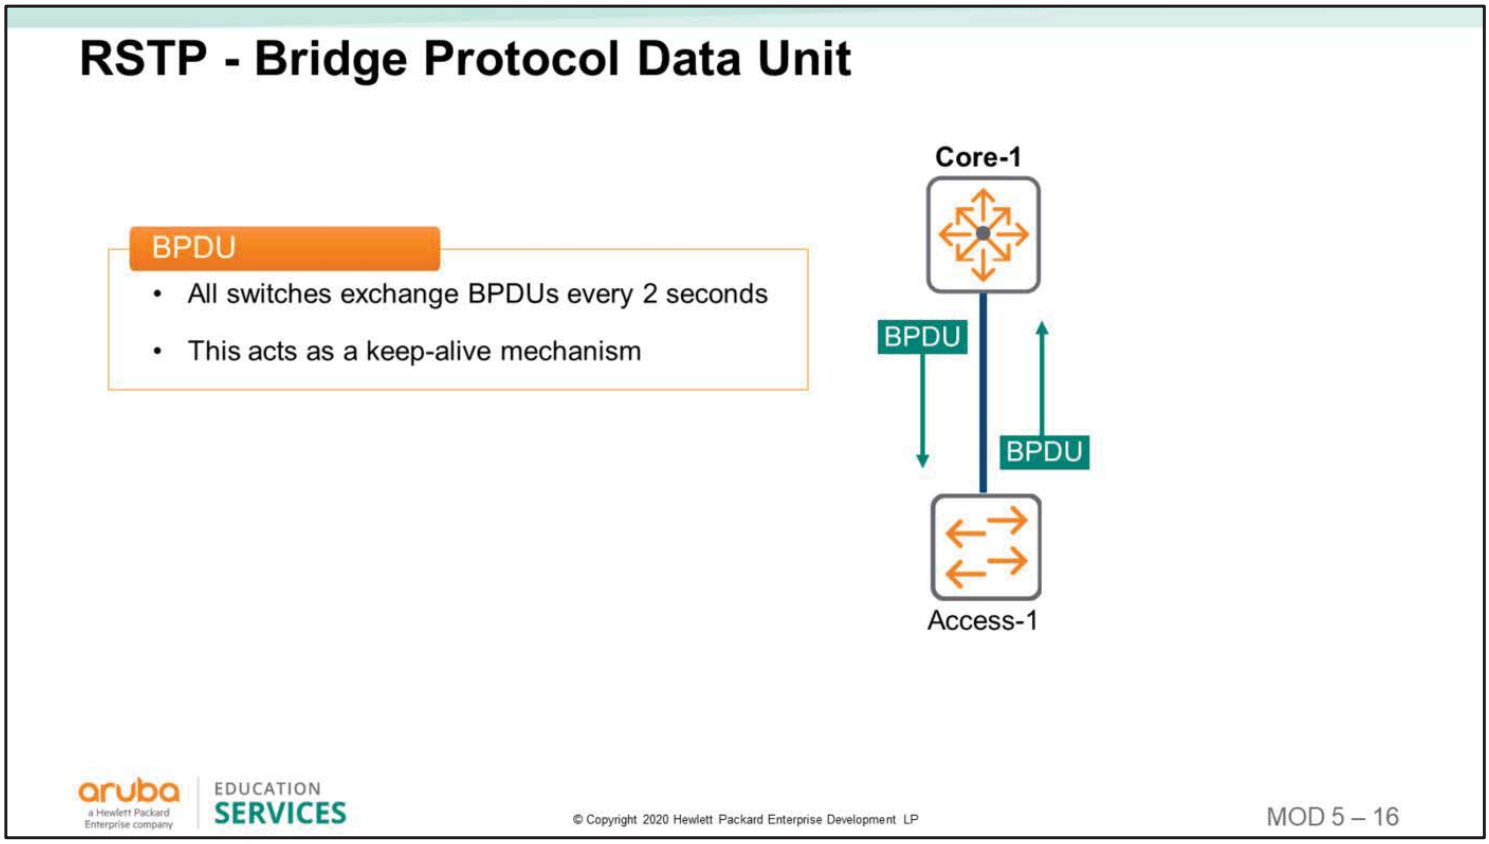

Bridge Protocol Data Unit

All switches that participate in the Spanning-Tree Algorithm exchange control messages called Bridge Protocol Data Units (BPDU).

In the original 802.1d standard, BPDUs were generated only by the root switch, and then had to “trickle down” to other switches. This led to the need for the slow MAX AGE and Forward Delay timers.
With RSTP, all switches originate BPDUs with its current information, every 2 seconds - the default hello-time period. Thus, if a port stops receiving BPDUs for three consecutive hello timers, the switch quickly knows that it has lost connectivity to its neighbor. It ages out the protocol information and begins to converge.

Because each switch originates BPDUs, this becomes a true keep-alive mechanism. Failure detection will take no longer than 6 seconds.

RSTP – Port States

STP (802.1d) Port State	RSTP (802.1w) Port state	Port Included in Active Topology	Port learning MAC addresses?
Disabled	Discarding	No	No
Blocking		No	No
Listening		No	No
Learning	Learning	Yes	Yes
Forwarding	Forwarding	Yes	Yes

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 17

During original tree establishment and any ensuing convergence, switches must decide which ports will forward data and which ports must be disabled to prevent Layer 2 Loops. Spanning Tree uses port states to transit from a Blocking port to a Forwarding port. The table in the slide summarize the port states and their specific task. This table lists the states that were used on 802.1d and compare them against the new port states in 802.1w.

To transition from Blocking to Forwarding, the original 802.1d standard takes 30 seconds - 15 seconds are spent in a Listening state, and 15 in the Learning state.
One of the reasons RSTP is more efficient than 802.1d is because the port transitions quickly from discarding to learning to forwarding.

Note. Blocking and Listening state are also deprecated in MSTP.

Note. In RSTP the Learning state is a transitory state and is only used during a re-convergence of the protocol when a change happens in the topology. Stable states will be Discarding or Forwarding states.

RSTP – Path cost

Link Speed	RSTP Port Cost in ArubaOS-CX
100 Mbps	200,000
1Gbps	20,000
10 Gbps	2,000
100 Gbps	200
1 Tbps	20

BPDUs advertise the path cost

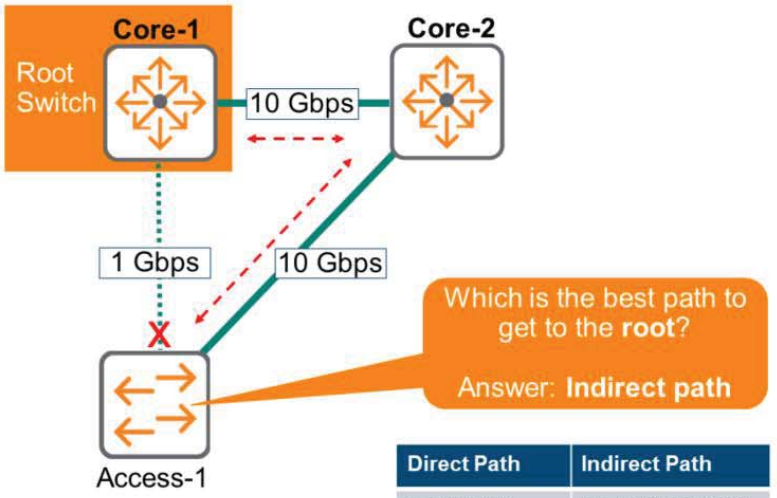

Direct Path	Indirect Path
0 + 20,000 = **20,000**	0 + 2,000 + 2,000 = **4,000**

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 18

RSTP may have several possible paths to get from the root switch to some non-root switch. It chooses the best path based on cost, which is based on link speed. The figure shows the AOS-CX default port costs.

Consider the example shown in the figure. Intuitively, you might think the root switch's best path to Access-1 is the direct path
However, this is a 1Gbps link, which has a cost of 20,000.

The indirect path to Access-1 via Core-2 requires two 10Gbps links, which only have a cost of 2,000 each. 2,000 + 2,000 = 4,000 – far less than 20,000. (Strictly speaking the root switch will add its cost to itself in this case the cost is equal to 0).
Thus, this is the preferred, best path, and the redundant path is disabled to avoid loops.

RSTP - Port Roles (Designated and Root)

Designated Port

- Closest to the root switch
- All root switch ports are designated
- Port state = **Forwarding**

Root Port

- Another switch in the link is closer to the root
- Best path to the root switch (for a non-root switch)
- Port state = **Forwarding**

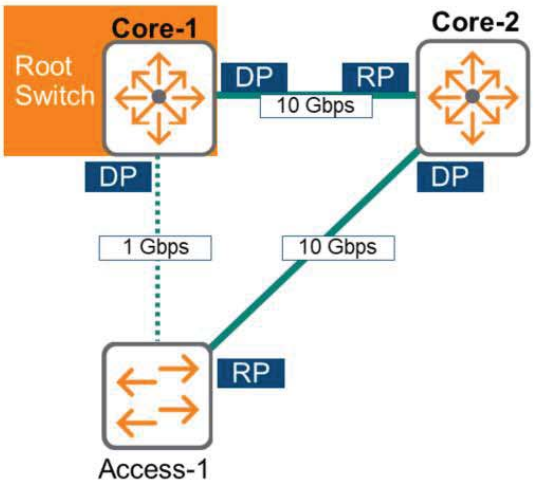

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 19

Switches that use 802.1w RSTP do not have a complete view of the topology. They build and maintain the loop-free topology by exchanging BPDUs, which indicate how close they are to the root switch. BPDUs help switches to calculate the correct port role for each of their ports.

All ports on the root switch are Designated ports. The root switch is like the “boss” of the domain. It does not block its ports – only non-designated ports must worry about that. Note also that there are no loops in the topology shown, so all ports are forwarding. Looking at the non-root switches, what is the difference between the Root Port and Designated port? They both forward frames. The Root Port is simply the port that is closest to the root switch – the best path toward the root switch. The designated port is designated to accept traffic from downstream root ports.

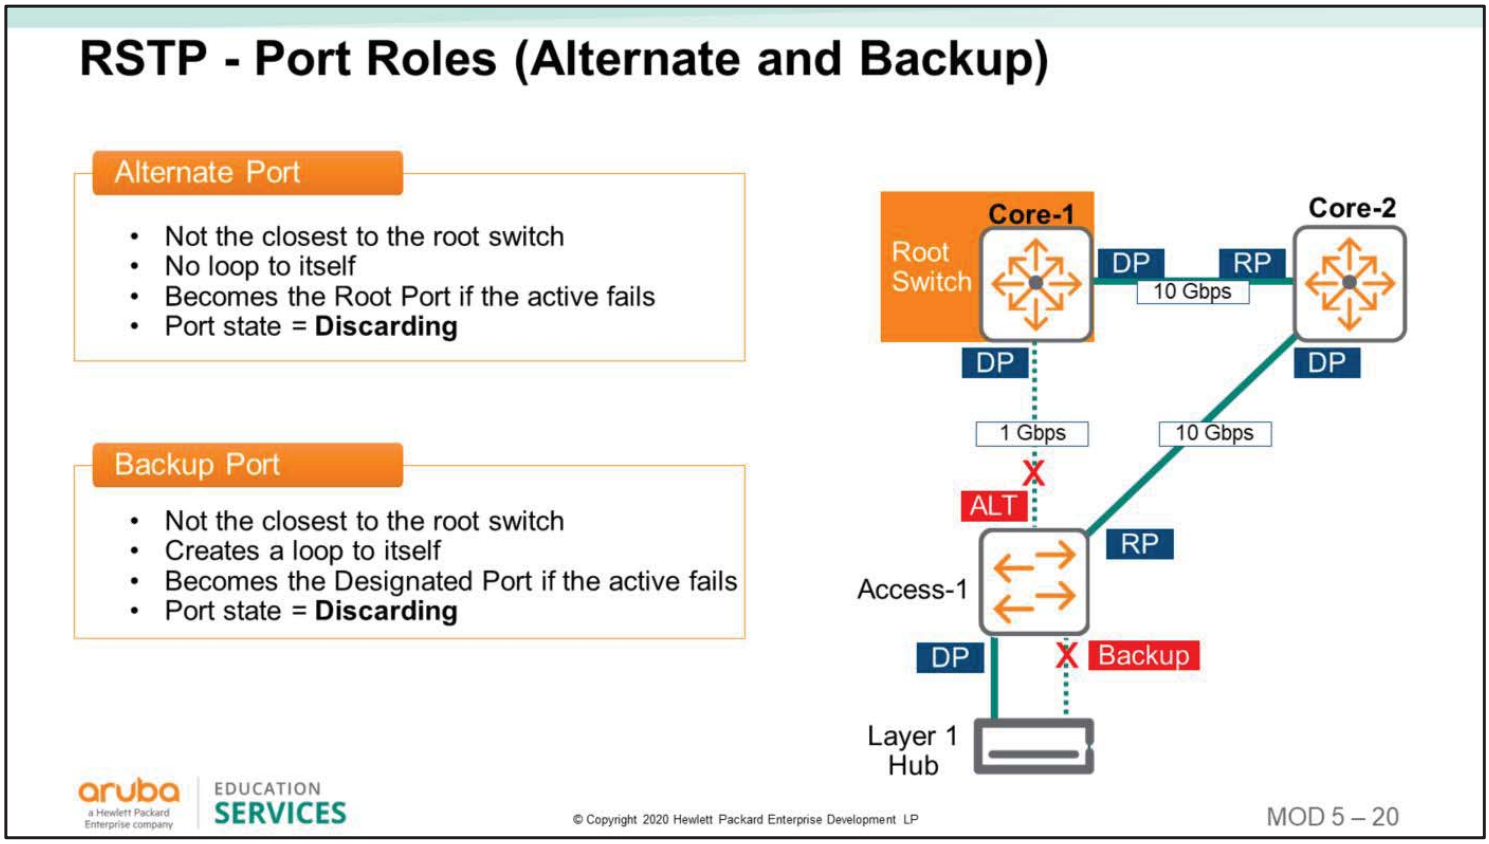

Alternate

Port is not the closest to the root on this link.
This link does not offer the switch its best path to the root.
Port is not connected to the same switch. No Loop to itself.
This port becomes the root port if the active port fails.

As an example, consider the following situation. Access-1 port 1/1/21 can become an Alternate port since it fulfills the requisites to it.

Backup

Port is not the closest to the root in this link.
This link does not offer the switch its best path to the root.
Port is connected to the same switch. The switch has a looped connection to itself or it has a loop created by a Layer 1 hub.
The backup port becomes designated port if the existing designated port fails

A loop could be caused when a Layer 1 hub is placed on the topology. As an example take in consideration the following image, in this situation ports 1/1/26 and 1/1/27 on Access-1 are in a loop situation because a Hub was introduced to the topology. If RTSP was running on Access-1 neither of these ports (1/1/25 and 1/1/26) will become an Alternate Port.

A backup port is considered when a Hub is connected to the network. In this case Port 1/1/26 becomes the backup port.
Note: Introducing Layer 1 Hubs into the topology must be avoided.

RSTP Operation

Building the Loop-free Topology

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Rapid Spanning Tree – Operational Overview

Building Loop-free Topology

1. Elect the Root switch
2. Select the Root Port on all non-root switches
3. Select the Designated Port
4. Other ports are moved to Discarded state

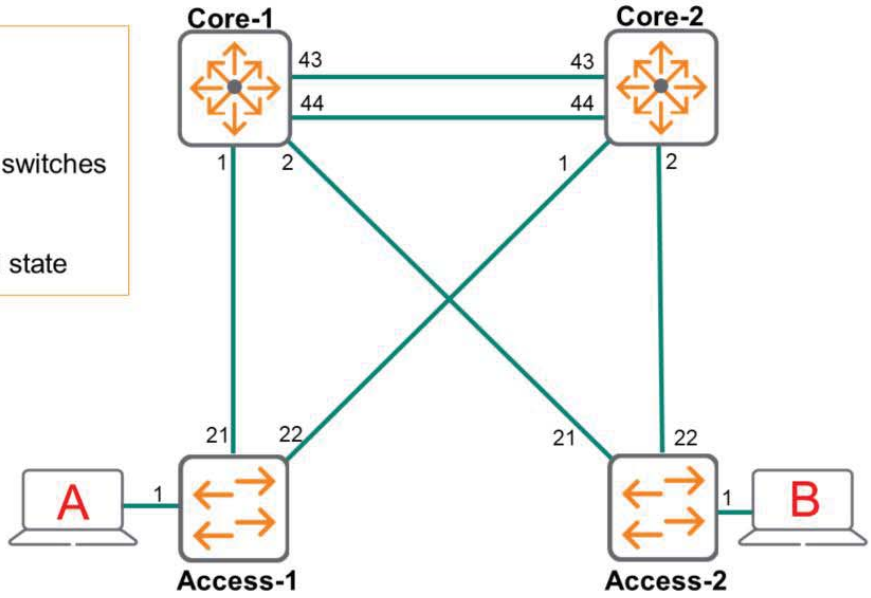

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 22

The Rapid STP algorithm converges on a loop-free topology by performing the following steps:

Elects Root switch.
Select the root port on all non-root switches or Non-designated bridges.
Select the designated port on each switch-to-switch links.

Other ports move to a blocking state to prevent loops

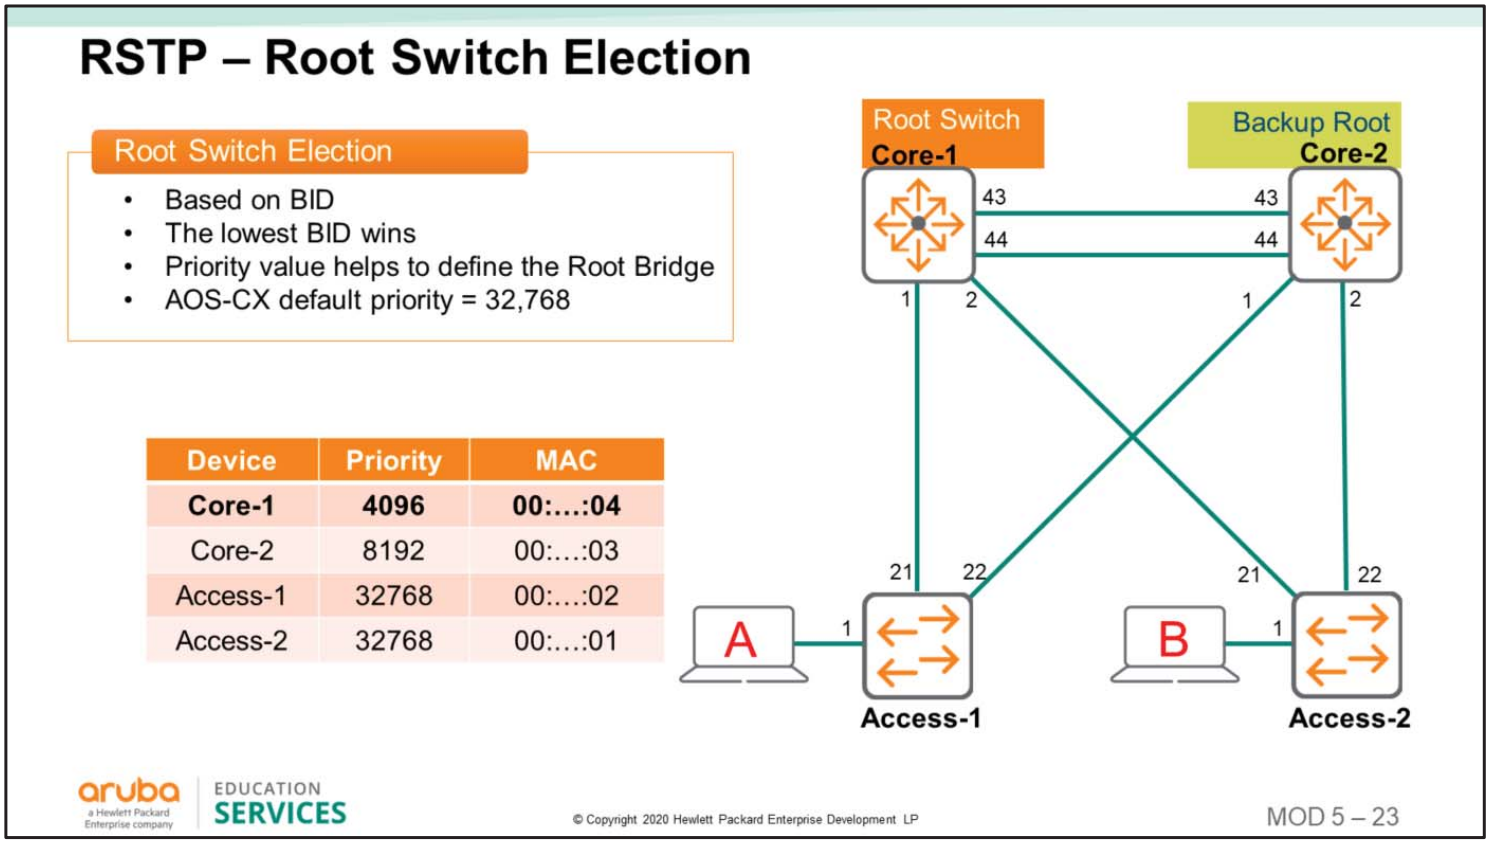

For both the 802.1d and 802.1w standards, root switch election is based on the Bridge ID (BID) - the lowest value wins. Recall that the two-part BID is composed of a priority, which defaults to 32,768 for all switches, and a globally unique MAC address.

You can rely on this default, rather random behavior, but it is not recommended. Some small, low-powered switch on the edge of your network might win the election. This makes for a less stable tree, with sub-optimal paths and poor resiliency. Consider the topology shown in the figure. If the priority is set to the default, then Access-2 will become the Root Switch - it has the Lowest MAC address.

You want to ensure that one of the high-powered, more centrally-located core switches wins the election. You have a more robust, resilient, and optimally pathed tree structure. To do this, you simply lower the priority value of the preferred switches, lower than the 32,768 default. In this example, Core-1 becomes the Root Switch simply by setting its priority to 4096. However, if Core-1 fails, then Access-2 will become the Root Switch, which would not be optimal. Therefore it is important that Core-2 becomes the second-best option. In this case Core-2 is setup with a priority of 8192.

Recall that the priority value must be configured in increments of 4096. You might lower the value of Core-1 to 4096, and the value of Core-2 to 8192.

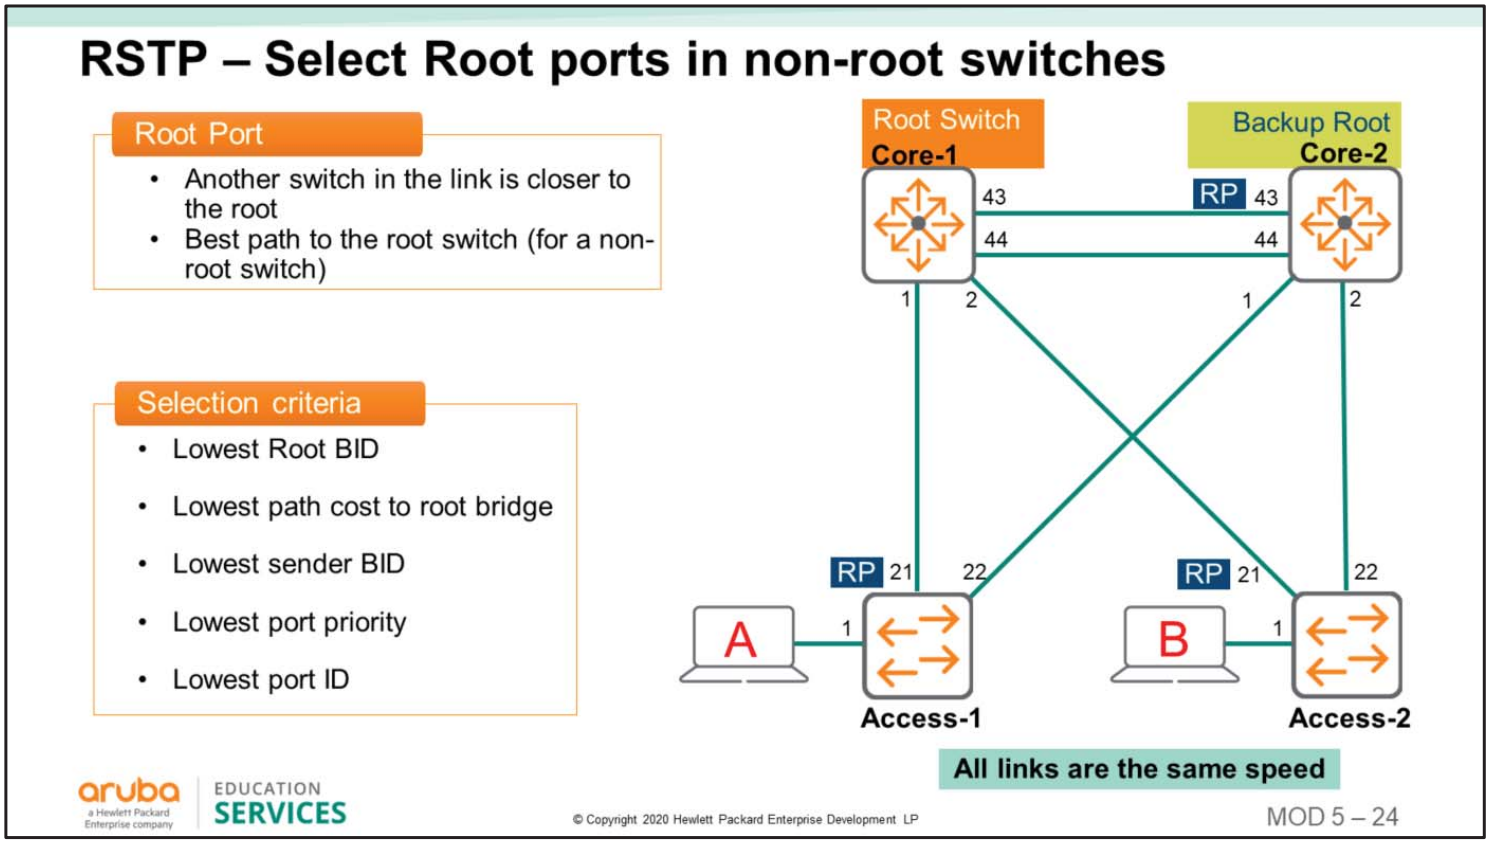

Non-root switches or Non-Designated bridges must select the best path to the root switch, by selecting its root port - the port connected to this best path. Root port selection criteria are as follows:

Lowest root bridge ID.
Lowest path cost to the root bridge.
Lowest sender bridge ID.
Lowest port priority.
Lowest port ID.

Core-2 analysis: This device will receive BPDUs on ports 1, 2, 43 and 44. Core-2 must decide which port is the best one:

Lowest Root Bridge ID. All devices agree that the root bridge is Core-1, this criteria does not help to select the best path.
Lowest path cost to the root bridge. Assuming that all links are the same speed then, indirect paths on port 1 and 2 are discarded simply because the cost is higher than direct paths. Ports 43 and 44 have the same cost, the next criteria must be used.
Lowest sender Bridge ID. BDPUs received on ports 43 and 44 are generated by the same device (Core-1), this criteria does not help to break the tie.
Lowest port priority. The sender (Core-1) will include a port priority on the advertisement. The port priority comes in range between 0-240 and in AOS-CX the default value is set to 128, a bridge will consider a lower number the winner. In this case assuming that default values are set this criteria does not help to break the tie.
Lowest port ID. The final decision is taken on the port ID, where the lowest the value the winner. Core-1 lowest port number is 43, so the port on Core-2 that is connected to this port

becomes the Root Port (in this case is Core-2's port 43).

Access-1 analysis: This device will receive BPDUs on ports 21, 22. Access-1 must decide which port is the best one:

Lowest root bridge ID. All devices agree that the root bridge is Core-1, this criteria does not help to select the best path.
Lowest path cost to the root bridge. Assuming that all links are the same speed then, indirect paths on port 22 is discarded simply because the cost is higher than the direct path. Ports 21 becomes the Root Port.

Access-2 analysis: This device will receive BPDUs on ports 21, 22. Access-2 must decide which port is the best one: A similar decision will take as Access-1, port 21 the direct path offers a best path so it will become the root port.

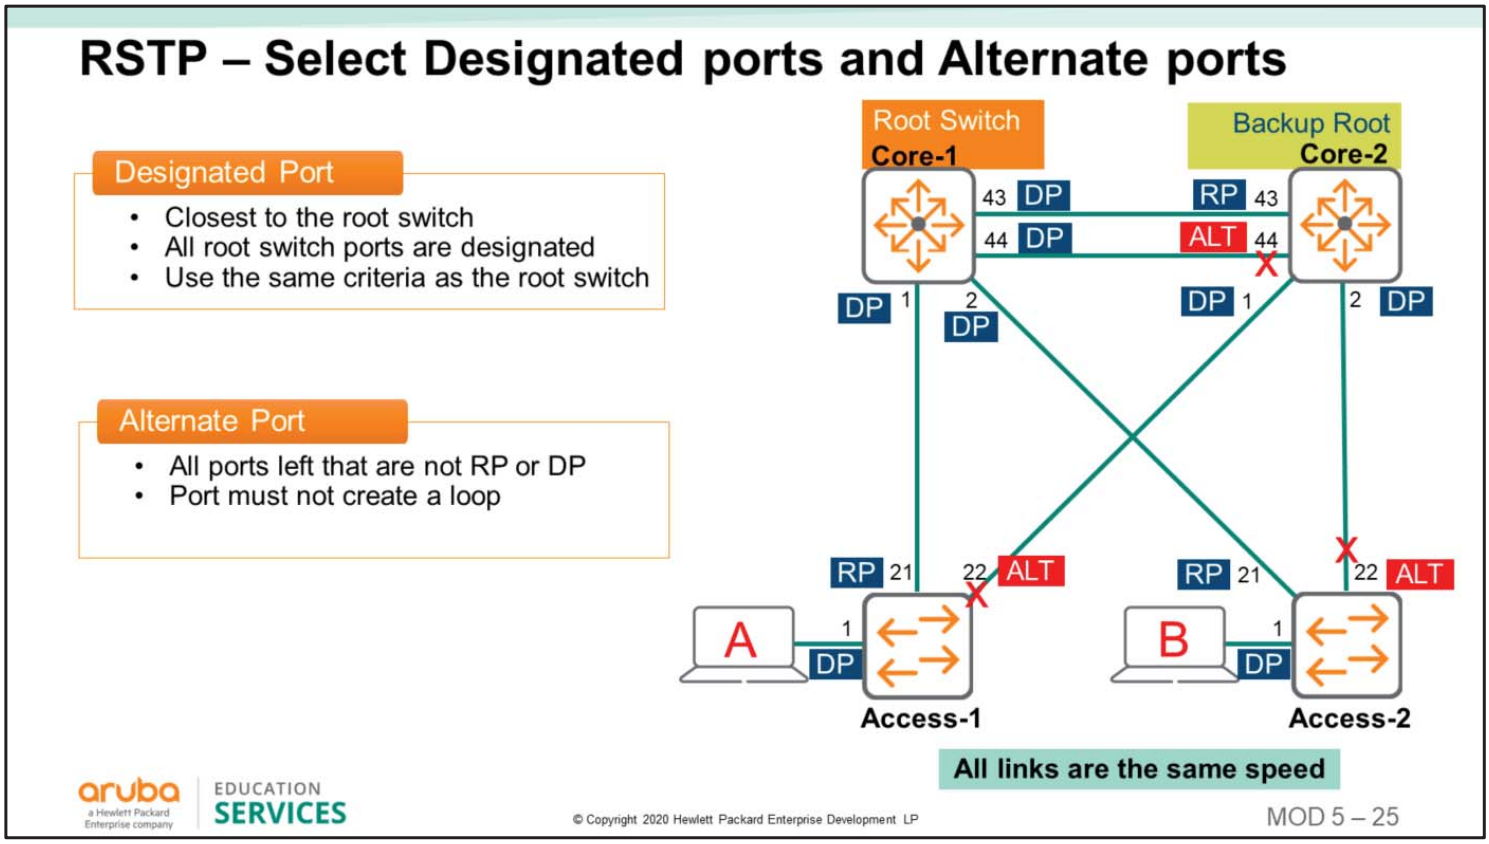

The criteria to select a Designated port is the same as root port.

Lowest root bridge ID.
Lowest path cost to the root bridge.
Lowest sender bridge ID.
Lowest port priority.
Lowest port ID.

The following statements are a shortcut to determine quickly a designated port on a link.

All ports in the root switch are always designated ports.
In a switch-to-switch link where previously a root port was elected on one side, the other side must always be a designated port.

Core-1 analysis: This is the Root switch and its ports will always be the closest to the root bridge, therefore ports 1, 2, 43 and 44 on Core-1 become Designated Ports.

Core-2 analysis: The analysis focus on ports 1, 2 and 44.

Port 1. Core-2 will evaluate if this port is closest to the root, this process is done by comparing the BPDU received from Access-1 on this port and the one that it sends out. This process follow the same rules as the root port, so let's start the analysis:

Lowest root bridge ID. All devices in the topology agree that Core-1 is the root, this criteria does not break the tie.
Lowest path cost to the root bridge. Access-1 and Core-2 advertises the same cost to reach the root switch (both are 1 hop away).

Lowest sender bridge ID. This criteria breaks the tie since Core-2 has a lower BID than Access-1. Port 1 in this link becomes the Designated Port.

Port 2. Core-2 will evaluate if this port is closest to the root, this process is done by comparing the BPDU received from Access-2 on this port and the one that it sends out.

Lowest root bridge ID. All devices in the topology agree that Core-1 is the root, this criteria does not break the tie.
Lowest path cost to the root bridge. Access-2 and Core-2 advertises the same cost to reach the root switch (both are 1 hop away).
Lowest sender bridge ID. This criteria breaks the tie since Core-2 has a lower BID than Access-2. Port 2 in this link becomes the Designated Port.

Port 44. This port cannot be Designated since the port on the other side of the link is closest to the root.

Access-1 analysis: This analysis focus on port 1 and 22

Port1. This port is the only RSTP speaker in the link, this port becomes the Designated Port.
Port22. This port cannot be Designated since the port on the other side of the link is closest to the root.

Access-2 analysis: This analysis focus on port 1 and 22

Port1. This port is the only RSTP speaker in the link, this port becomes the Designated Port.
Port22. This port cannot be Designated since the port on the other side of the link is closest to the root.

In this example, the topology does not include any Layer 1 hub which implies that there will not exist a Backup port only Alternate Ports. Simply all ports that were not elected as Designated or Root ports will become Alternate and their state will be discarding. In the slide you can notice that Access-1’s port 22, Access-2’s port 22 and Core-2’s port 44 become Alternate ports.

RSTP Edge Ports and Link Types

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

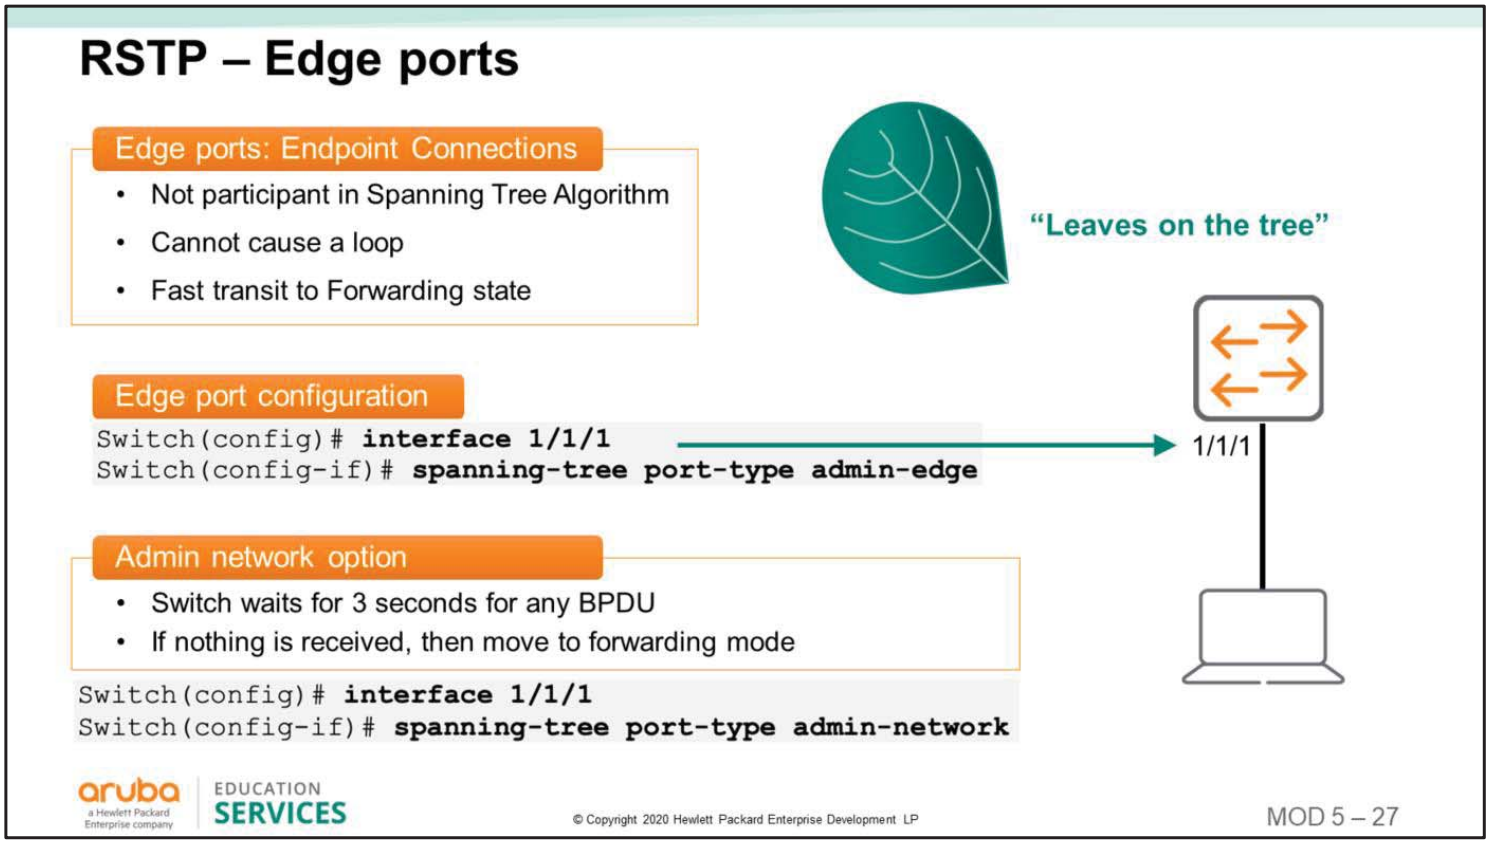

Edge ports connect to endpoint, and so they do not receive BPDUs. They need not participate in the spanning-tree algorithm, because they cannot cause a loop – they are “leaf on the spanning tree”, so to speak. Since they cannot cause loops, they can quickly transition to a forwarding state, with no intermediate steps.

If BPDUs are received on an edge port, then the port will act as a normal Spanning Tree port and participate in the algorithm to prevent Layer 2 Loops.

You must manually configure edge ports, as shown here.

Switch(config)# interface 1/1/1
Switch(config-if)# spanning-tree port-type admin-edge

An alternative to the admin-edge option is the AOS-CX administrative-network option. With this option, the port looks for BPDUs for the first 3 seconds after the link is up. If no BPDUs are received, the port becomes an edge port and immediately forwards frames. If BPDUs are detected, the port becomes as a non-edge port. It participates in normal STP operation.

Switch(config)# interface 1/1/1
Switch(config-if)# spanning-tree port-type admin-network

RSTP - Topology Change Mechanism

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

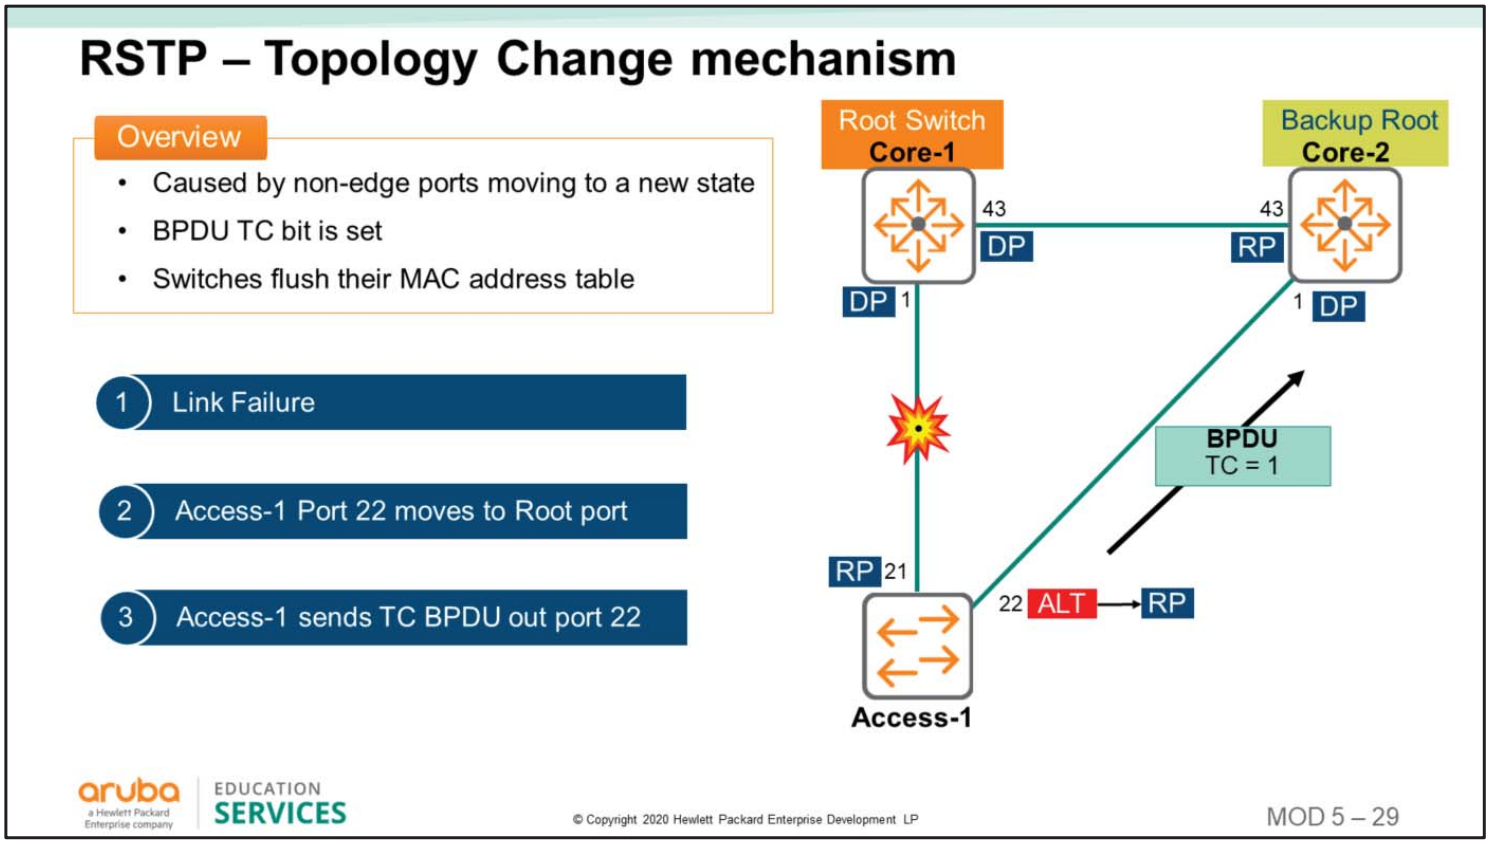

In RSTP, a topology change occurs when non-edge ports move to a different state, or when BPDUs are no longer received. A switch detects this topology change. The switch actively informs the rest of the switches in the network of the Topology Change. It sets the TC bit in BPDUs and transmits these BPDUs. It flushes the MAC address table entries associated with all non-edge ports.

Other switches receive a BPDU with the Topology Change (TC) bit set from a neighbor. The switch clears the MAC address table entries learned on all its ports, except the ingress port for the TC BPDU. These switches in turn send BPDUs with TC set.

Note. In the original 802.1d standard, any switch could send TC BPDUs to notify others, but the instruction to clear the MAC address table must always come from the Root switch. This two-step process takes more time to be completed.

To understand how a Topology change is handle by Rapid STP consider the following scenario:

|A link failure occurs between Core-1 and Access-1.
|As soon as Access-1 detects that its root port is down, it moves its Alternate port 22 to be the new Root Port. Port 22 immediately moves to a forwarding state. The rest of the ports on Access-1 are unaffected.
|Access-1 generates a TC BPDU and sends it out port 22.

RSTP – Topology Change mechanism

4. Core-2 flushes MAC address table
5. Core-2 generates a TC BPDU
6. Core-1 receives TC-BPDU, flushes MAC table

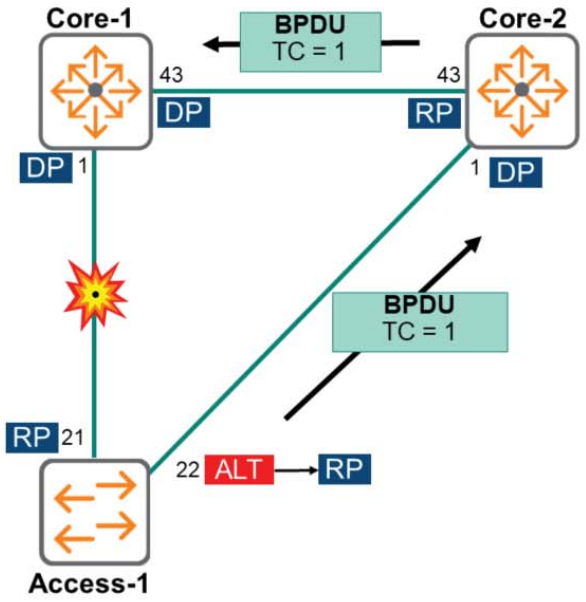

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

|Access-2 flushes its MAC address tables.
|Access-2 sends a TC-BPDU.
|Core-1 receive the TC-BPDU and flushes its MAC address table.

Note. Core-1 also detects the link failure and does the same process as Access-1, generating a TC and advertise it to its neighbors.

Common Spanning Tree

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

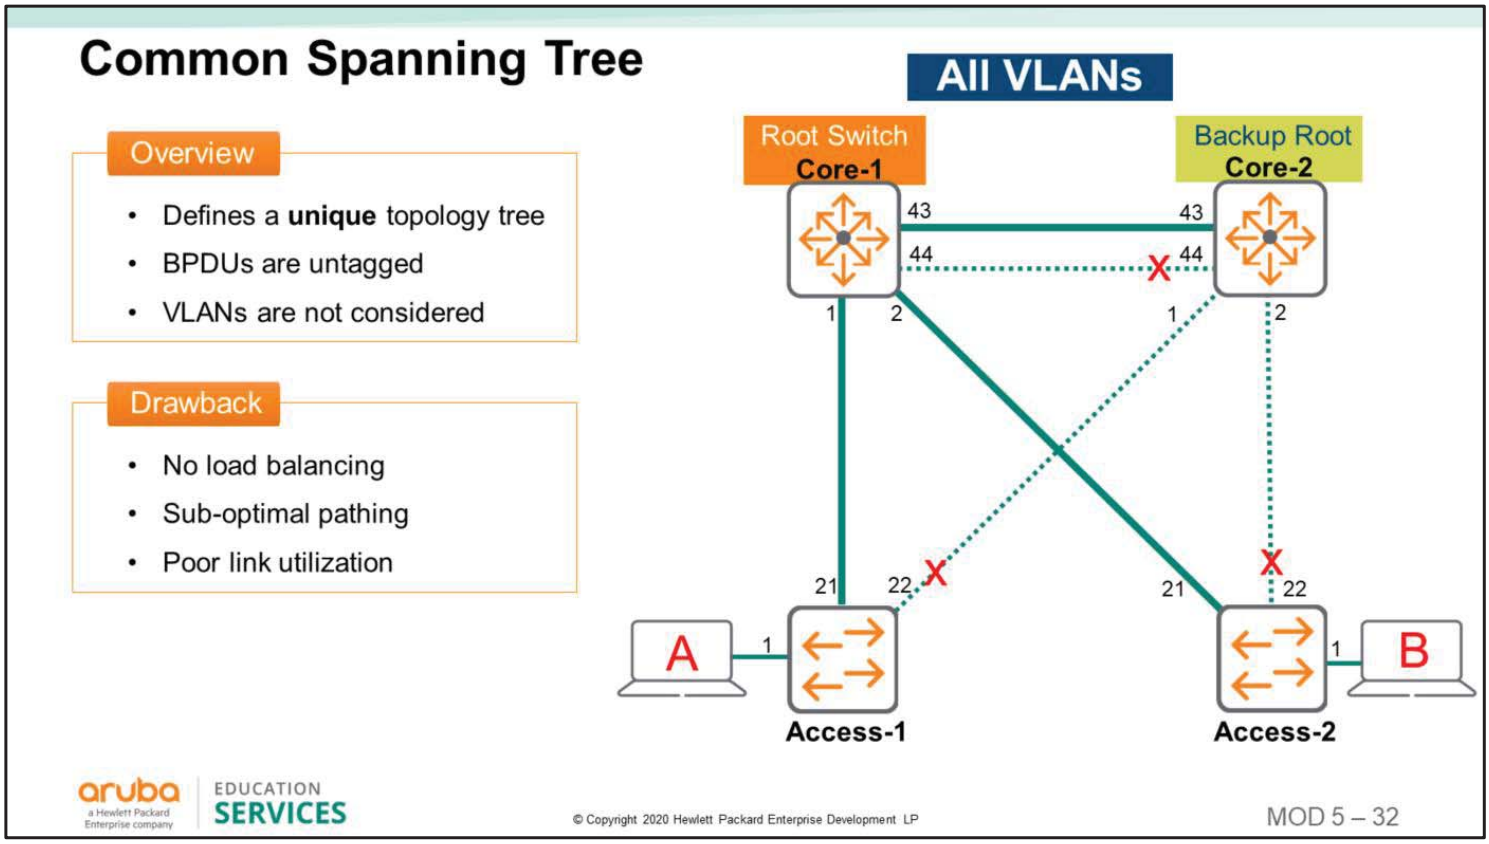

Common Spanning Tree (CST) defines one instance of Spanning Tree for the entire network, regardless of the number of VLANs. This is the sad, original scenario that you just explored. All BPDUs use untagged VLAN 1, and so VLANs are not considered.

There is no load balancing, so you get the sub-optimal pathing, poor link utilization, and potentially slower performance. You will now learn about these issues and how to resolve them with Multiple Spanning Tree.

Multiple Spanning Tree

Theory and Region configuration

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

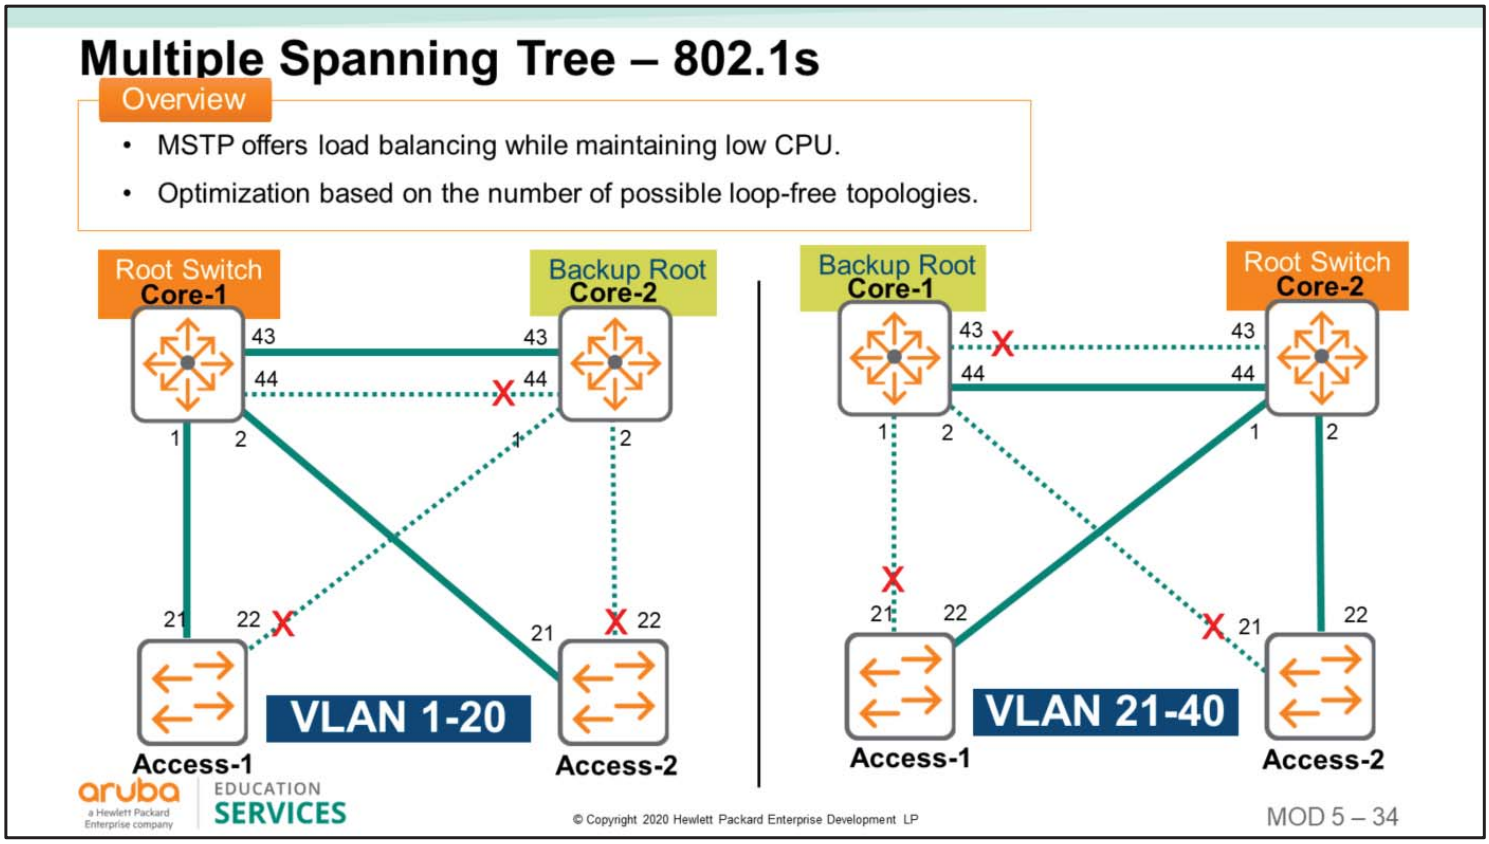

The solution to the previous problem (Single STP Topology) can be accomplished by IEEE 802.1S – the Multiple Spanning Tree Protocol (MSTP). MSTP enhances RSTP by building a new topology or instance of the spanning tree protocol. Instances then can carry different VLAN traffic allowing Load balancing and maintaining optimal paths.

The example in the figure illustrates this. The figure shows two topologies created on the same physical network. In the left-hand example, VLANs 1-20 all share a single iteration of RSTP. Core-1 is the root, and the tree grows from there. VLANs 1-20 share these forwarding links, with the unused links available for redundancy.

In the right-hand example, VLANs 21-40 share a different iteration of RSTP. Core-2 is the Root, and the topology emanates from that Root. The solid links are shared by VLANs 21-40, and the other links are available for backup. The dormant links for VLANS 1-20 are the active links for VLANs 21-40, and vice versa. Nice!

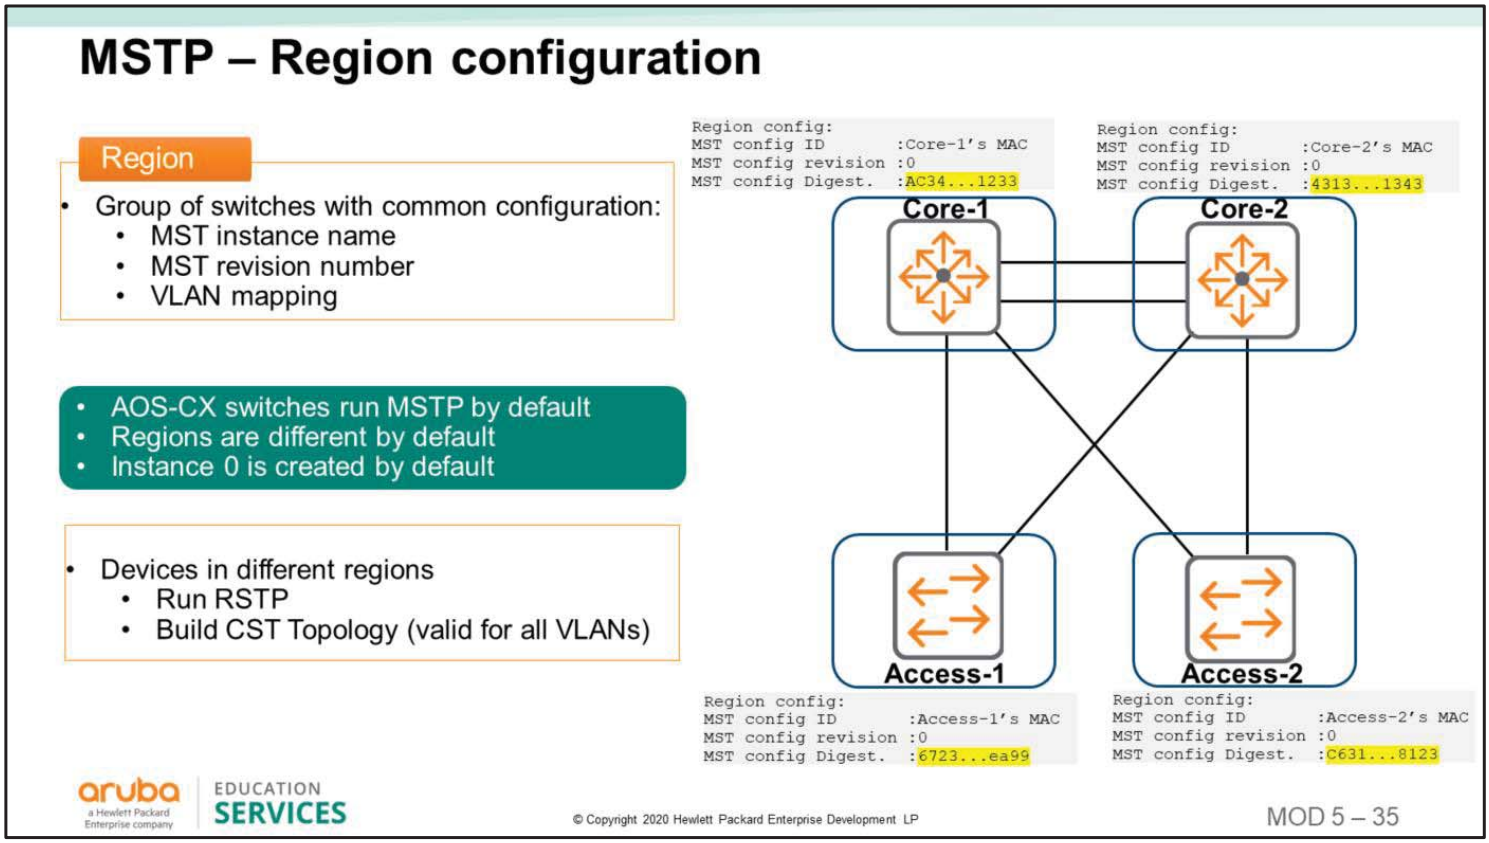

MSTP introduces the concept of a region - a group of switches that have been configured with the same instance, which is mapped or associated with a group of VLANs. Again, the figure shows an MSTP Region for Instance 1, which serves VLANs 1-2000.

To make sure all devices have mapped the same VLANs to the instance, MSTP uses a hash algorithm (sometimes referred as Digest).

Note. A hash algorithm is a mathematical function that is applied to an input (x). A hash function is interesting because from the output (y) you cannot derive or infer what the input was. Thus, it is referred to as a one-way function. Another characteristic of a hash is that if the same input (x) has been entered, the result will be always the same. MSTP takes advantage of this characteristic.

AOS-CX switches run MSTP by default, however each device has different configuration parameters therefore each device runs in its own region. When this happens, the topology behaves like the Common Spanning Tree (CST) creating a single topology or instance common for all VLANs. RSTP exchange messages between the individual devices, as you have learned in previous sections of this module.

In the figure, notice that each Switch has MSTP enabled but the Digest or hash is different therefore these devices are in different regions.

Instance 0 or Internal Spanning Tree (IST) is used as both: a regular instance in MST and the creation of the CST in a multi-region deployment for backward compatibility with RSTP speakers, for this reason Instance 0 is known as CIST (Common Internal Spanning Tree)

MSTP – Region configuration

Common configuration for all switches

```
(config)# spanning-tree config-name CXF
(config)# spanning-tree config-revision 1
(config)# spanning-tree instance 1 vlan 1-20
(config)# spanning-tree instance 2 vlan 21-40
```

Core-1 configuration

```
(config)# spanning-tree priority 1
(config)# spanning-tree instance 1 priority 1
(config)# spanning-tree instance 2 priority 2
```

Core-2 configuration

```
(config)# spanning-tree priority 2
(config)# spanning-tree instance 1 priority 2
(config)# spanning-tree instance 2 priority 1
```

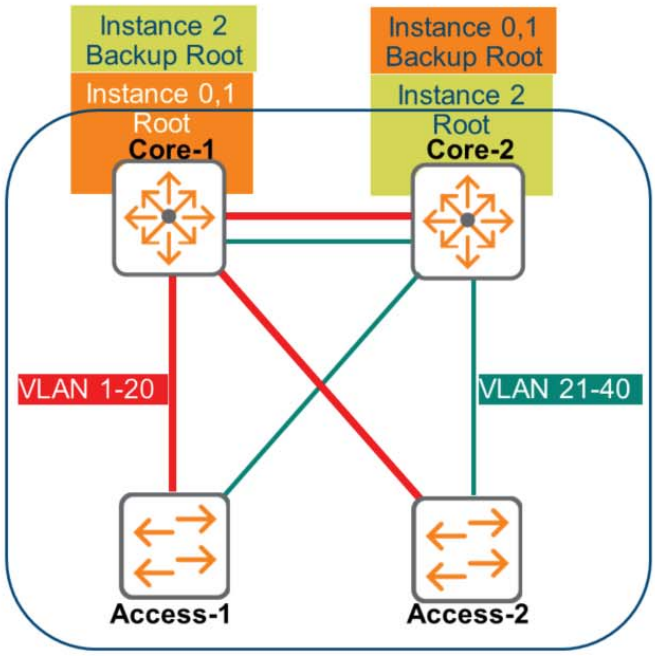

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 36

In the example, notice that instance 1 and 2 are created. Along with the default instance 0, you have a total of 3 instances. The common configuration creates a single region where you can create multiple instances - for optimal pathing and load balancing.
Instance 1 is mapped VLANs 1-20 and uses Core-1 as the root switch. Instance 2 is mapped VLANs 21-40 and uses Core-2 as the root switch.

In ArubaOS-CX, use the following steps to configure MSTP:

Configure the Spanning Tree mode as MSTP.

The command spanning-tree mode sets the mode to either MSTP or Rapid Per VLAN Spanning Tree (RPVST).

Switch(config)# spanning-tree mode mstp

Note. To change the spanning-tree mode, you must first disable spanning-tree with the command no spanning-tree in the global context.

Set up the MST region name.

The spanning-tree config-name command sets the configuration name for the MST region in which the switch resides.

All switches within an MST region must have identical configuration names, note that the name that you provide is case sensitive.

Switch(config)# spanning-tree config-name <Config-name>

Set up the configuration revision number.

Use the revision number to differentiate between region configurations. This is valuable, for example, when changing configuration settings within a region, and you want to track the configuration version you use. It is also valuable when you create a new region from a subset of switches in a current region and you want to maintain the same region name.

All switches within an MST region must have identical revision numbers.

Switch(config)# spanning-tree config-revision <revision-number>

Set up the MSTP instance and map the VLAN IDs to it.

The spanning-tree instance command creates a new instance with VLAN mapped to an existing instance.

Each instance must have at least one VLAN mapped to it. When VLANs are mapped to an instance, they are automatically unmapped from the instance they were mapped before. Any MSTP instance can have all the VLANs configured on the switch.

Switch(config)# spanning-tree instance <instance-id> vlan <vlan-ids>

Optionally, set the switch priority for an instance.

Switch(config)# spanning-tree instance <instance-id> priority <priority-value>

The same priority values can be used as in RSTP, the range to set up is 0-15.

Enable Spanning Tree globally.

Switch(config)# spanning-tree

Lab Activity

5.1 Configuring Rapid Spanning-Tree Protocol
5.2 Deploying MSTP
5.3 Explore broadcast storm effects (optional)

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Additional Content

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

RSTP Proposal and Agreement Sequence

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

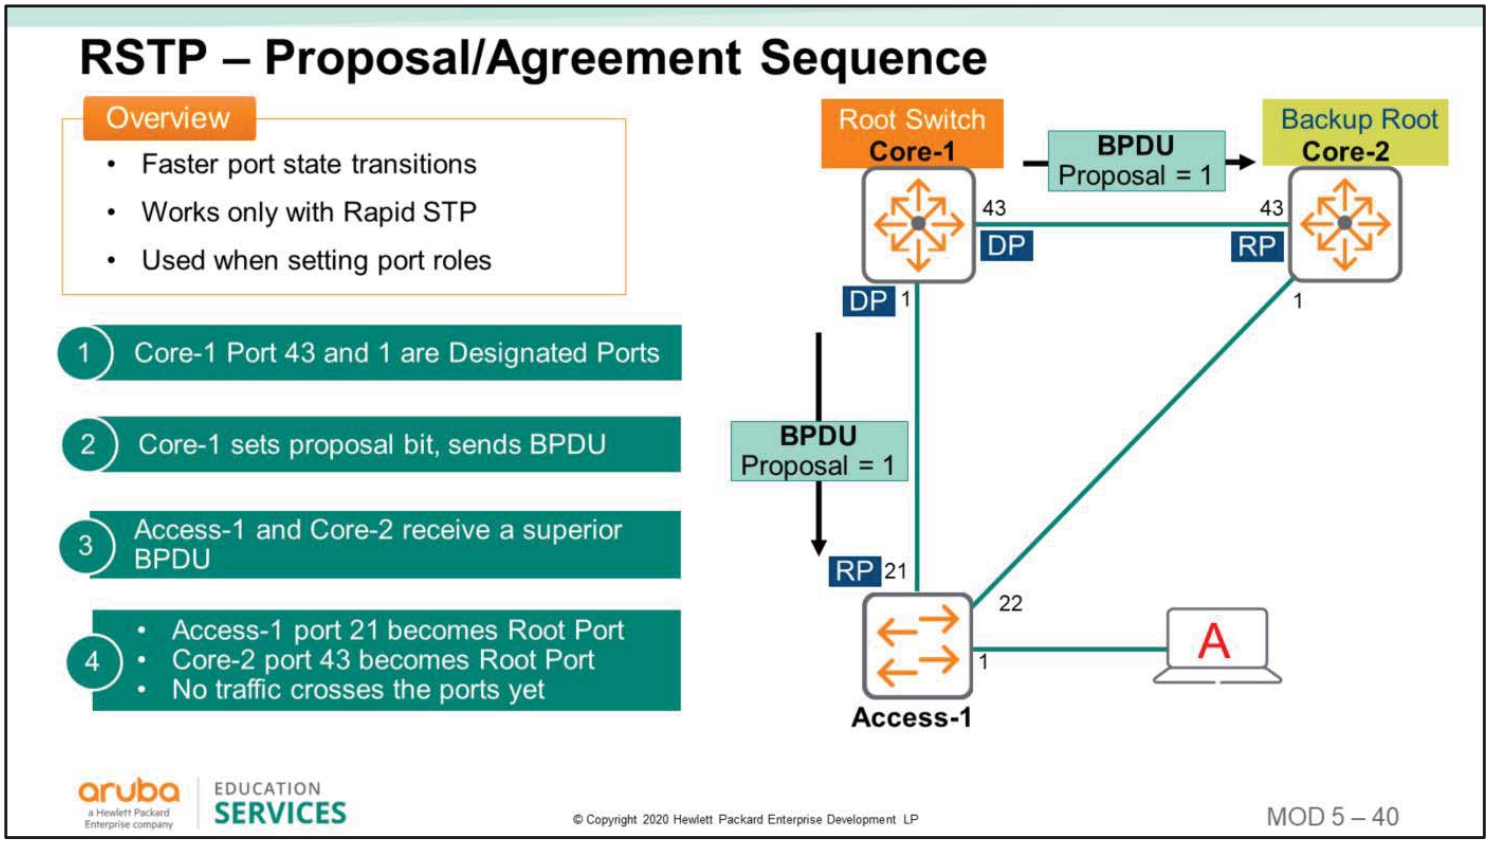

Consider the original 802.1d standard. When a port is selected as a designated port the transition to the forwarding state is not immediate. The port waits until the forwarding delay timer expires two times - one time for the Listening port state and another on the Learning state. That is 30 seconds by default.

For faster transitions, Rapid Spanning Tree relies on a proposal-agreement mechanism, which takes place during the election of port roles. Here's how it works:

When the root switch (Core-1) elects a Designated Port, that port is in Discarding or Learning state.
The switch sets a proposal bit in the BPDU and sends it out port 1 and on port 43 - a Designated Port. At this moment, this port does not forward any traffic.
Access-1 switch receives the BPDU. It understands that this is a superior BPDU - the BPDU's Bridge ID a better (lower) priority value.
Access-1 port 21 becomes the new root port but no traffic will cross the port until the synchronization process concludes

Note. To simplify the explanation, the scenario will consider only three switches as shown in the topology.

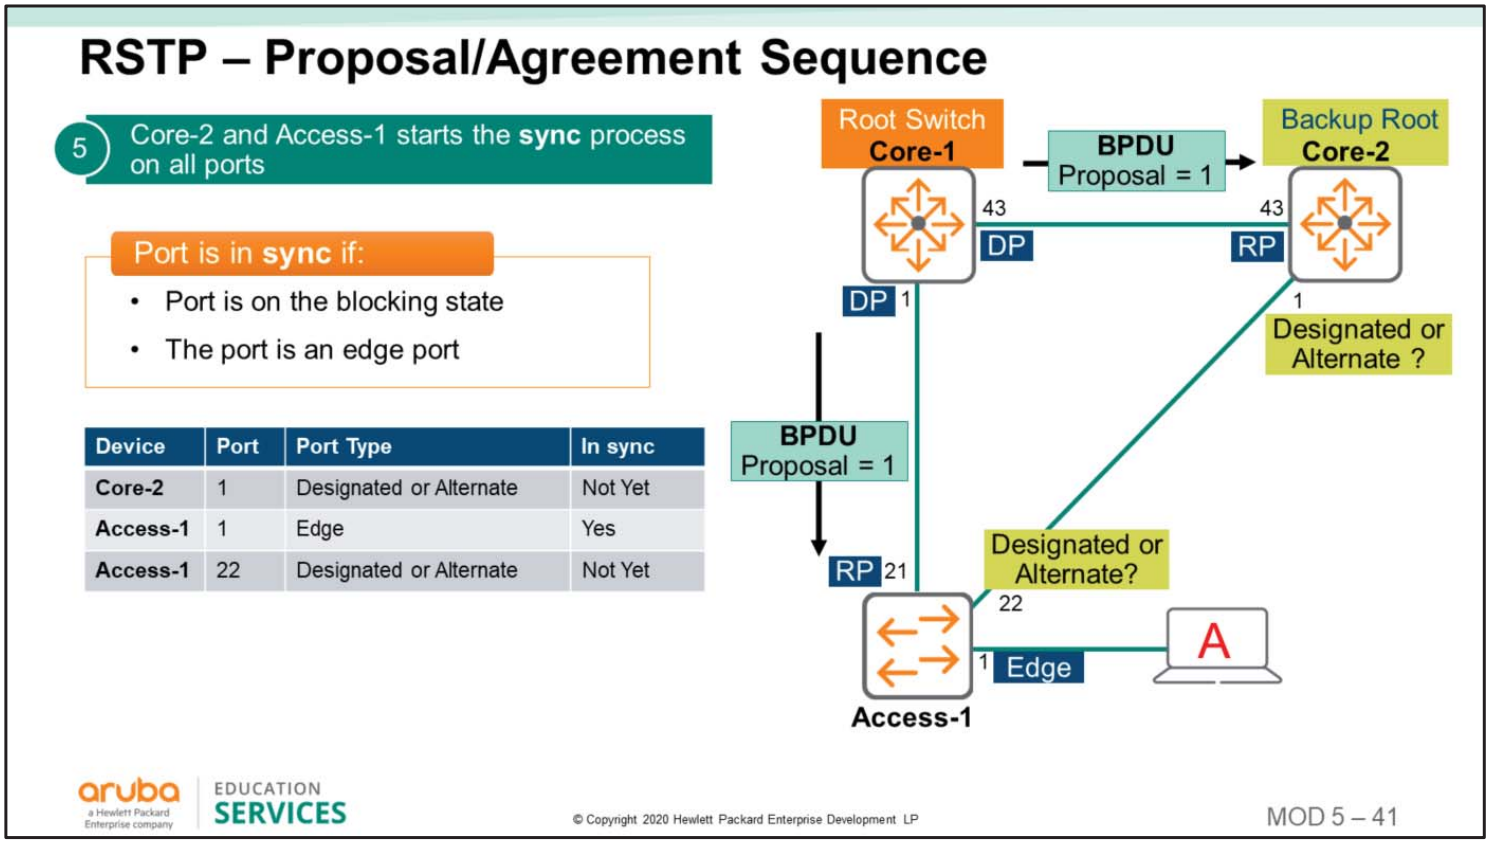

Core-2 starts a synchronization process on all its ports.

A port is in sync if it meets either of these criteria:
The port is on the blocking state.
The port is an edge port.

|In the figure you can conclude the following:
Core-2 port 1 does not offer the best path to the root. This port could either become a Designated Port or an Alternate Port

- If the port is an Alternate port, then this port by default is in a blocking state which means that the port is in Sync.
- If the port is a Designated port, then Core-2 forces the port into the discarding state, which means that the port is in sync.

Access-1 port 22 does not offer the best path to the root. This port could either become a Designated Port or an Alternate Port

- If the port is an Alternate port, then this port by default is in a blocking state which means that the port is in Sync.
- If the port is a Designated port, then Core-2 forces the port into the discarding state, which means that the port is in sync.

Access-1 port 1 is an edge port which implies that this port is already in Sync.

Note. By default a Designated port is in a forwarding state.

RSTP – Proposal/Agreement Sequence

6 Core-2 and Access-1 exchange BPDUs

Device	Port	Port Type	In sync
Core-2	1	Designated	Not Yet
Access-1	22	Alternate	Yes

7 Core-2 and Access-1 sends agreement to Root switch

8 Core-1 Port 1 and 43 moves to **Forwarding** state

9 Core-2 repeats the process and sends a proposal BPDU

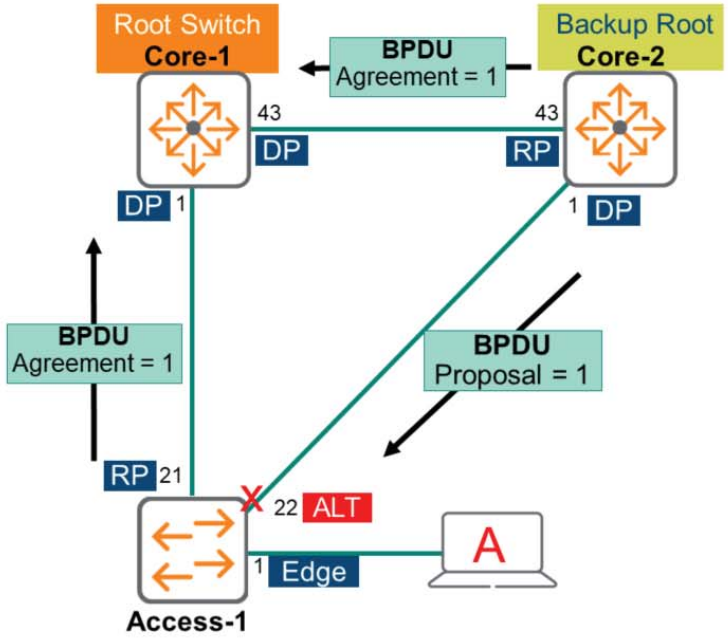

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 42

Core-2 and Access-1 exchange BPDUs and since Core-2 has a superior BPDU its port becomes the Designated Port.
Access-1 sends an agreement message to the Root switch. This message is a copy of the proposal BPDU with the agreement bit set. Thus, the root switch knows that this agreement correspond to a specific proposal.
The root switch receives the agreement BPDU and immediately moves port 1 to a forwarding state.
Core-2 port 22 is a Designated port, then this switch repeats the same process and sends a Proposal BPDU.

RSTP – Proposal/Agreement Sequence

10 Access-1 port 22 is in sync

Device	Port	Port Type	In sync
Access-1	22	Alternate	Yes

11 Access-1 sends an agreement to Core-2

12 Core-2 moves port 1 to **Forwarding** state

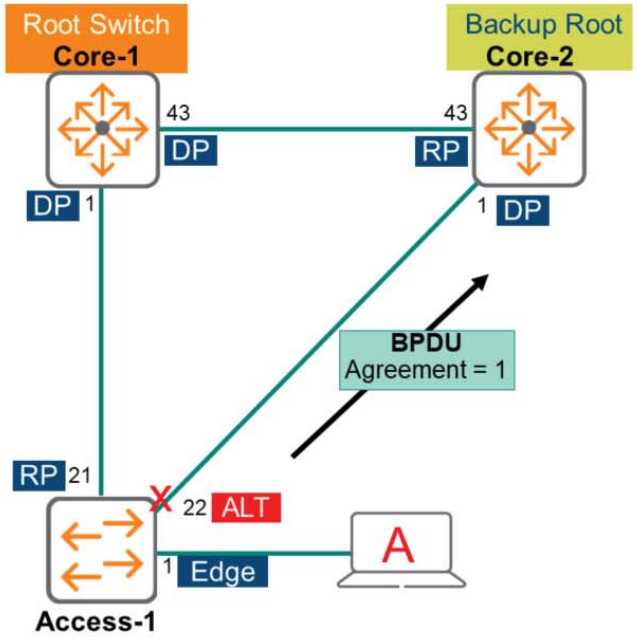

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 43

Access-1 port 22 is already in sync since is an Alternate port
Access-1 sends an agreement to Core-2
Core-2 moves port 1 to forwarding mode

Rate Limit

Goal

- Specify the amount of traffic that can ingress to a port.
- Helps to protect CPU utilization.
- Independent feature from STP
- Type of traffic supported:
  - Broadcast
  - Multicast
  - Unknown Unicast

Configuration

```
switch(config)# interface 1/1/1
switch(config-if)# rate-limit broadcast 4000 pps
```

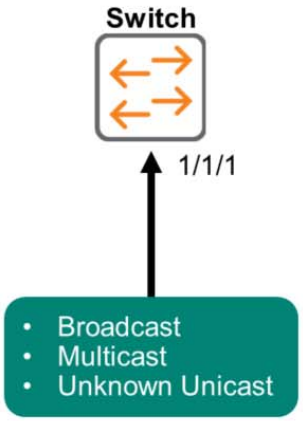

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 44

A misconfiguration in Spanning-Tree can lead to a Layer2 Loop creation. In these situations, AOS-CX helps to specify the amount of Broadcast, Multicast or Unknown Unicast traffic can be accepted to ingress on an Ethernet port or LAG.

This feature is independent of Spanning-Tree but is a mechanism to protect the CPU utilization and make the switch always reachable to the administrator.

The command must be applied in the interface level.

Note. You will learn about LAG ports in the next module.

Loop protection

Problem

- Edge ports are unprotected and could create loops

Loop protection

- Mechanism enabled on the port or LAG.
- Sends loop protection packets (5 seconds by default)
- When a loop is detected:
  - Only trigger a SNMP notification
  - Transmission is disabled on the port
  - Transmission and reception is disabled on the port

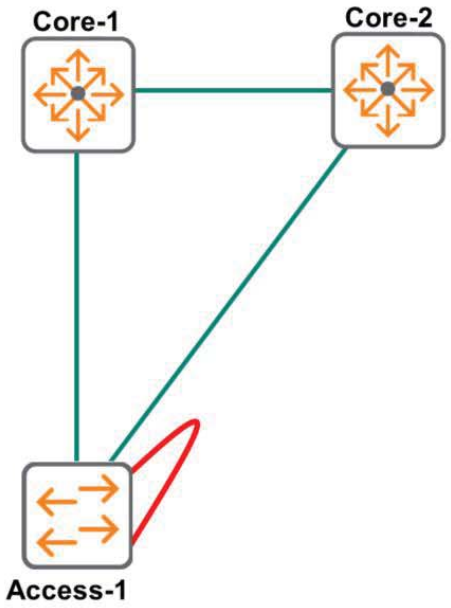

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 5 – 45

So far, the focus has been on preventing loops on ports that are intentionally redundant – by design. However, there are situations when a loop can be created on ports that were not intended to be connected. This most commonly occurs on switch ports that face users - edge ports. For example, a user connects two Ethernet jacks on the same switch together, as shown in the figure.

The loop protection feature protects edge switches from downstream loops by transmitting loop protocol packets out of the ports. The enabled ports automatically send loop protection packets at specified intervals (5 seconds by default). When a loop-protection-enabled port receives a loop protocol packet, the port knows there is a downstream loop, and so takes certain remedial actions.

The actions supported on AOS-CX are:

Do not disable. Only a notification is triggered and reported using SNMP.
Transmission disabled. The port that transmitted the loop detection packet is disabled. This is the default action.
Transmission and Reception disabled.

In AOS-CX this feature can be enabled on layer 2 interface or LAG

AOS-CX configuration:

```
switch# config
switch(config)# interface 1/1/1
switch(config-if)# loop-protect
```

switch(config-if)# loop-protect action {do-not-disable | tx-disabled | tx-rx-disabled}

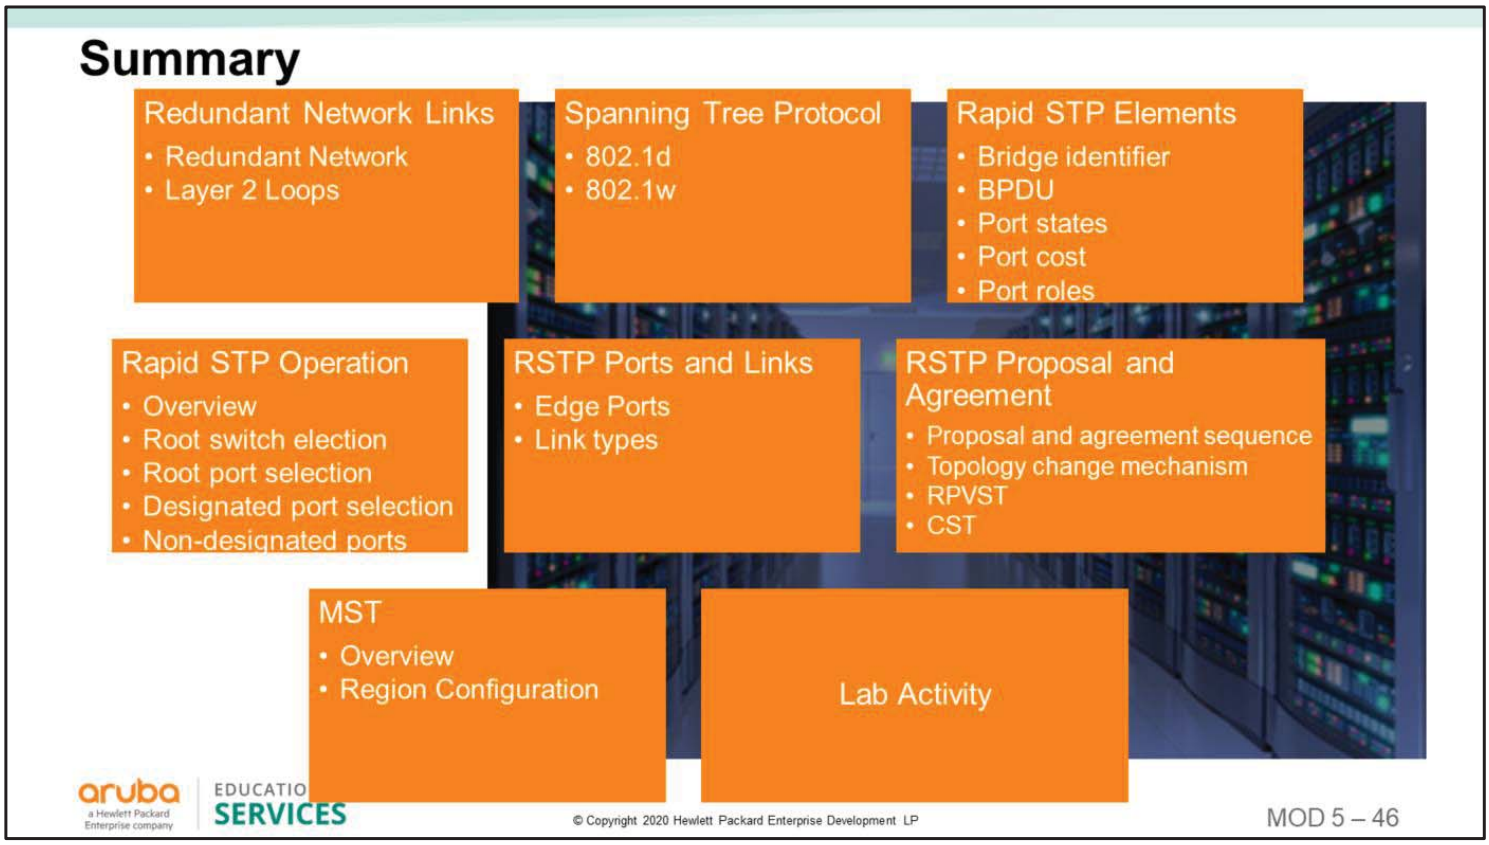

You explored a vital aspect of designing Layer 2 switched networks, especially as relates to resiliency, path optimization, and network efficiency – the Spanning-Tree protocols.

You began by learning how single points of failures can be mitigated by connecting redundant switches and redundant links. However, this causes the network-breaking problem of Layer 2 loops – broadcast storms and MAC table instability.
You learned to resolve these issues with the 802.1d Spanning-Tree Protocol (STP). But this is too slow for modern networks, so you learned to upgrade to the newer IEEE 802.1w standard – the Rapid Spanning-Tree Protocol (RSTP). You learned about RSTP elements, and how they work together to create a functional system that provides layer 2 resiliency, while avoiding loops.

You went deeper into RSTP operation. You explored how edge ports and link types, as well as the RSTP proposal and agreement process can be leveraged to further increase the efficiency and uptime of your systems.
As always, this new RSTP solution brings new problems and challenges. In large networks, this solution can waste bandwidth and CPU resources, potentially preventing your network from reaching its true potential.

You learned how to resolve this issue by deploying the IEEE 802.1s standard Multiple Spanning Tree (MST) solution.

Finally, you engaged in some hands-on lab activities to solidify your knowledge.

aruba
a Hewlett Packard Enterprise company
Link Aggregation
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 6 – 1

Objectives

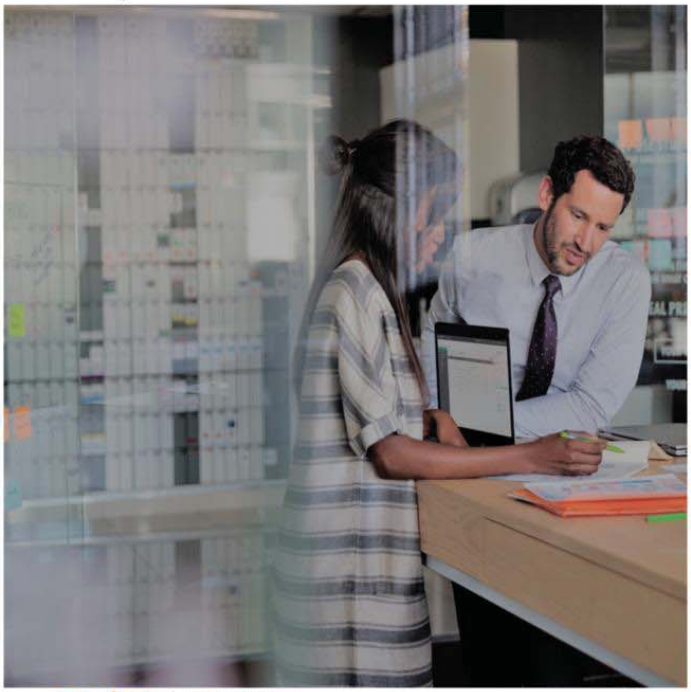

–Learning Objectives

- Describe LAG and its interface requirements
- Explain, configure, and verify LAG
- Describe, configure, and verify load balancing

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 6 – 2

After completing this module, you will be able to describe how Link Aggregation Groups (LAG) can improve performance and resiliency for switch-to-switch links. You will be able to explain, configure, and verify LAG.
You will also be able to describe, configure, and verify LAG load balancing.

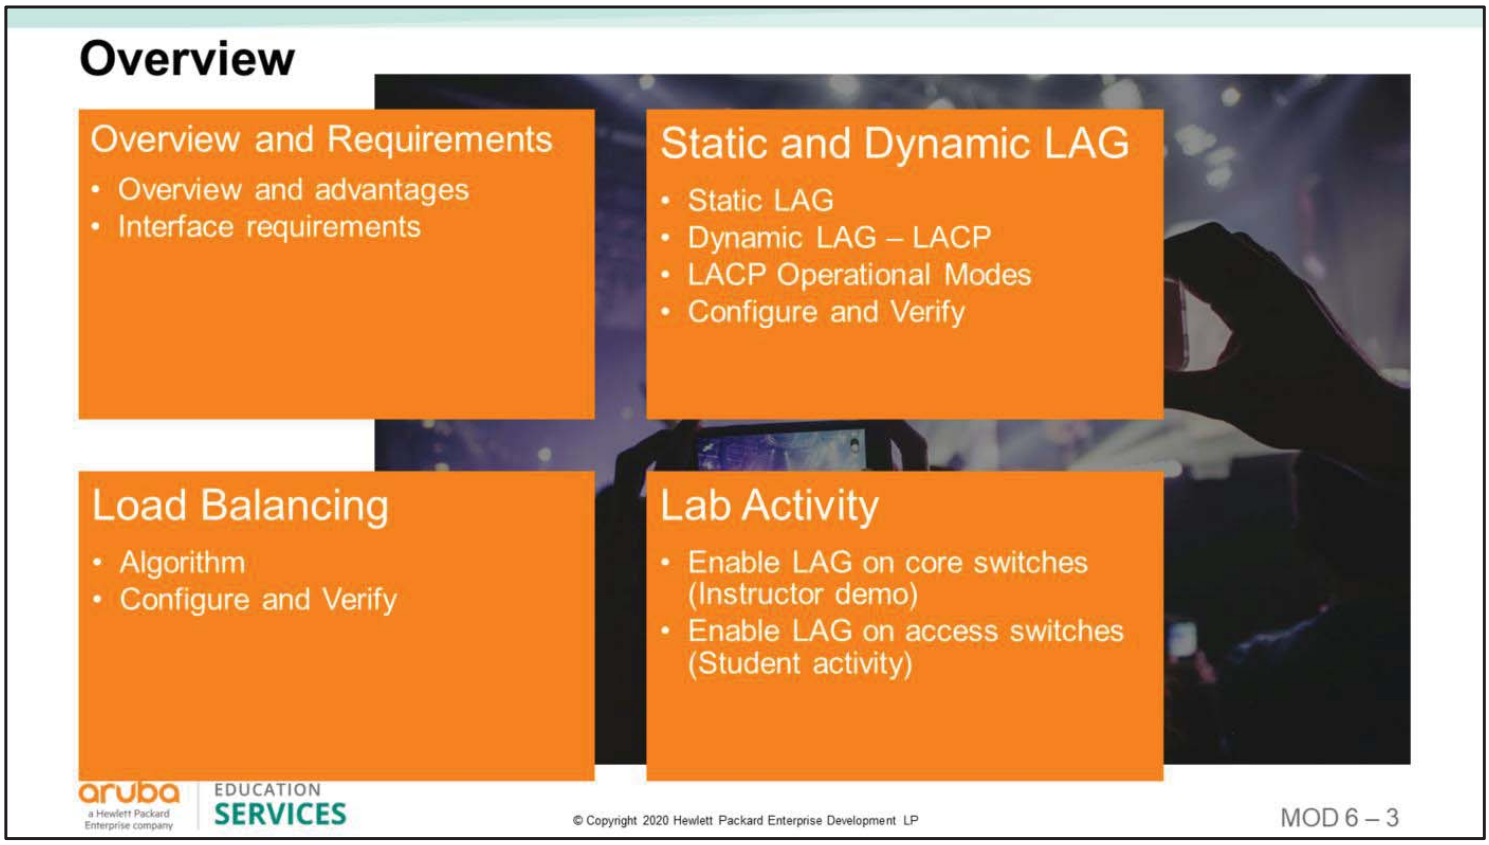

Switch-to-switch links are busy links! Without knowledge of LAG, you could easily oversubscribe these links – poor performance with no resiliency.
You will explore Link Aggregation Group advantages and requirements, before learning about the difference between static and dynamic LAG, which relies on the Link Aggregation Control Protocol (LACP). You learn about LACP operation modes and then apply that knowledge to configure and verify it.
You will learn about load balancing algorithms and inputs, and then learn how to configure and verify LAG load balancing.

Overview and Interface Requirements

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

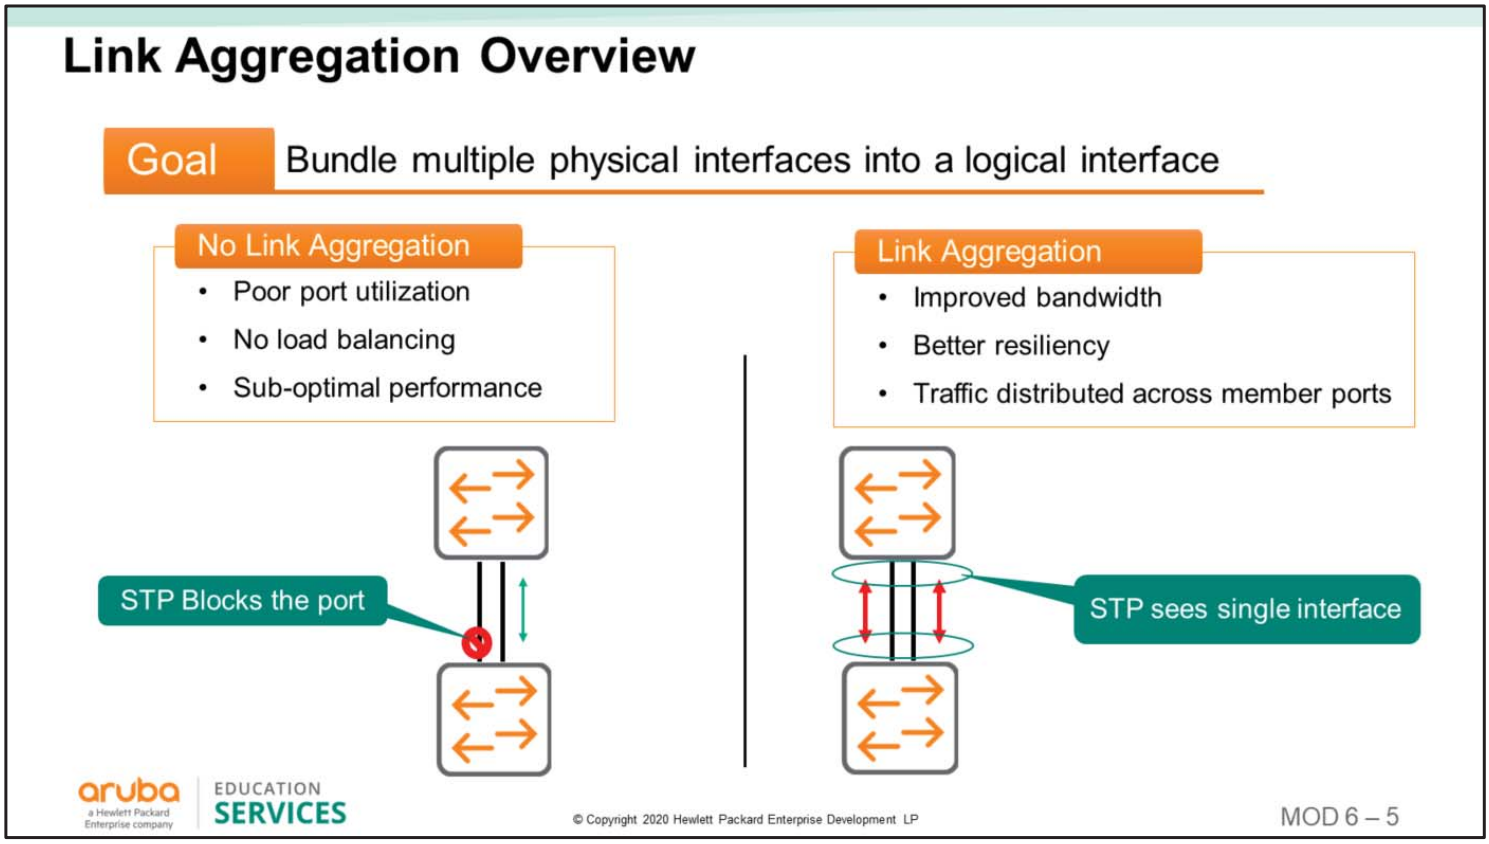

You know that switches provide connectivity for many devices, on many different VLANs. All that traffic must share the switch-to-switch links. So you add multiple switch interconnects, as shown in the left-hand example. While this does add some resiliency, it does not add additional bandwidth – the two links would create a loop, and so STP will automatically block one of the ports. You have two links, but only one of them are used – poor link utilization, no load balancing, and sub-optimal performance.

The solution is Link aggregation, which bundles multiple physical interfaces into a single logical interface, as shown in the figure's right-hand example. Since protocols like STP perceive this bundle as a single interface, there is no blocking. All switch interconnects carry traffic. You might go from one 10GBPS link to two, four or more bundled links. You get far more bandwidth, because traffic is distributed across member ports. You get better resiliency because if one member fails, remaining links carry the load. Convergence is faster than spanning tree, because there is no need for an STP Topology Change notification.

Note. Link aggregation can be used not only for switch-to-switch links, but it can be used on other links to servers and routers. This text however finds an easy way to introduce this topic by talking about switch-to-switch links that were discussed on previous modules.

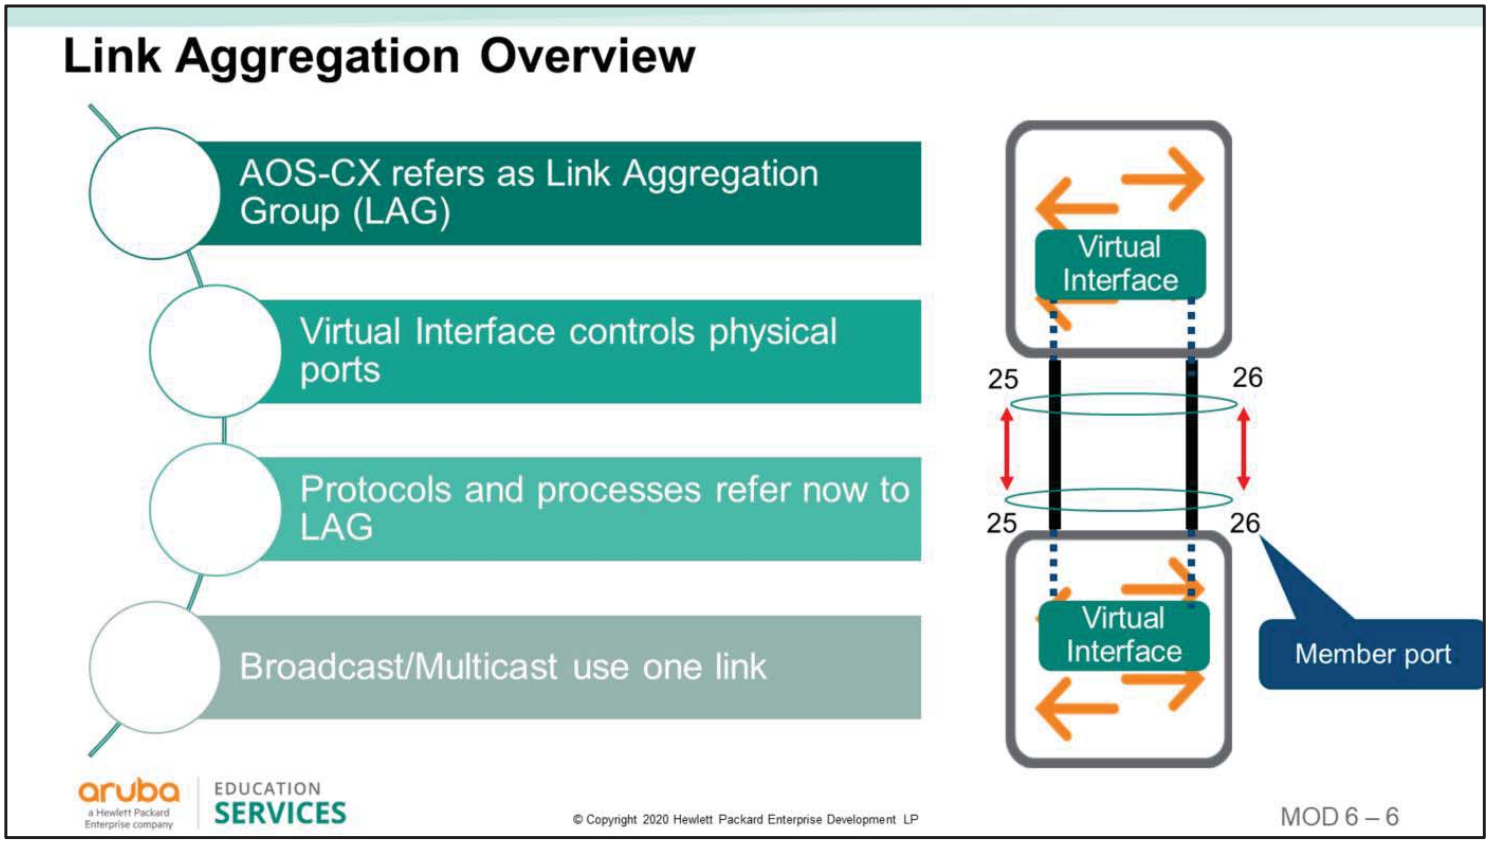

When you enable Link aggregation on a switch the protocol creates a virtual interface. You then configure physical ports to be members of that virtual interface. The switch's various protocols and processes will only perceive and refer to the virtual interface. They no longer perceive the individual physical interface members.

In AOS-CX the virtual interface is referred as LAG (Link Aggregation Group) and the interfaces are called member ports.

Is important to mention that broadcast and multicast traffic is sent across only one physical link in the bundle. This behavior ensures that Link Aggregation does not create a Layer 2 loop.

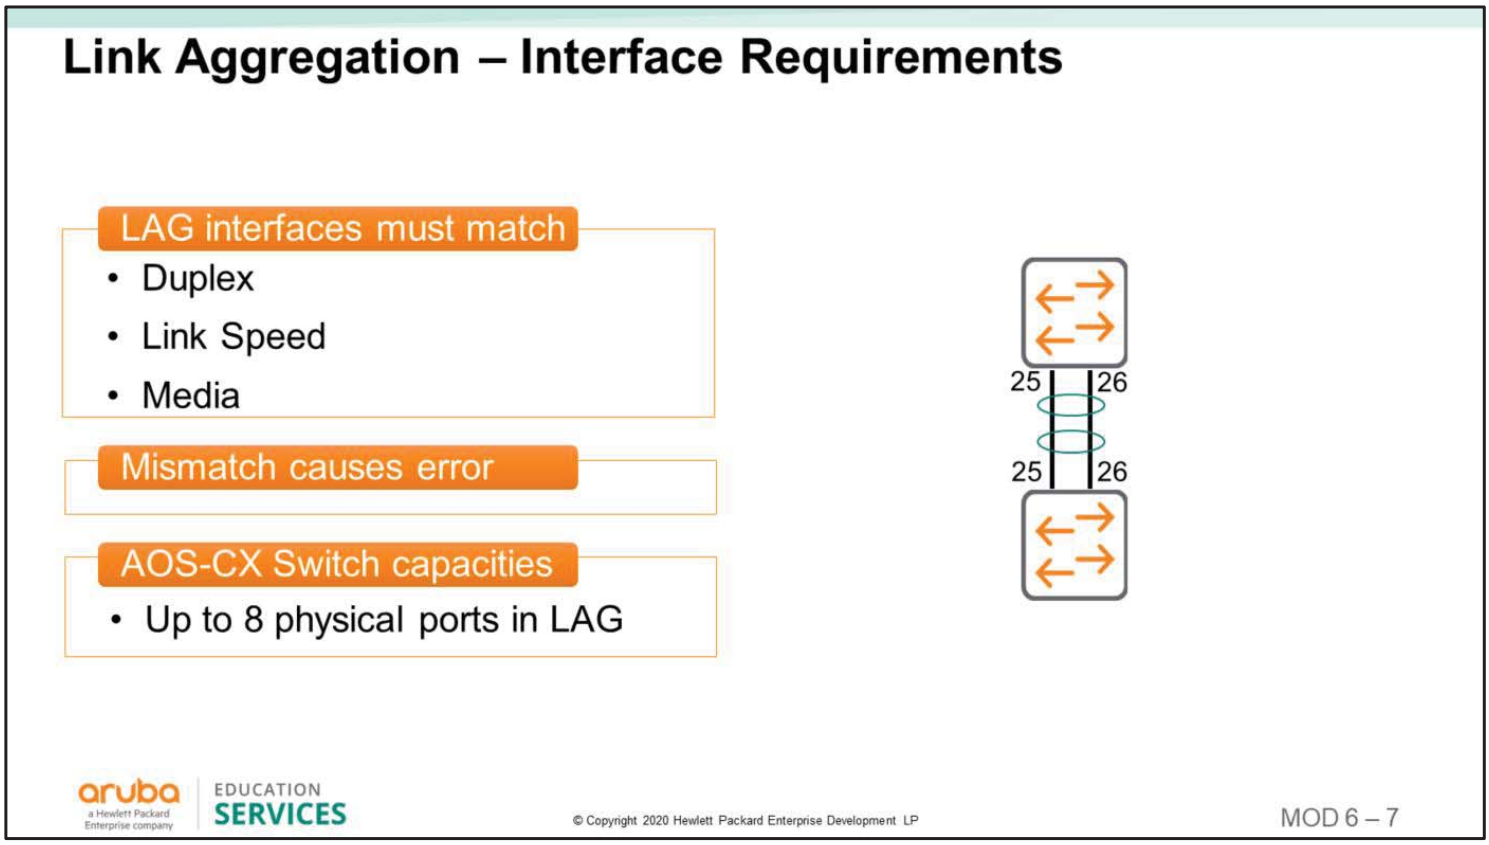

Requirements for interfaces

Interfaces that are mapped to the same Link Aggregation Group must be configured in a consistent manner. The following items must match:

Duplex mode (Full-duplex or Half-duplex)
Link speed
Media

AOS-CX displays a warning if you attempt to map mismatched interfaces to a LAG. For example, Interface 5 with speed of 10 Gb/s cannot be added to lag10 with base speed of 1 Gb/s.

Maximum number of ports

In AOS-CX each Link Aggregation Group can be have up to 8 individual ports. Use the command show capacities to verify your switch capacity.

Now that you understand LAG, you should learn about two operating modes - Static LAG and Dynamic LAG, which uses the Link Aggregation Control Protocol (LACP).

Static and Dynamic LAG

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

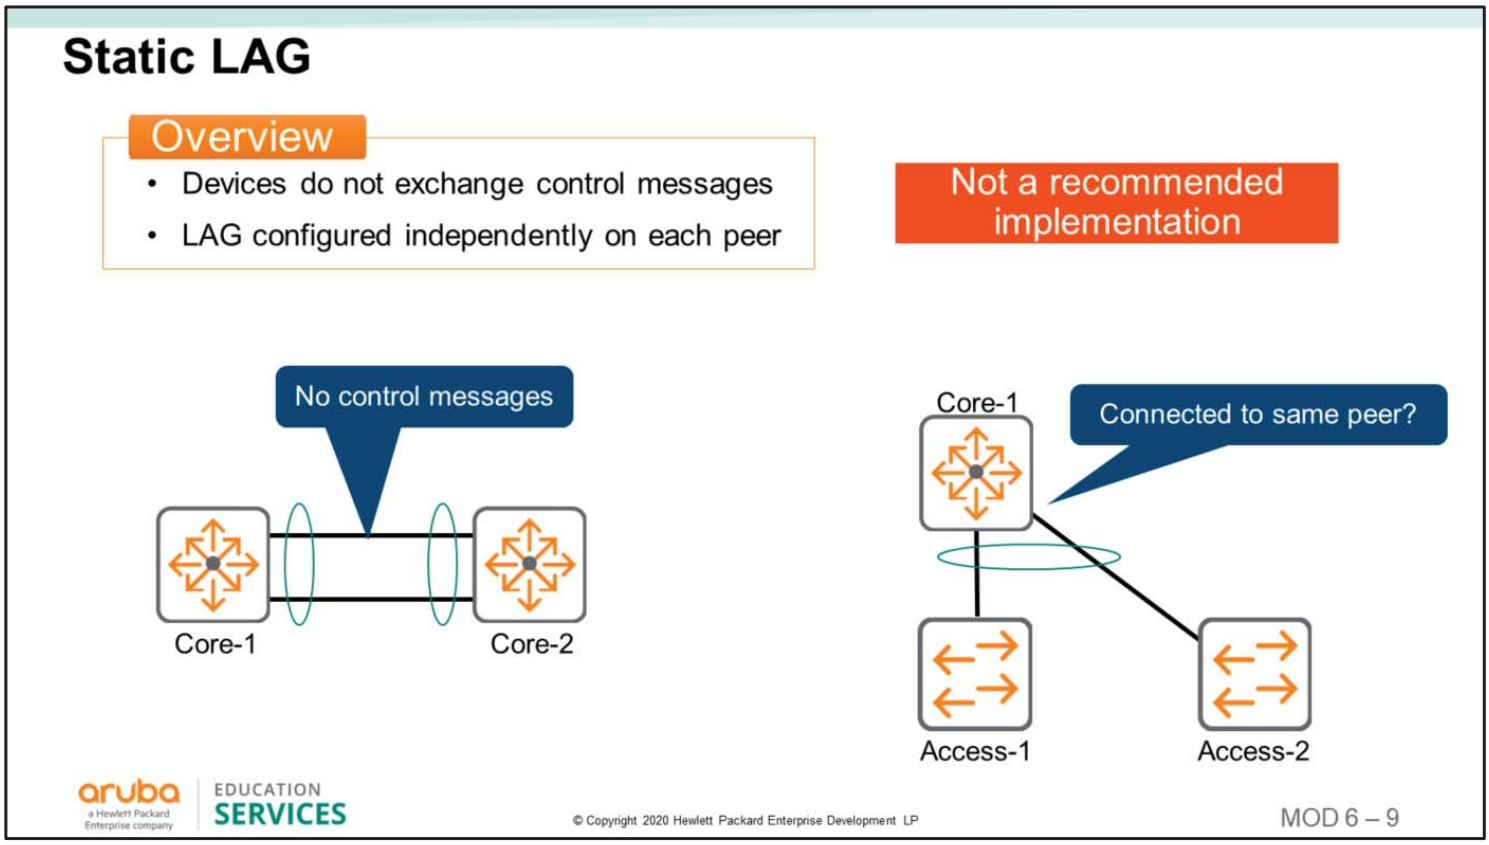

In Static Link Aggregation mode devices do not exchange any control information. There is no signaling between switch peers about LAG. You simply configure LAG independently on each peer. If your configuration is good, LAG works. However, the switch peers have no knowledge of who they are connected to, or whether they are connected to the same peer.

This mode is not recommended because a misconfiguration on one side is not detected by the peer. This can lead to unexpected behavior, which can be challenging to detect and troubleshoot. Another scenario that could cause problems is shown in the figure's right-hand example. Switch Access-1 is not aware that you have erroneously connected LAG member ports to two separate physical switches.

Configuring Layer 2- Static Link Aggregation Group

```
Switch(config)# interface lag <id>
Switch(config-lag-if)# no routing
Switch(config)# interface {member/slot/port}
Switch(config-if)# lag <id>
```

Create LAG interface
Disable routing on LAG
Map the interface to LAG

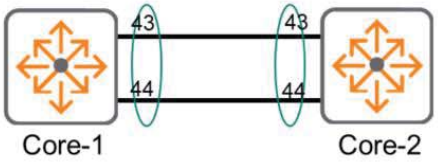

Configuration example

```
Core-1(config)# interface lag 20
Core-1(config-lag-if)# no routing
Core-1(config)# interface 1/1/43
Core-1(config-if)# lag 20
Core-1(config)# interface 1/1/44
Core-1(config-if)# lag 20
```

```
Core-2(config)# interface lag 20
Core-1(config-lag-if)# no routing
Core-2(config)# interface 1/1/43
Core-2(config-if)# lag 20
Core-2(config)# interface 1/1/44
Core-2(config-if)# lag 20
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 6 – 10

The following steps create a Layer 2 Dynamic Aggregation Group:

1. Create a LAG interface, represented by an identifier – any unused value between 1 to 256.

Switch(config)# interface lag <id>
2. Disable routing (Layer 3 capabilities) in the new LAG interface. This step limits the LAG to only process Layer 2 frame headers.

Switch(config-lag-if)# no routing

3. Map a port member to the LAG.

In the interface context level:

Switch(config)# interface 1/1/1
Switch(config-if)# lag <id>

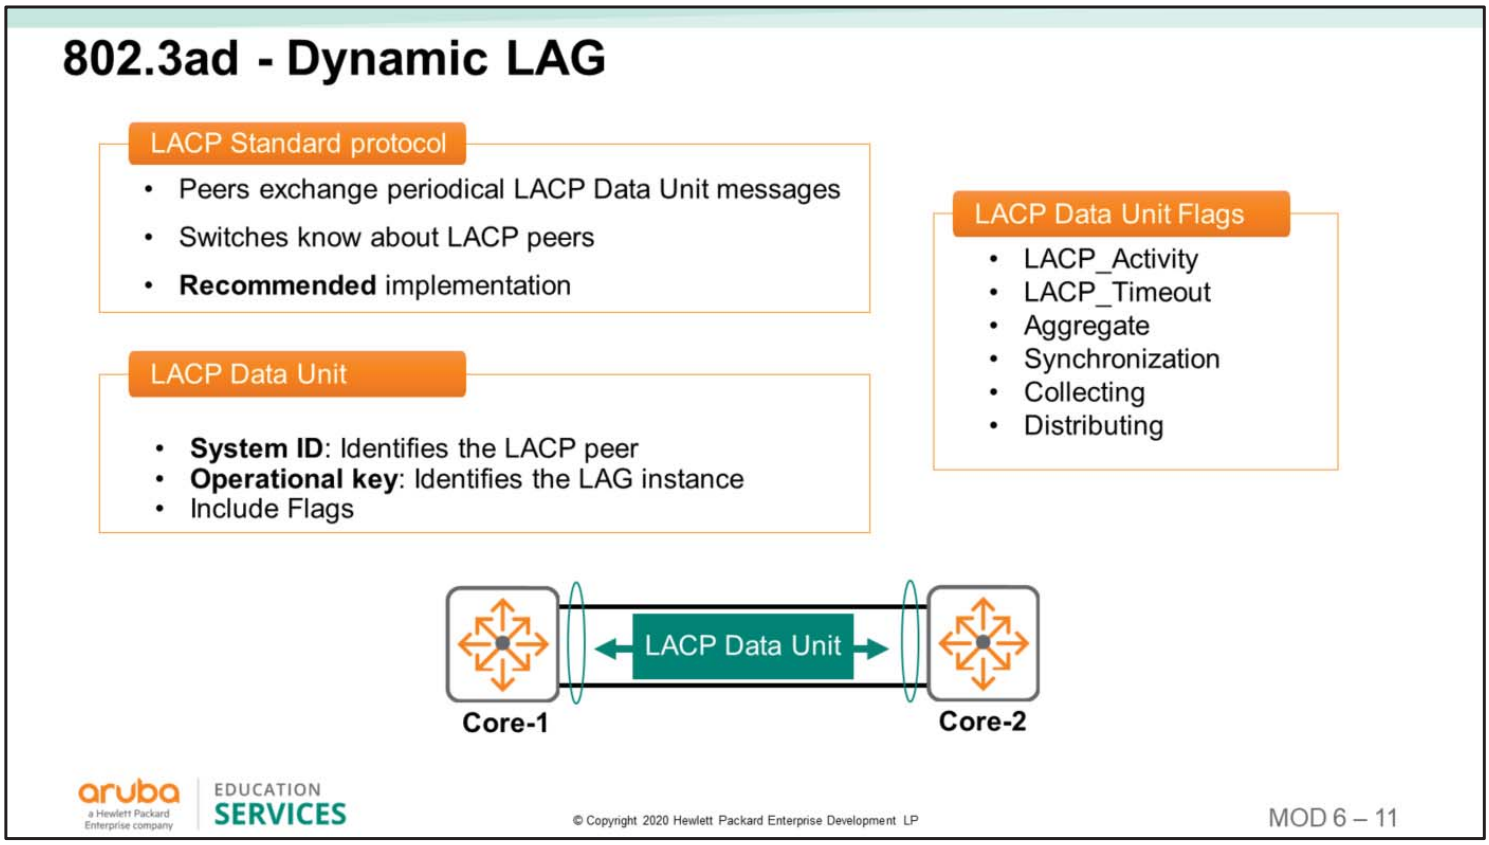

On 802.3ad peer devices that use Dynamic Link Aggregation exchange control messages to establish and maintain the LAG. This mechanism will also detect link failures and ensure that LAG port members terminate on the same device.

The standard is commonly known as Link Aggregation Control Protocol (LACP). LACP exchanges periodic messages called LACP Data Units. These messages include:

- **System ID**, which uniquely identifies the switch.
- **Operational key,** which uniquely identifies the Link aggregation group.

Dynamic LAG or LACP is the recommended method to implement Link Aggregation, to avoid unexpected network problems.

There are some flags that are included in the LACP-DU. The extra content section shows the meaning of these flags.

With 802.3ad standard you can bundle up to 16 ports but only 8 active ports can be used at any given time.

802.3ad Operational Modes

Passive	Device waits to receive a LACP Data Unit
Active	Device transmits LACP Data Unit

Device A port mode	Device B port mode	LAG formed?
Passive	**Passive**	**No**
Passive	Active	Yes
Active	Passive	
Active	Active	

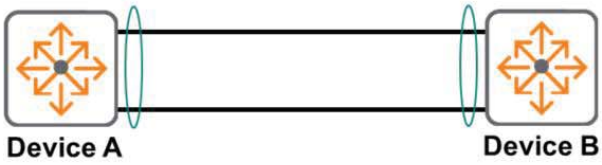

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 6 – 12

LACP can be configured in one of two modes, which controls how peer negation proceeds.

With Passive mode, the device passively waits to receive an LACP Data Unit message from the peer, to dynamically create the Link Aggregation Group. This mode places the LAG in a listening state. It is as if the switch is thinking, “I’ll just sit here and wait until I hear from my peer.” If you configure the peer switch to also be in passive mode, then it is thinking the same thing. Both switches wait to hear from the other, and so a functional LAG connection is never formed. At least one peer must be in Active mode.

With Active mode, the device actively transmits LACP Data Unit messages over its member ports, “Hey! Let’s form a LAG”. Whether the peer is Passive or Active, it responds, negotiation continues, and the LAG successfully forms. This is shown in the figure.

Let’s talk about configuration.

Configuring Layer 2- Dynamic Aggregation Group

```
Switch(config)# interface lag <id>
Switch(config-lag-if)# no routing
Switch(config-lag-if)# lacp mode {active | passive}
Switch(config)# interface {member/slot/port}
Switch(config-if)# lag <id>
```

Create LAG interface

Disable routing on LAG

LACP Mode

Map the interface to LAG

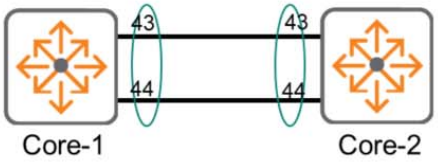

Configuration example

```
Core-1(config)# interface lag 20
Core-1(config-lag-if)# no routing
Core-1(config-lag-if)# lacp mode active
Core-1(config)# interface 1/1/43
Core-1(config-if)# lag 20
Core-1(config)# interface 1/1/44
Core-1(config-if)# lag 20
```

```
Core-2(config)# interface lag 20
Core-2(config-lag-if)# no routing
Core-2(config-lag-if)# lacp mode active
Core-2(config)# interface 1/1/43
Core-2(config-if)# lag 20
Core-2(config)# interface 1/1/44
Core-2(config-if)# lag 20
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 6 – 13

The configuration for Dynamic LAG is the same as static except for and extra line where you need to specify the LACP mode, couldn't be simpler?

Verifying LAG Interface

```
Switch# show interface lag

 Aggregate lag1 is up
 Admin state is up
 Description :
 MAC Address                 : 88:3a:30:97:59:c0
 Aggregated-interfaces       : 1/1/25 1/1/26 2/1/25 2/1/26
 Aggregation-key             : 1
 Aggregate mode              : active
 Speed                       : 40000 Mb/s
 L3 Counters: Rx Disabled, Tx Disabled
 <<<<Omitted output>>>>
```

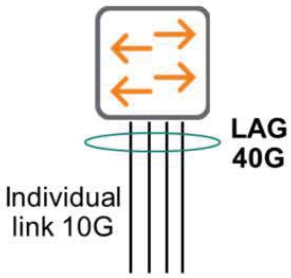


```
Switch# show interface lag brief
--------------------------------------------------------------------------------
Port      Native  Mode   Type        Enabled Status  Reason          Speed
          VLAN                                                       (Mb/s)
--------------------------------------------------------------------------------
lag1      1       trunk  --          yes     up      --              40000
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 6 – 14

Use show interface lag to verify your efforts. You can also use show interface lag brief to see the LAG's total available bandwidth.

Load Sharing

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

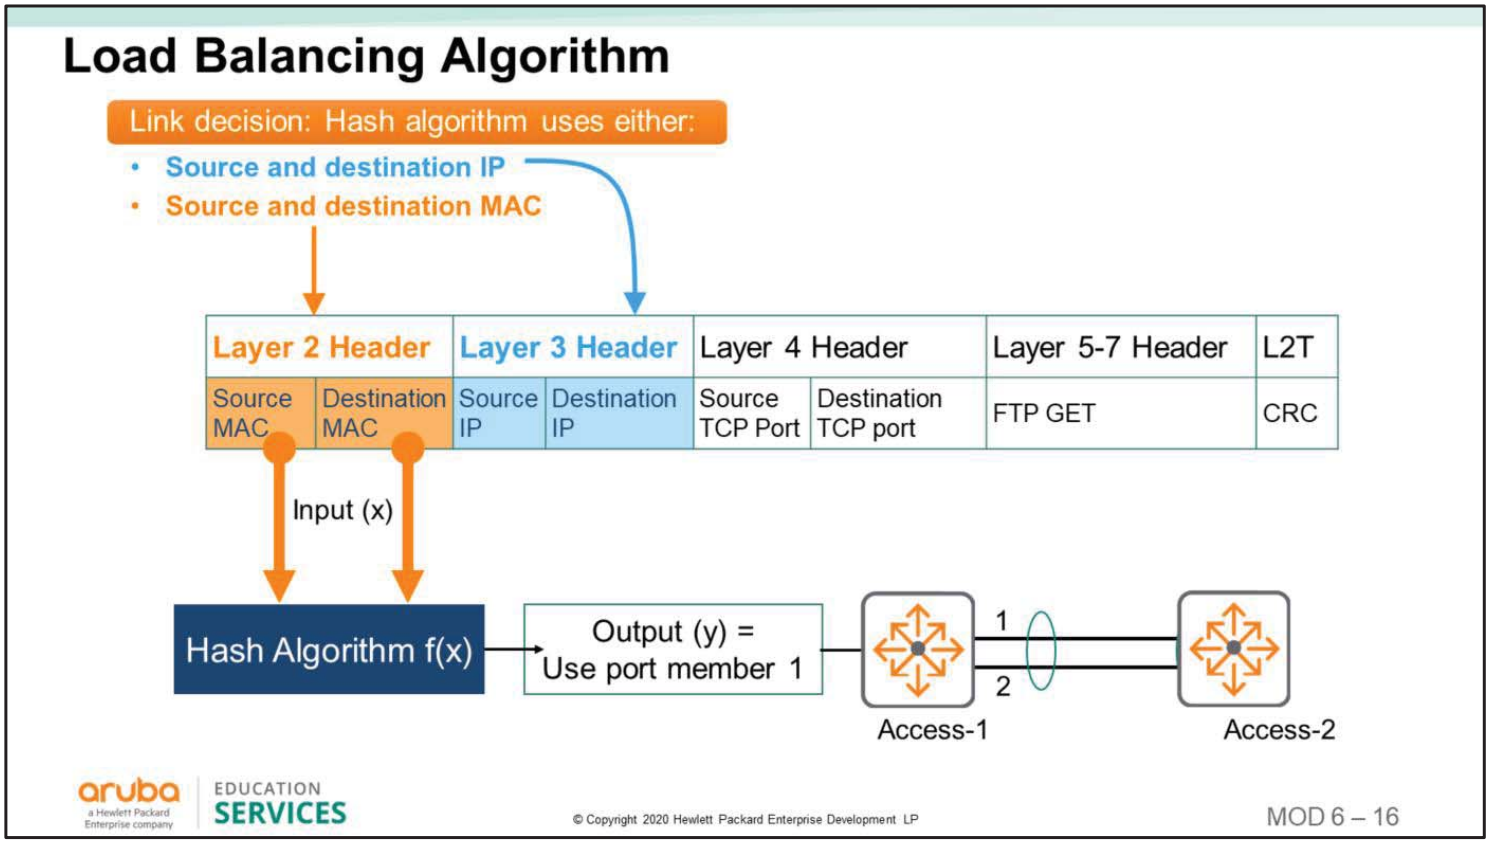

Peer LAG devices must balance the traffic load across multiple physical interfaces. How does the switch decide which member port should be used for a given packet? It uses a Hash algorithm.

A hash algorithm is a mathematical function that is applied to an input (x). Interestingly, if you see the output (y) of this function, you cannot derive or infer input (x). Therefore it is referred to as a one-way function. Another characteristic of a hash is that if the same input (x) has been entered, the result will be always the same. So, what does the switch use as inputs to this function?

The switch uses packet header information as hash function inputs, and the result is the port member to be used for that particular packet. Depending on switch algorithm used, inputs for the hash could be:

Layer 4 TCP/UDP ports
Layer 3 Source and Destination IP addresses
Layer 2 Source and Destination MAC addresses

The figure shows how the Layer 2 source and destination MAC addresses are taken as inputs for the algorithm. The output is to use physical port 1 for all packets with this source and destination MAC address combination - always. Only a link failure will cause port 2 to be used for these packets.

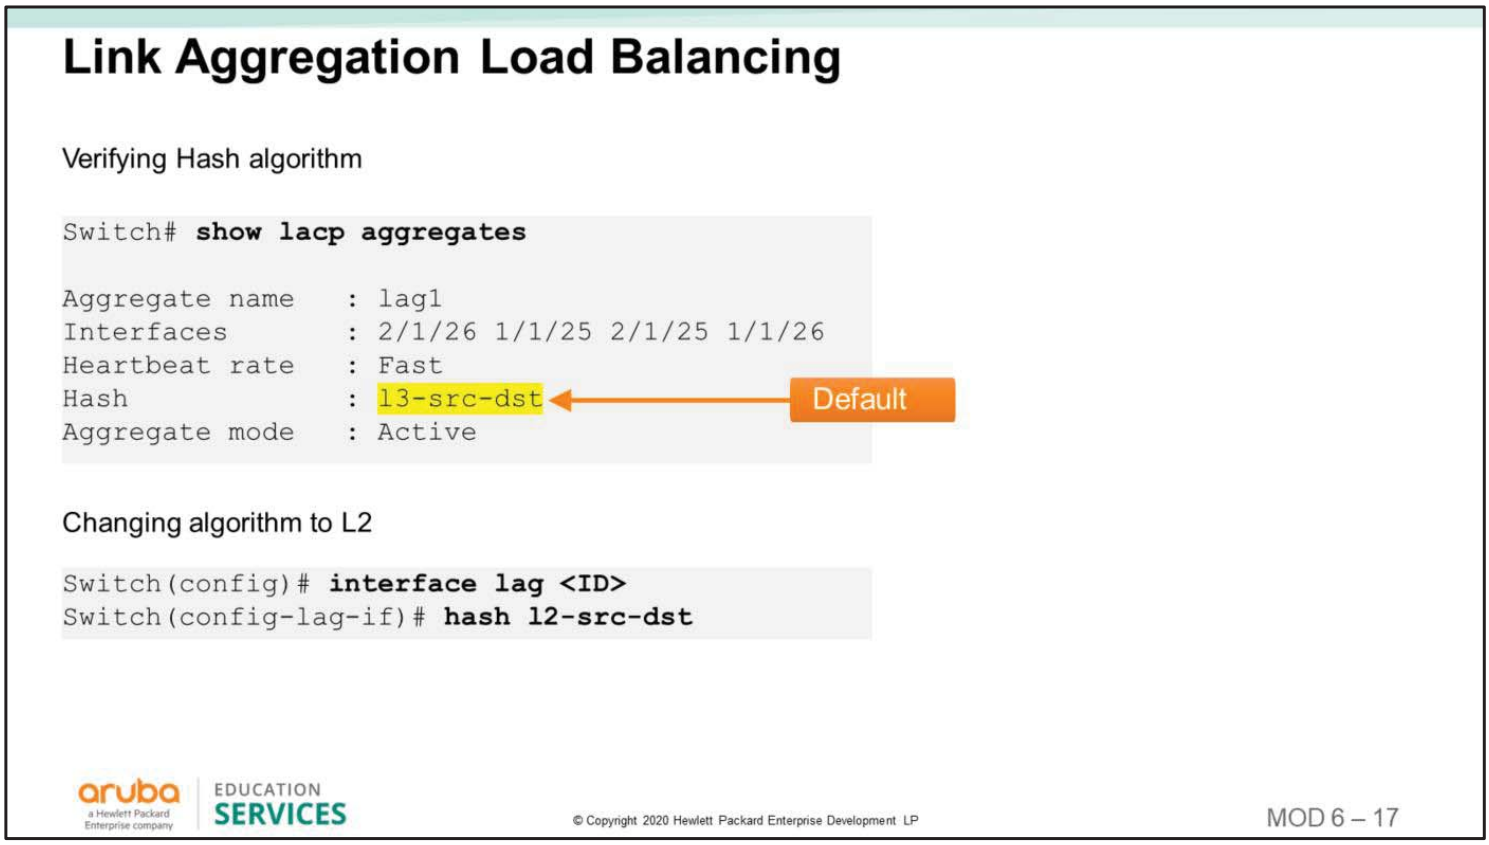

In AOS-CX the default input information for the hash algorithm are the Source and Destination IP addresses. You can verify with the show lacp aggregates command. Notice that the highlighted hash information is “l3-src-dst”. You know the “L3” means Layer 3, which is IP addresses.

In some cases, the hash algorithm’s use of source/destination IP addresses may not properly distribute the load equally among port members. This could lead to traffic congestion and poor performance.

In that case, you can modify the hash input data. As shown in the figure, you have decided to use Layer 2 source and destination MAC addresses as inputs to the hash function. This might help to avoid port member oversubscription.

Lab Activity

6.1 Link Aggregation between Core switches (instructor demo)
6.2 Deploying Link Aggregation using LACP (optional)

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

LACP Data Unit Flags

Flag	Meaning
LACP_Activity	Indicates a participant's intent to detect or maintain the LAG • Set = Active LACP • Not Set = Passive LACP
LACP_Timeout	The participant wishes to receive periodic transmissions. • Set = Short timeout • Not Set = Long Timeout
Aggregate	The participant allow the link to be used as part of LAG
Synchronization	The participant multiplexor component is in sync with the system ID and key information
Collecting	The reception component of the multiplexor is on
Distributing	Participant distributor is not off

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 6 – 19

Table shows the flags used in the LACP-DU and their meaning.

Congratulations! Now you know how to leverage LAG to avoid poorly performing switch-to-switch links.

You explored Link Aggregation Group advantages and requirements, before learning about the difference between static and dynamic LAG, which relies on the Link Aggregation Control Protocol (LACP). You learned about LACP operation modes and then applied that knowledge to configure and verify it.

You learned about load balancing algorithms and inputs, and then learned how to configure and verify LAG load balancing.

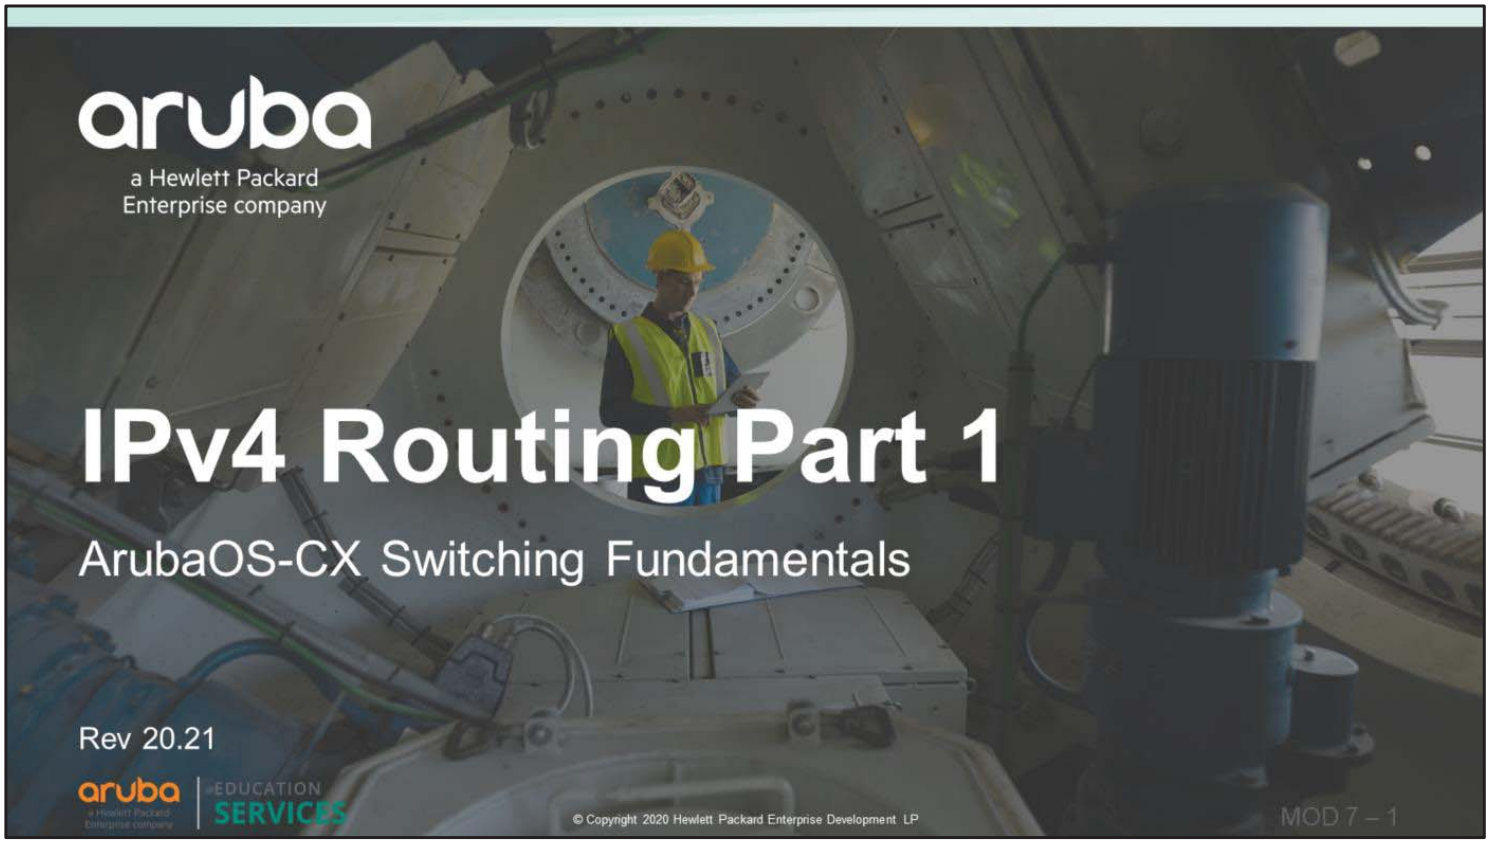
aruba
a Hewlett Packard Enterprise company
IPv4 Routing Part 1
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 7 – 1

Objectives

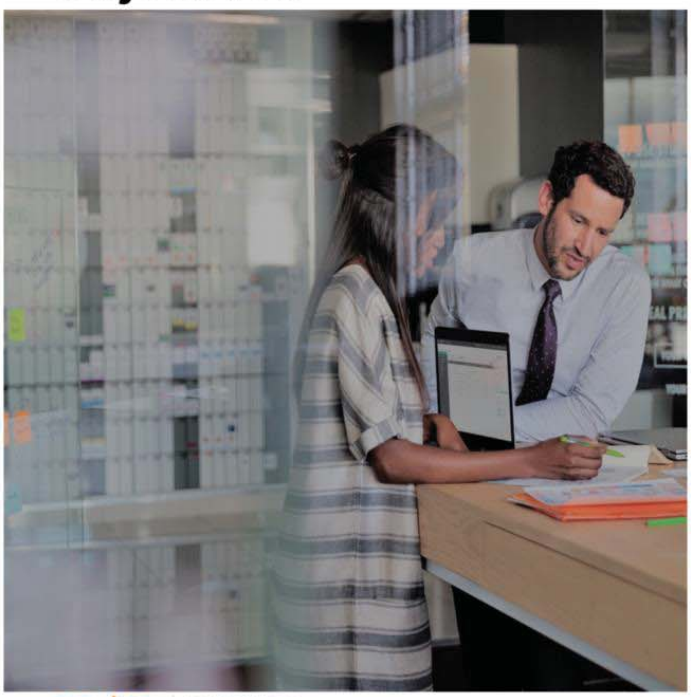

– Learning Objectives
  - Describe routing, IP addressing, and masking
  - Explain IP routes and the default gateway
  - Describe inter-VLAN routing
  - Describe the packet delivery process
  - Explain Virtual Routing and Forwarding

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 7 – 2

After completing this module, you will be familiar with fundamental Layer 3 routing concepts. You will be able to describe routing, IP addressing, and the addressing mask.
You will be able to explain IP routes and how endpoints use a Default Gateway (DG) configuration. You will also be able to describe Inter-VLAN routing, and the packet delivery process. Finally you will be able to explain Virtual Routing and Forwarding (VRF).

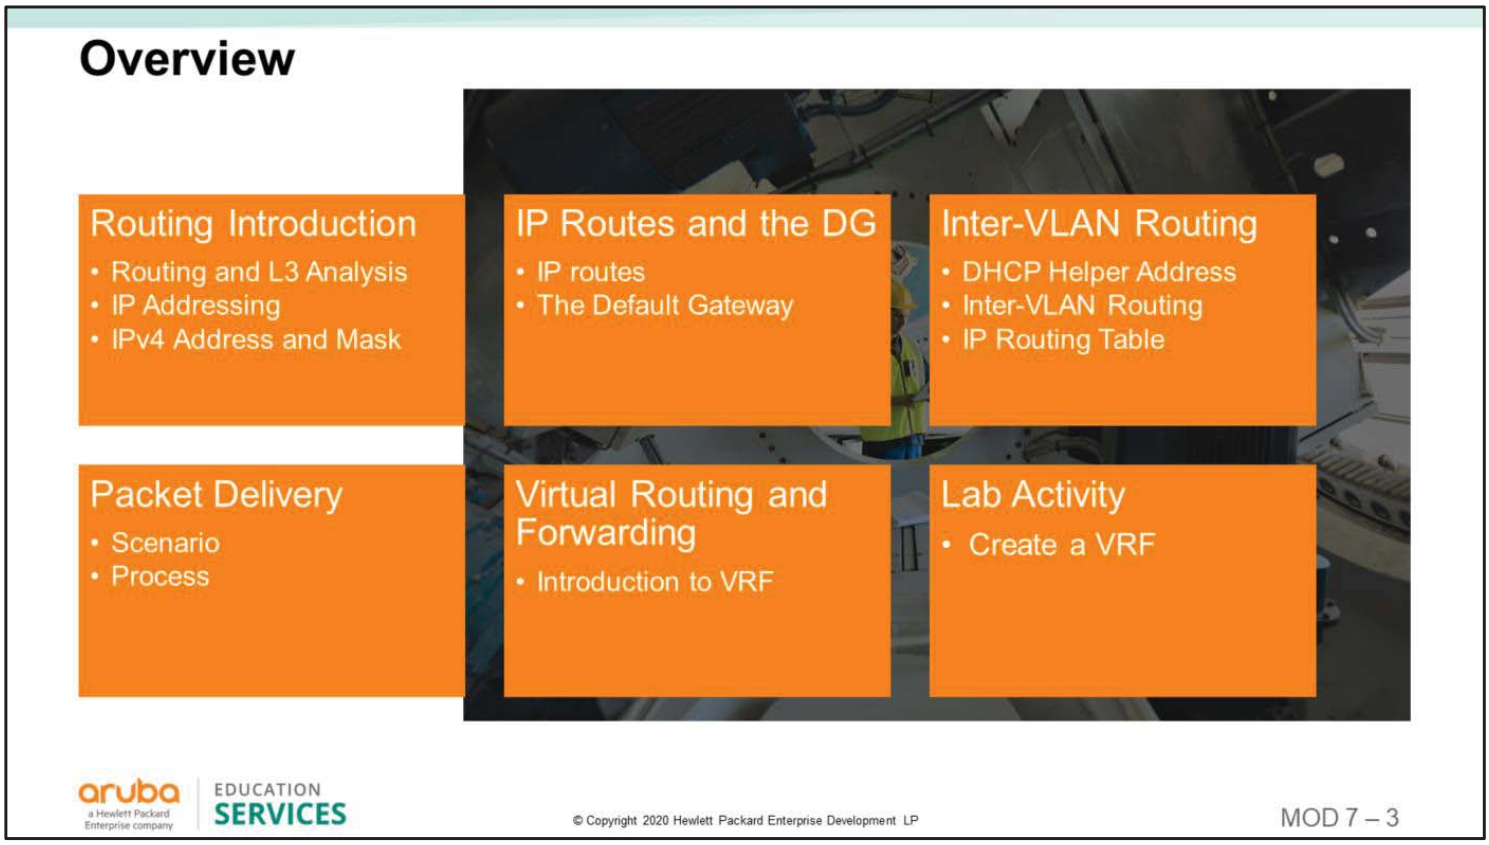

You have learned quite a bit about Layer 2 processes and Layer 2 frames – communications within a single VLAN. Now you learn how to connect those VLANs together and route between them. You see how routing devices use Layer 3 packet analysis to forward Layer 3 IP packets, and explore IP addressing and masking – vital skills to design, deploy, and diagnose scalable routed networks.

You will see how IP routes and Default Gateways (DG) benefit end systems, before exploring Inter-VLAN routing and DHCP helper addresses. Armed with this information, you will explore a Layer 3 packet delivery scenario.

You will learn how a single physical router can be divided into multiple virtual routers using Virtual Routing and Forwarding (VRF). Then you will apply this knowledge in a lab activity.

Routing Introduction

Basic concepts

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

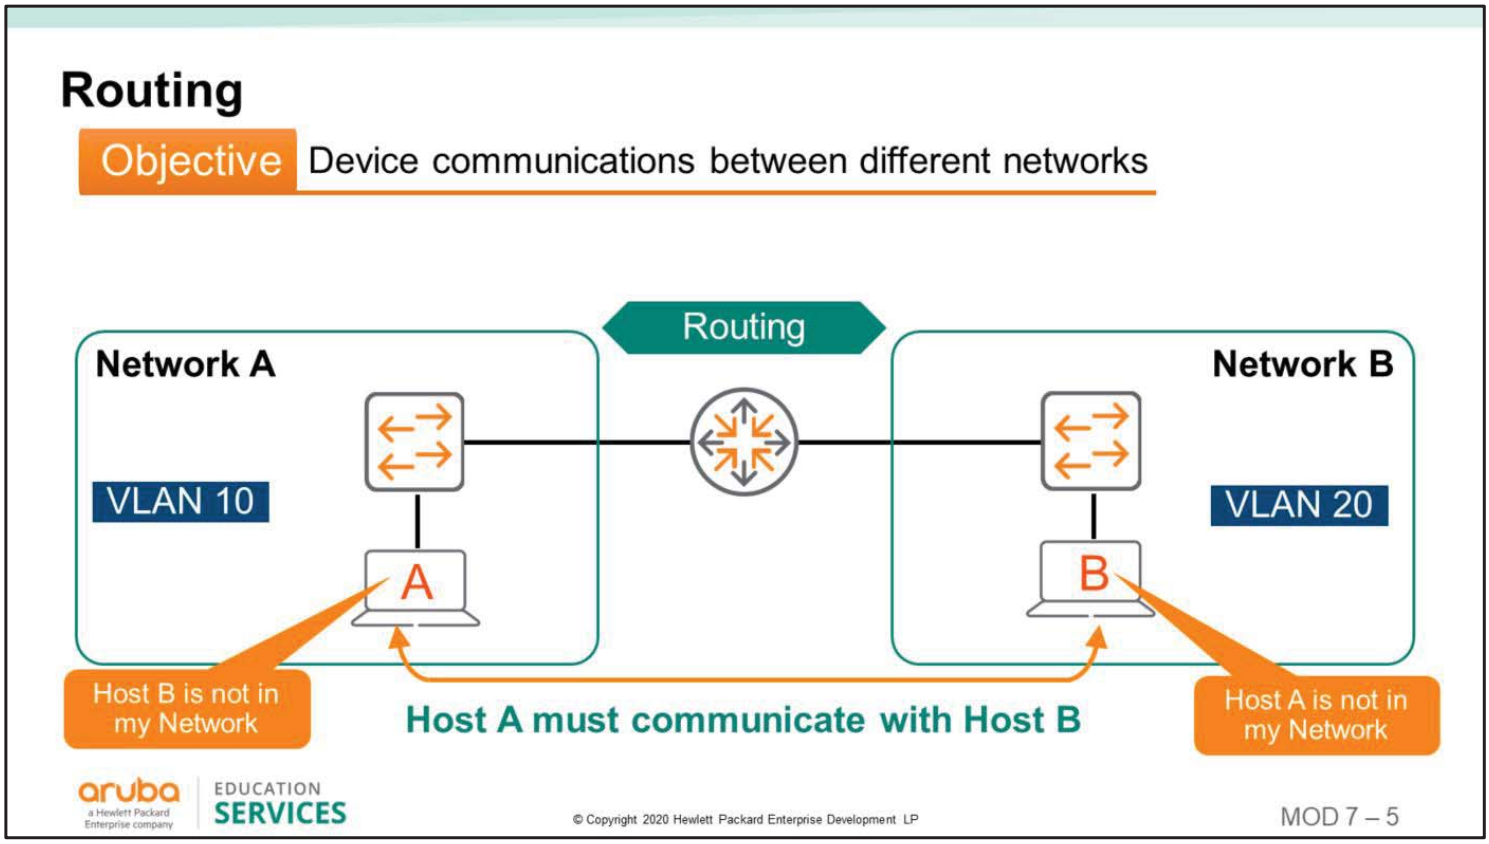

You have learned how devices communicate within the same network (broadcast domain), using Layer 2 switching. Layer 2 switches forward frames among devices in the same LAN by processing Layer 2 frame headers. Recall that switches build a MAC address table based on source MAC addresses, and forward frames based on destination MAC addresses.

Know you will learn how Layer 3 routers move packets between different LANs or broadcast domains, based on Layer 3 IP addresses.

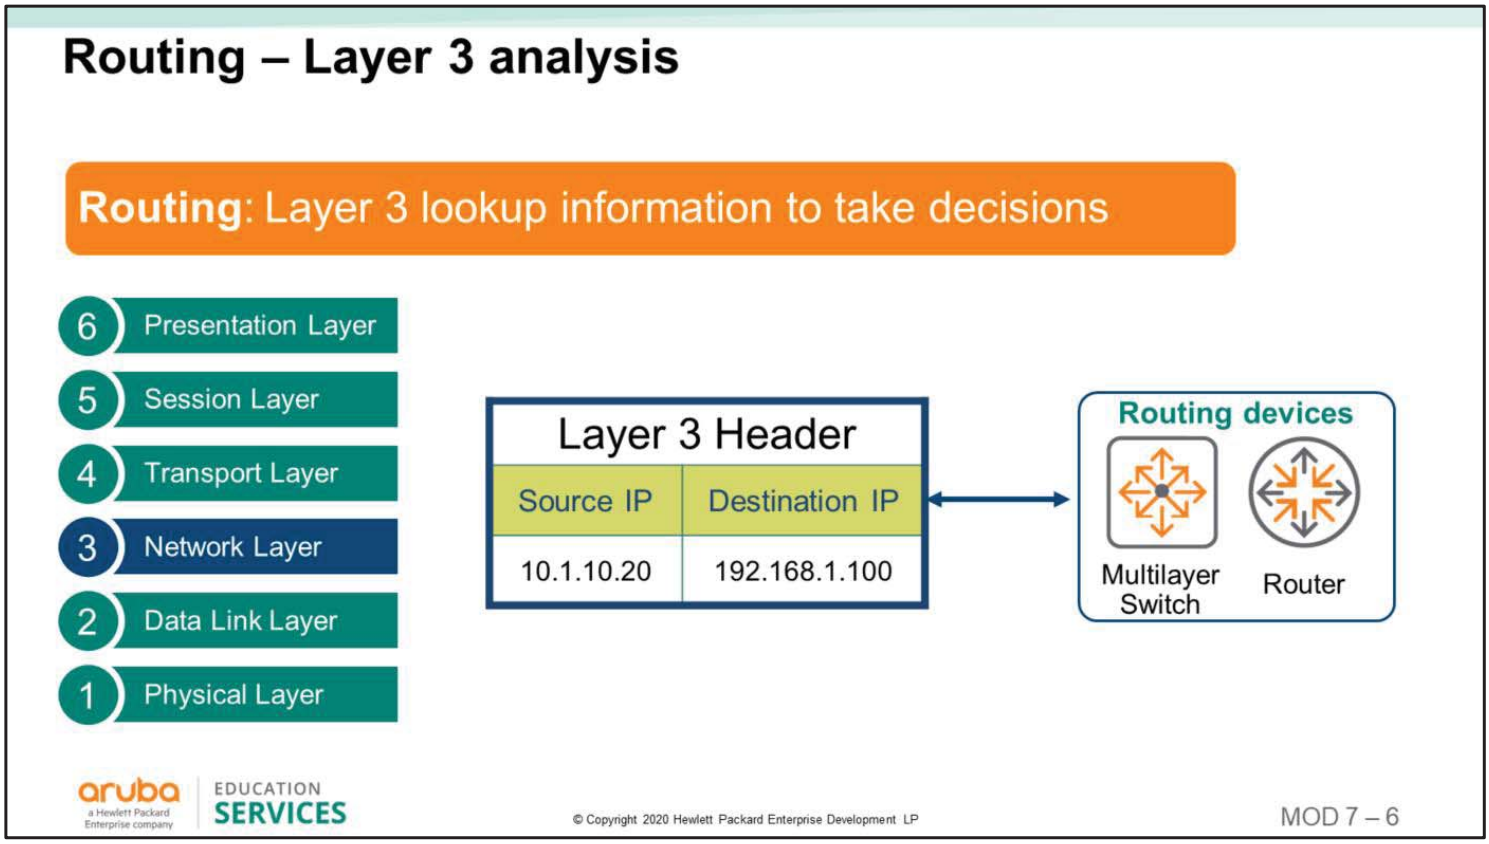

The ability to analyze the Layer 3 header is known as routing.

Layer 3 devices perform routing – they analyze Layer 3 IP addresses, select the best path to get from original source to ultimate destination, and the forward packets along that path.

IP Addressing

Provides identification for each device in a network

Feature	IPv4	IPv6
Size of address	32 bits – 4 octets	128 bits – 16 octets
Example of address	192.168.10.250	fe80:0000:0000:0000:844d:a9ff:febb:62c3
Number of possible addresses	4.29 Billion	340 Undecillion (3.4×10^{38})

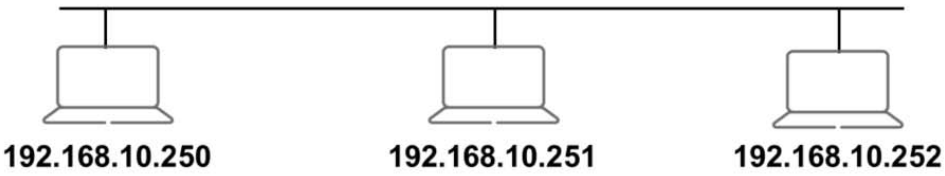

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 7 – 7

Basic routing decisions are based on the analysis of the Layer 3 addressing. The Internet Protocol (IP) provides an identification or address for each device in a network.

Currently there are two versions of IP that are widely used - IP version 4 (IPv4) and version 6 (IPv6). The main difference between these two protocols is the addressing space, version 4 can allocate approximately 4.29 billion addresses, while version 6 can allocate 340 undecillion addresses. This module is focused on IPv4 addressing and routing.

An IPv4 address consists of 32-bits expressed in a dot-decimal notation. This notation divides the address in four sections called octets. As the name implies, each octet is composed by 8-bits – a byte. This dot-decimal notation makes it easy for humans to work with IPv4 addresses. The figure shows three hosts in the same LAN, each with a unique IP address.

Note. Since an octet is composed by 8 bits, this implies that valid decimal values in an octet are from 0 to 255.

IPv4 address

IPv4 binary notation	11000000101010000000101000000011			
Octet	11000000	10101000	00001010	00000011
Decimal	192	168	10	3
IPv4 dot-decimal notation	**192.168.10.3**			

How many bits?
Network mask defines the portions

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 7 – 8

An IP address consists of two parts: Network ID and the Host ID. The network ID is the most significant part of the address (Left side) that identifies the network. The host ID on the other hand is the less significant part of the address (Right side) that identifies an individual host.

To help you understand the relationship between these two concepts, you can think that the Network ID is analogous to the street name and the Host ID is analogous to the house number in that street.

IPv4 addresses are always 32 bites long. Sometimes, 16 bits represent the network, and 16 bits represent the host. Sometimes there are 24 bits to specify the network, and 8 bits for hosts. There can be nearly any combination of network and host bits. This is all controlled by the subnet mask.

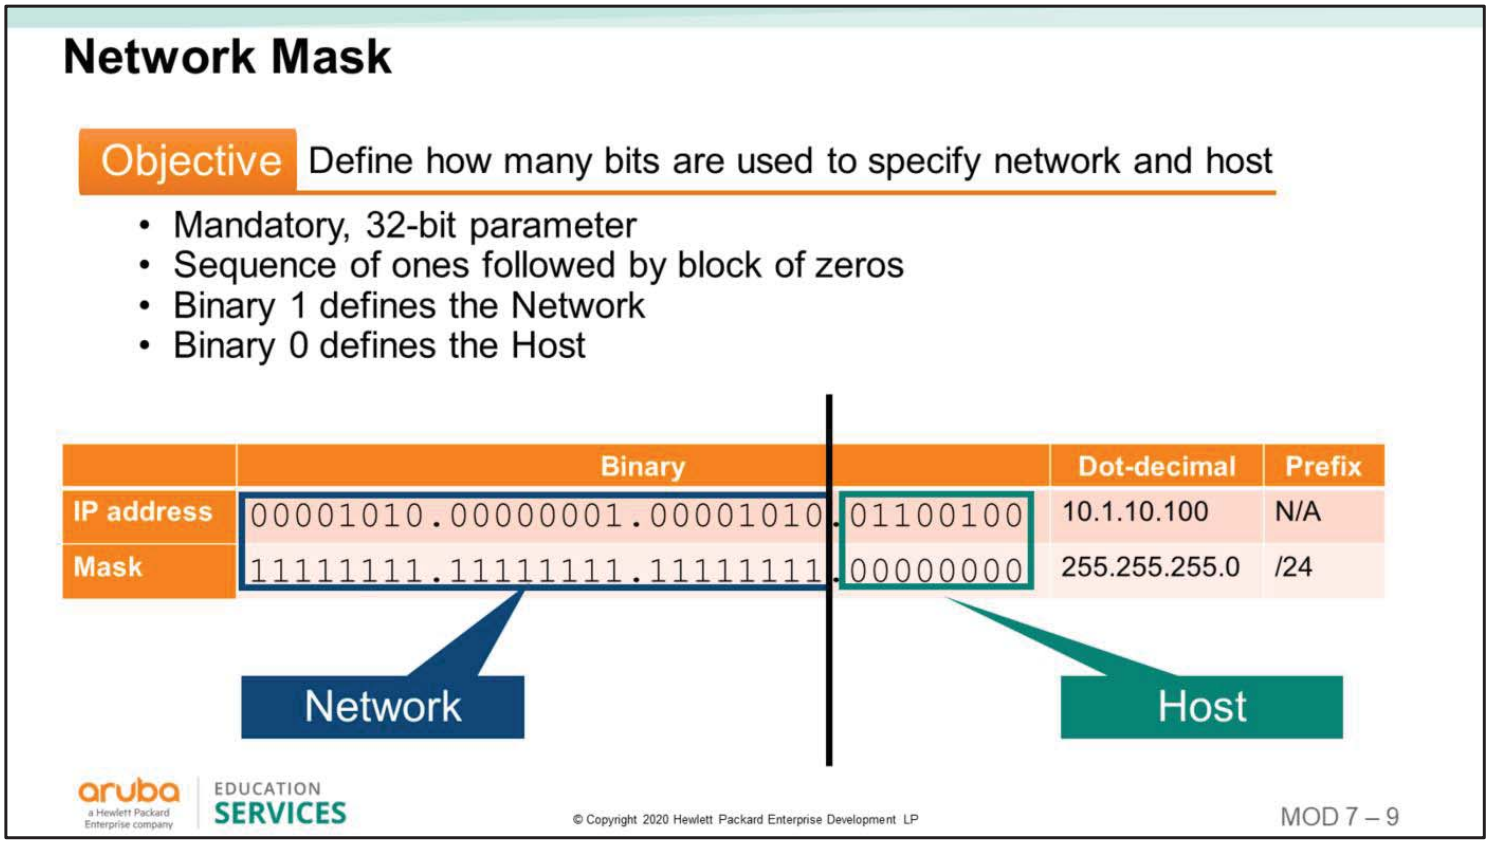

The network mask is an IP parameter that indicates how many bits represent the network portion of an address, and how many bits represent the host portion of an address. The 32-bit Network Mask is a mandatory IP parameter for all IP network devices.

The network mask determines if two endpoints are on the same network or on different networks. This process is done by a simple comparison.

If the Network ID for the source and destination is the same, then both devices are in the same broadcast domain. Layer 2 switching is enough to complete the communication.
If the Network ID for the source and destination is different, then the devices are in different networks. Layer 3 routing is required for this communication.

A Network Mask is 32-bits long, just like an IPv4 address. The mask is simply a contiguous string or block of binary ones, followed by a block of zeros. The ones indicate the portion of the IPv4 address that is assigned to the Network ID and the zeros will represent the portion of the IPv4 address assigned to the Host ID.

The image shows the relationship between an IP address and the Network mask. Where the binary 1's in the mask end and the 0's begin defines the line between the network and host portions of an IP address.

The Network Mask usually is represented in two different ways:

Decimal-dot notation. Same as an IPv4 address, the mask uses four different octets and each one is separated by a dot. For example: 255.255.255.0
Prefix notation. This notation in decimal indicates the number of bits that are set to one. This

notation uses a slash + number and is commonly placed next to an IP address. For example: 10.1.10.100/24. This notation indicates that the first 24 bits are set to one.

Question #1

IP Network Mask

Given IP address 172.20.3.54, and a mask of 255.255.255.0, what can be accurately stated about this addressing?

a. The host portion of the address is .3.54
b. This network portion of the address is 172.20.3
c. Host 172.20.3.89 would be on the same network.
d. A switch might use this address to forward the packet.
e. You could also indicate this mask as "/24".

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 7 – 10

Correct answers:
b.
c.
e.

IP Routes and Default Gateway

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

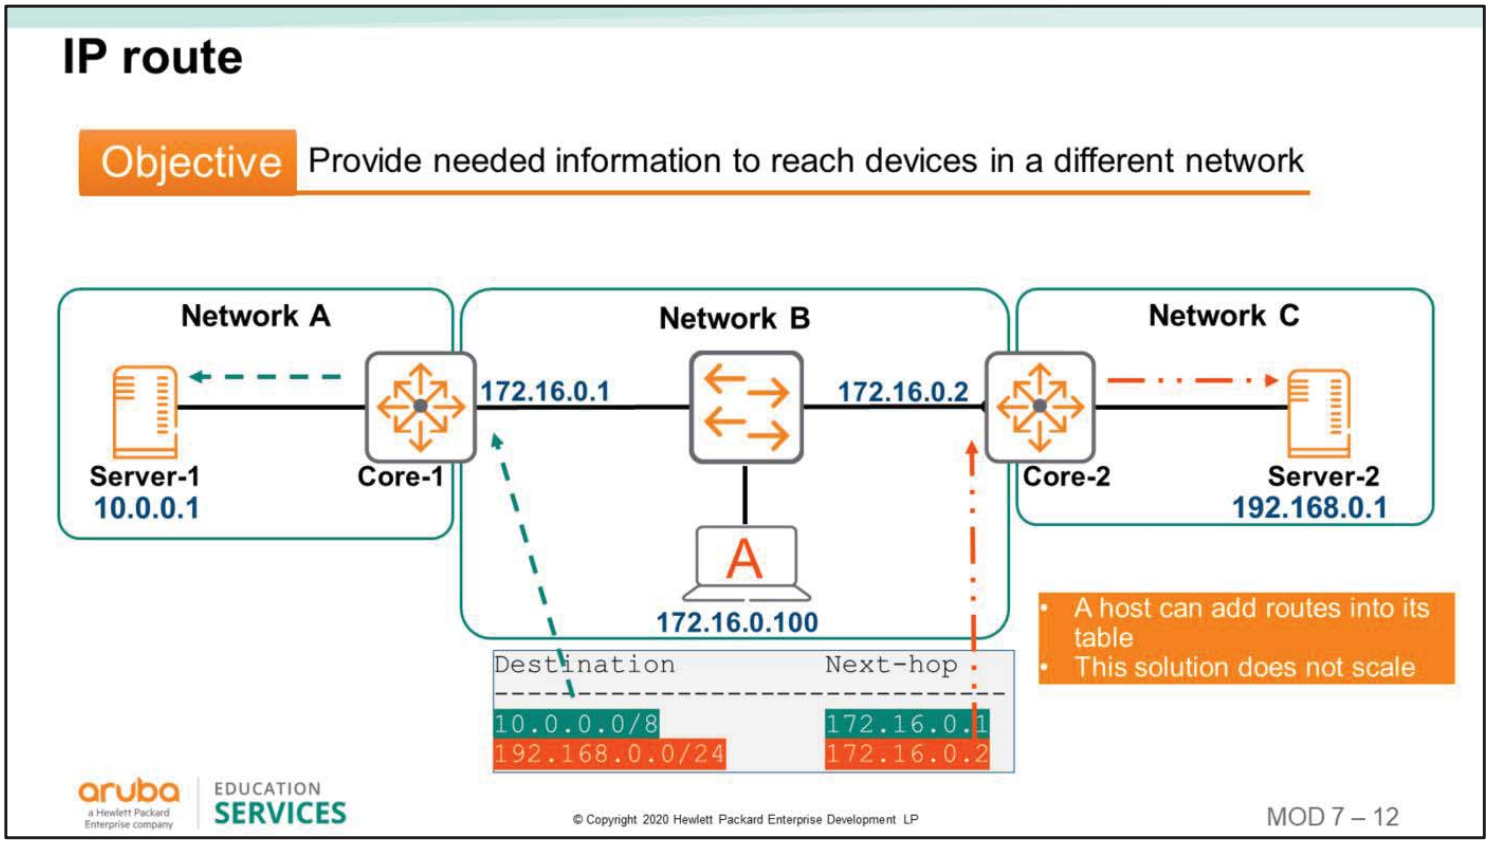

When a device must communication with others in a different network, it must know which local network device on its broadcast domain can route the traffic toward the destination network. This information is provided to computers using IP Routes.

For endpoints, this information must be manually added, in the form of a so-called “static route”. However, routers and multilayer switches can use manually-added static routes or they can dynamically and automatically determine the best routes to each destination, using a routing protocol. You will soon learn about dynamic routing protocols in a later module.

A static route must specify the following information, as shown in the figure for Router 1 and Router 2:
Destination IP address: where you want to go
Subnet mask: How that destination IP address is split into a network portion and a host portion
Next hop IP address: The IP address of the router on your network that can route packets to the destination.

Look at the scenario shown in the figure, which shows three networks connected by two routers. Core-1 connects Networks A and B, while Core-2 can route between networks B and C.
If Host A (in network B) must communicate with Server-1 in Network A, it must use Route 1. Route 1 says, “To get to destination 10.0.0.1, which has a mask of 255.0.0.0 (/8), you must use next-hop router at IP address 172.16.0.1.

What if Host A must use Server-2? Route 2 says, to arrive at 192.168.0.1, which uses a 255.255.255.0 (/24) mask, you must send the packets to next-hop router at IP address 172.16.0.2.

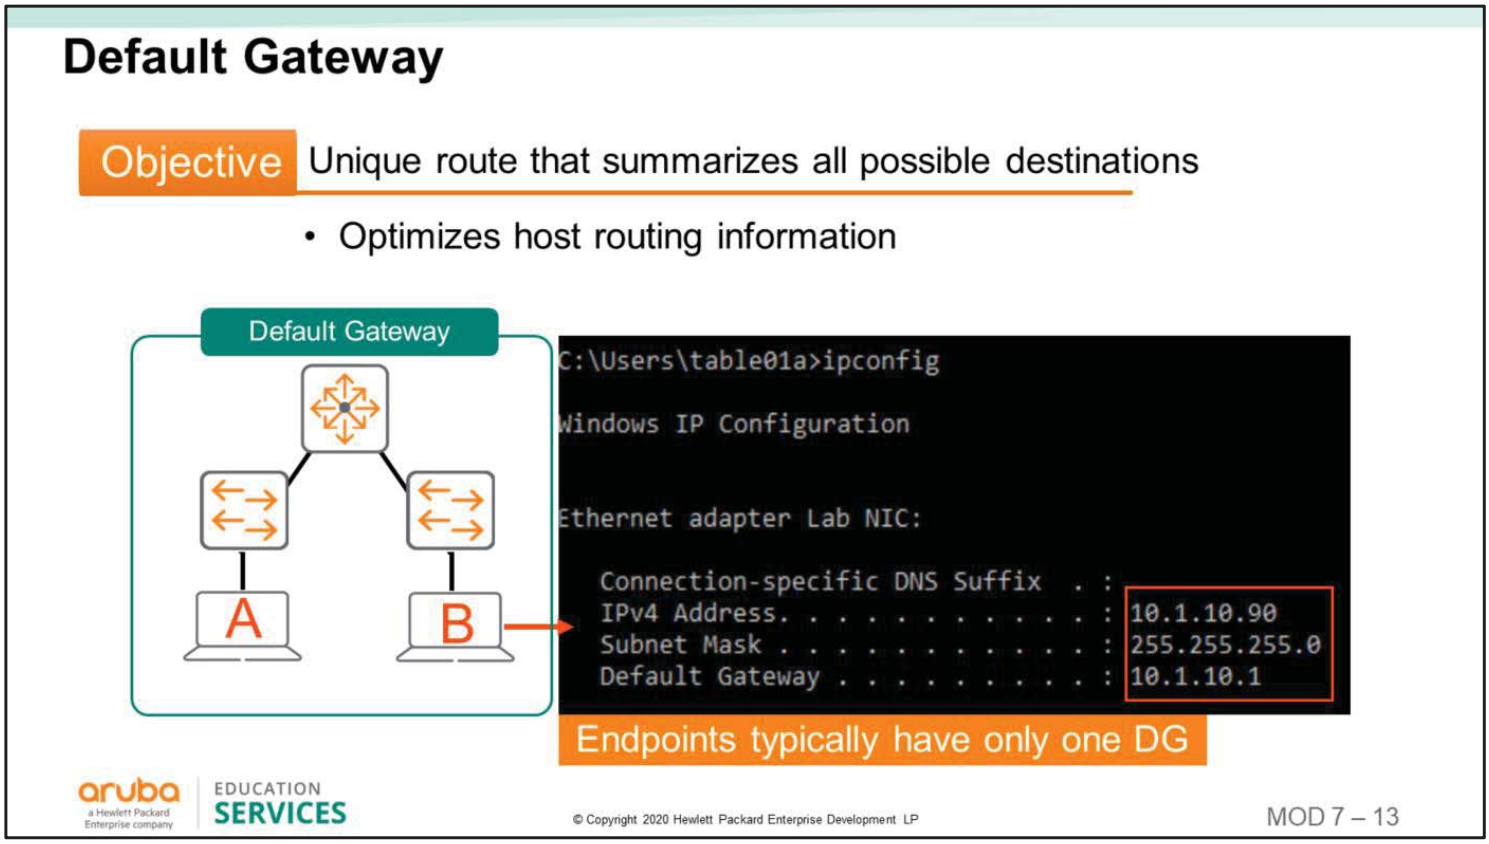

A default gateway (DG) is the device that routes all network destinations for the endpoint devices in a broadcast domain or network. It is like telling a host, “To get to everywhere in the known universe, go to this next-hop address”. The default gateway optimizes and simplifies endpoint routing decisions, since only a single route is required.

The figure shows how router Core-1 acts as the DG for all devices in Network B. These devices must only install a single route.

Note. Different endpoints in the same subnet could have different default gateways – if there is more than one router on the network.

In a windows device a static IP address and Default Gateway can be set it up in the same place. For windows 10 devices you can follow the next instructions:

Click on the windows start button and type Control Panel.
Navigate to Network and Internet → Network and Sharing Center → Change adapter settings.
Right click on the Network Interface Card (NIC) that you want to set up and select Properties.
Double click on Internet Protocol Version 4 (TCP/IPv4).
Enter the proper information for the IP address, subnet mask, Default Gateway. Optionally you can enter DNS server information.
Click OK twice.

Note. You have not yet learned about the subnet mask. This parameter will be discussed on module 8. For now, you should simply understand that a subnet mask must always be

configured, along with an IP address.

To verify the previous set up you, do the following:

Click on the windows start button and type Command Prompt.
Type ipconfig command.

Inter-VLAN Routing

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

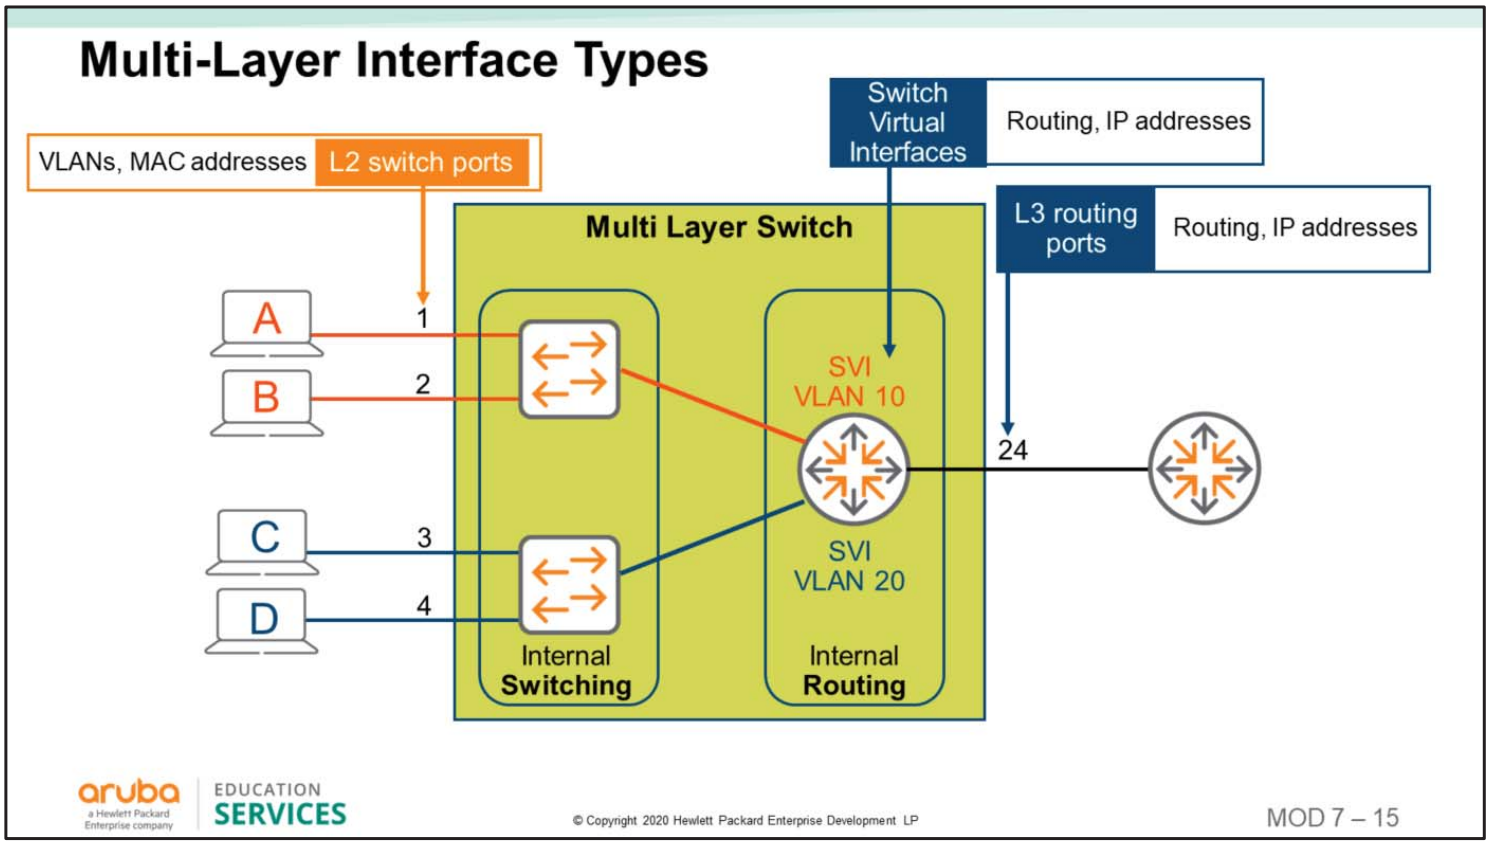

A switch port can only run Layer 2 features and protocols – like being a member of a VLAN, for example.
|The figure shows ports 1 – 5 being used as Layer 2 interfaces. They attach to end systems, accept L2 frames as being a member of some Layer 2 VLAN, and forward them – based on their destination Layer 2 MAC address. You learned that AOS-CX ports are Layer 3 interfaces by default, and so you must configure ports 1 – 5 with the command no routing.

|But what if you want to route between these VLANs. The AOS-CX switches are multi-layer switches - they have both internal Layer 2 switching functions, and internal Layer 3 routing capabilities. You need a way to connect each Layer 2 VLAN to the internal routing functions. To do this, you must create Switch Virtual Interfaces (SVI). This is a virtual Layer 3, routed interface that exists only inside the device, as a virtual construct.

|Suppose that you define SVI 10. Because it is an SVI, by definition it connects to the internal routing construct. Because it is SVI “10”, by definition it connects to VLAN 10, and so services routable traffic from VLAN 10 to other destination networks.

|Similarly, you might define SVI 20. With some routing configuration, which you will soon learn, your switch can now route traffic between your VLANs.

|Now suppose that you need to connect your multi-layer switch to an external router, perhaps using port 24 , as shown in the figure. Since all ports are Layer 3 interfaces by default, Port 24 connects to the internal routing functions by default. You merely need to configure it with typical layer 3 parameters, such as an IP address. You will soon learn about these concepts and syntax.

The SVI's are virtual Layer 3 interfaces, for internal routing, and port 24 is a physical layer 3 interface, for external routing. Both are layer 3 interfaces, and so perform routing functions. They accept routable Layer 3 packets and forward them based on their destination IP address.

Now you know about three important interface types – L2 switch ports, L3 SVI's, and L3 physical routed ports. You are ready to learn about another very important interface type – a trunk port.

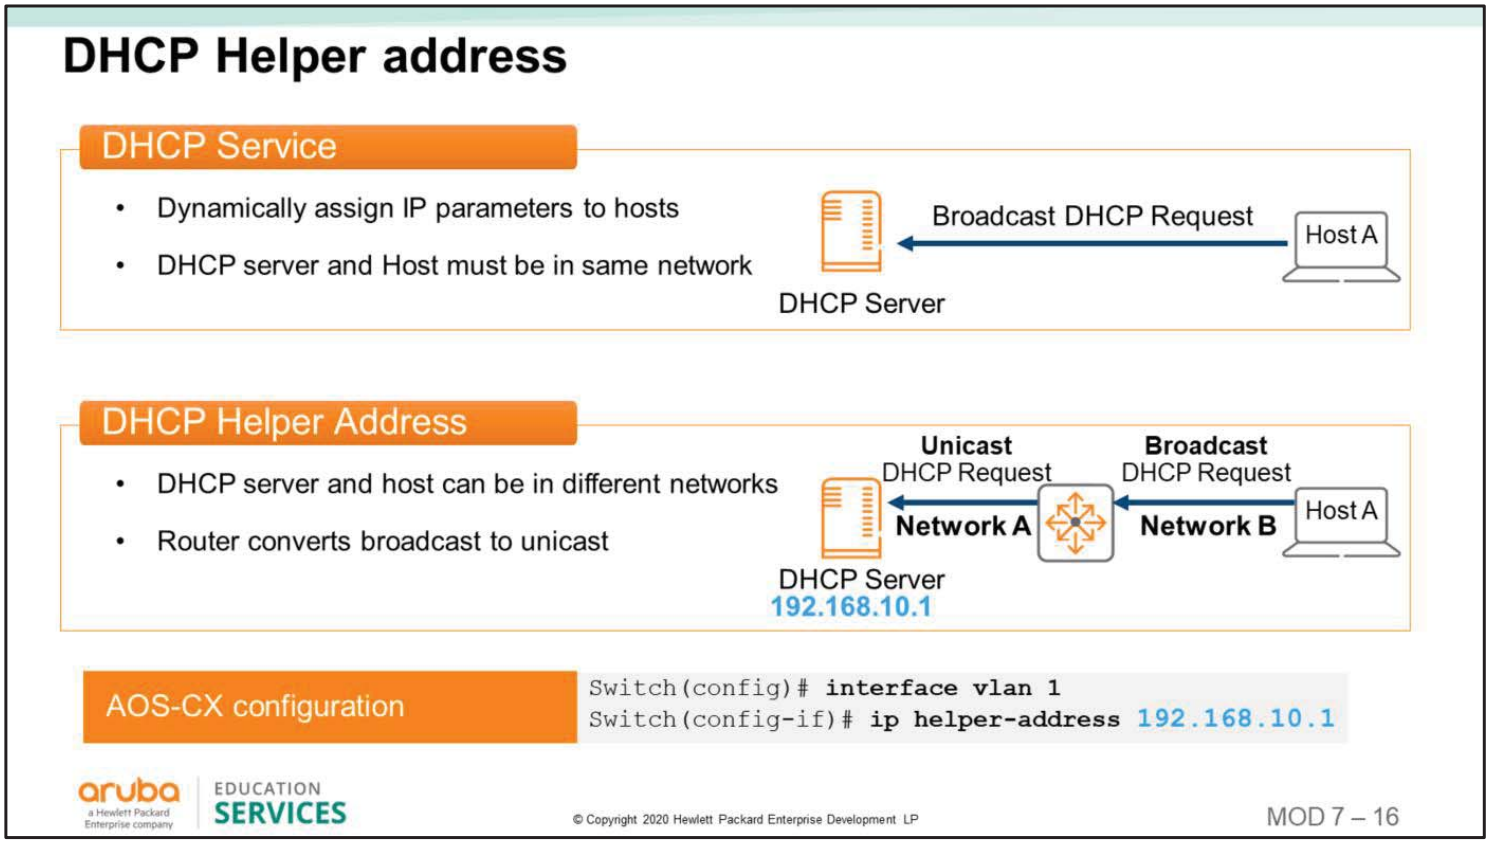

You have learned to set up static IP parameters on an endpoint. However, static IP address assignment is not typical. More commonly, hosts will automatically get an IP address by using the Dynamic Host Configuration Protocol (DHCP).

Endpoints broadcast a DHCP request, “Hey everyone, I need an IP address, a mask, and a DG. Because it is a broadcast, the host and server must be on the same subnet. Remember, a router defines the edge of a broadcast domain, and does not forward broadcast. Routers on different broadcast domains (VLAN) do not hear the request, and so no address is assigned.

Now wait a minute. If an organization deploys thousands of broadcast domains (VLANs), then you would need thousands of DHCP servers – one per VLAN! This is not realistic. You need a central DHCP service for all VLANs.

The solution is to configure a DHCP Helper address on each router interface that serves as the Default Gateway for endpoints. The following process describes how the solution works, when a router is properly configured with a helper address:

The client broadcasts a normal DHCP query.
The router that is on the client’s network (the client’s DG) receives this broadcasted DHCP query.
Instead of discarding the broadcast, as is normal, the router “helps” this broadcast by forwarding it on to the DHCP server. It converts this broadcast into a unicast, with the destination address specified in the ip helper-address command (192.168.10.1 in this example). Now that the message is a unicast, the router forwards it as it would any other unicast packet, towards its final destination - the DHCP Server.

Thus, the DHCP Server receives the DHCP request and replies with a DHCP offer - a unicast message sent to the requesting host, via the router.

In AOS-CX the command ip helper-address defines the address of a remote DHCP server. Up to eight addresses can be defined. When more than one DHCP server has been defined the switch will send the client request to all defined servers.

This command has to be used under a Layer 3 interface.

Switch(config)# interface vlan 1
Switch(config-if)# ip helper-address <IP-Address>

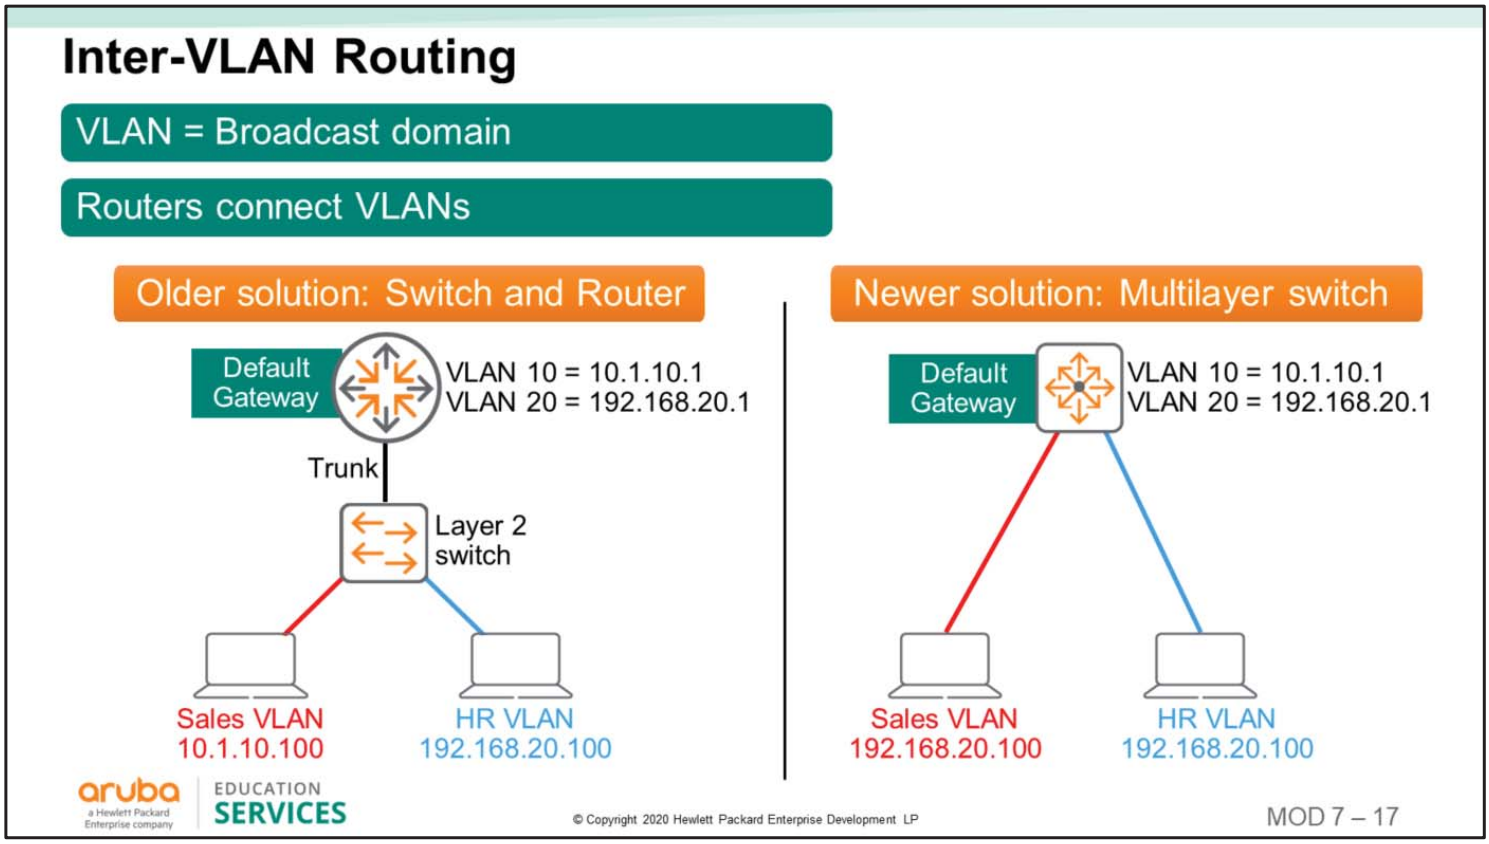

You learned about VLANs in Module 3. A VLAN is a broadcast domain, with a unique IP network number. In other words, all devices in the same VLAN have the same network address. All devices in the Sales VLAN are 10.1.10.x, where x is some unique host value. So IP addresses might be 10.1.10.100, 10.1.10.101, 10.1.10.102, and so on. Everyone in the HR VLAN might be 192.168.20.100, 192.168.20.101, and so on.

Recall that devices in different VLANs cannot communicate – unless you connect them with a router. Inter-VLAN routing connects separate VLANs into a routed internetwork of communicating devices.

In years past, multi-layer switches did not exist. Older environments used Layer 2 switches for host connectivity, and then routed between them with an external router. You see this in the figure's left-hand example. A potential problem with this deployment is that the Switch-to-router link can become oversubscribe, although LAG can alleviate this problem to some extent. Performance can also be suboptimal, because sending frames to the router requires an additional routing decision.

Multilayer switches are more efficient devices. The switching and routing functions of the device are connected via a high-speed internal backplane. Initial routing decisions and other processes all happen "in the box". This can reduce latency and increase performance.

All AOS-CX switches are multilayer switches which have routing enabled by default.

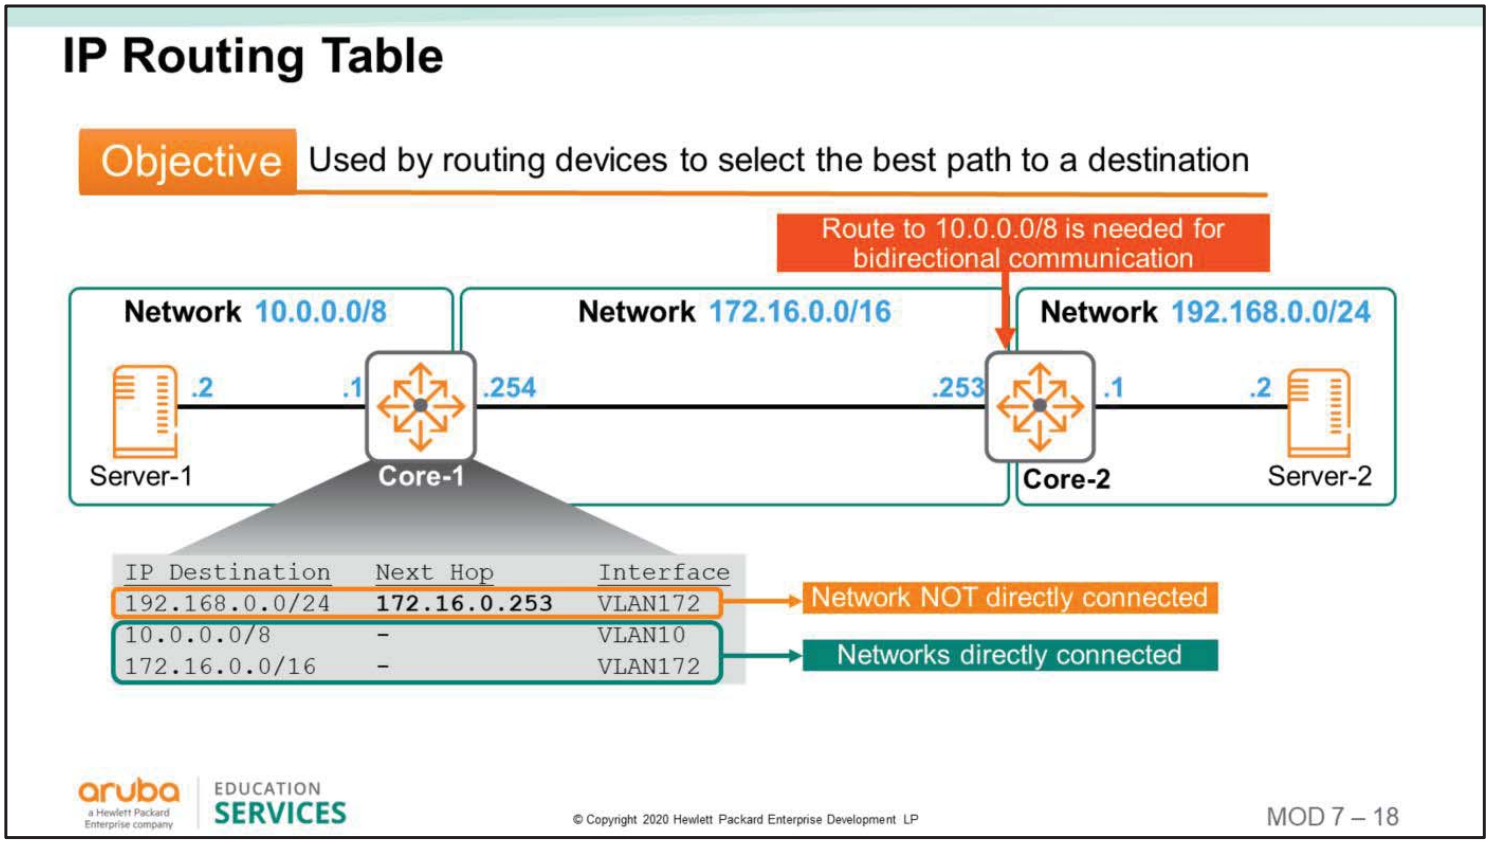

Routing devices (Routers and Multilayer switches) build and maintain a routing table that informs them of the best path to any given destination.

You can manually add entries to the route table, in the form of static routes. Alternatively, you can configure a routing protocol, which automatically builds and maintains this table. Typically, entries in the routing table do not expire unless a change in the topology causes an update. This differs from the MAC address table in Layer 2 switches where an entry expires after 5 minutes if the switch stops receiving traffic from the endpoint.

You will learn about all the entries that can exist in an actual route table soon. Meanwhile, the figure shows a slightly simplified view of the route table on the Core-1 and Core-2 routers. There are three networks, 10.0.0.0/8, 172.16.0.0/16 and 192.168.0.0/24. Recall that the subnet mask determine the network portion of an IP address.

OK. Now let us analyze Core-1's route table. The first entry says, "To get to any host on network 192.168.0.0/24, send packets to next hop router 172.16.0.253 (Core-2). To get to that next-hop address, forward the packet out local VLAN172".
The next entry says, "To get to any host on network 10.0.0.0/8, there is no next-hop...I'm directly connected to that network. Simply forward the packet out my local VLAN interface 10". Finally to connect to the network 172.16.0.0/16 Core-1 is directed connected to this network no there is no need to assign a Next-hop, simple forward the packet out of VLAN interface 172.

Consider at this moment that you configure only Core-1, when Server-1 connects to Server-2 Core-1 will properly route the packet and send it to Core-2. This device will have no problem to send the packet to Sever-2, this network is locally connected to Core-2. So

communication in one-way is successful. However the response from Server-2 is sent to Core-2, this device receives the packet but since the destination (10.0.0.2) is not its table route the packet will be dropped. Simply bidirectional communication cannot be done. To solve the problem remember that Core-2 will need to be configured with a route to reach the non-directed network (10.0.0.0/8). Try to think about how this route will look like.

Answer:

Destination = 10.0.0.0/8
Next-hop = 172.16.0.254
Interface = VLAN that is connected to network 172.16.0.0 – for this example VLAN172

Question #2

IP Routing Table

A router's IP routing table has an entry with a Next-Hop IP Address of 10.30.233.1. What does this number represent?

a. The packet destination
b. The next Layer 2 switch that should receive the packet
c. The cached ARP entry for the destination MAC address.
d. The next Layer 3 device that should receive the packet.
e. The router's own egress interface for routing.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 7 – 19

Correct answer:
d.

Packet Delivery

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

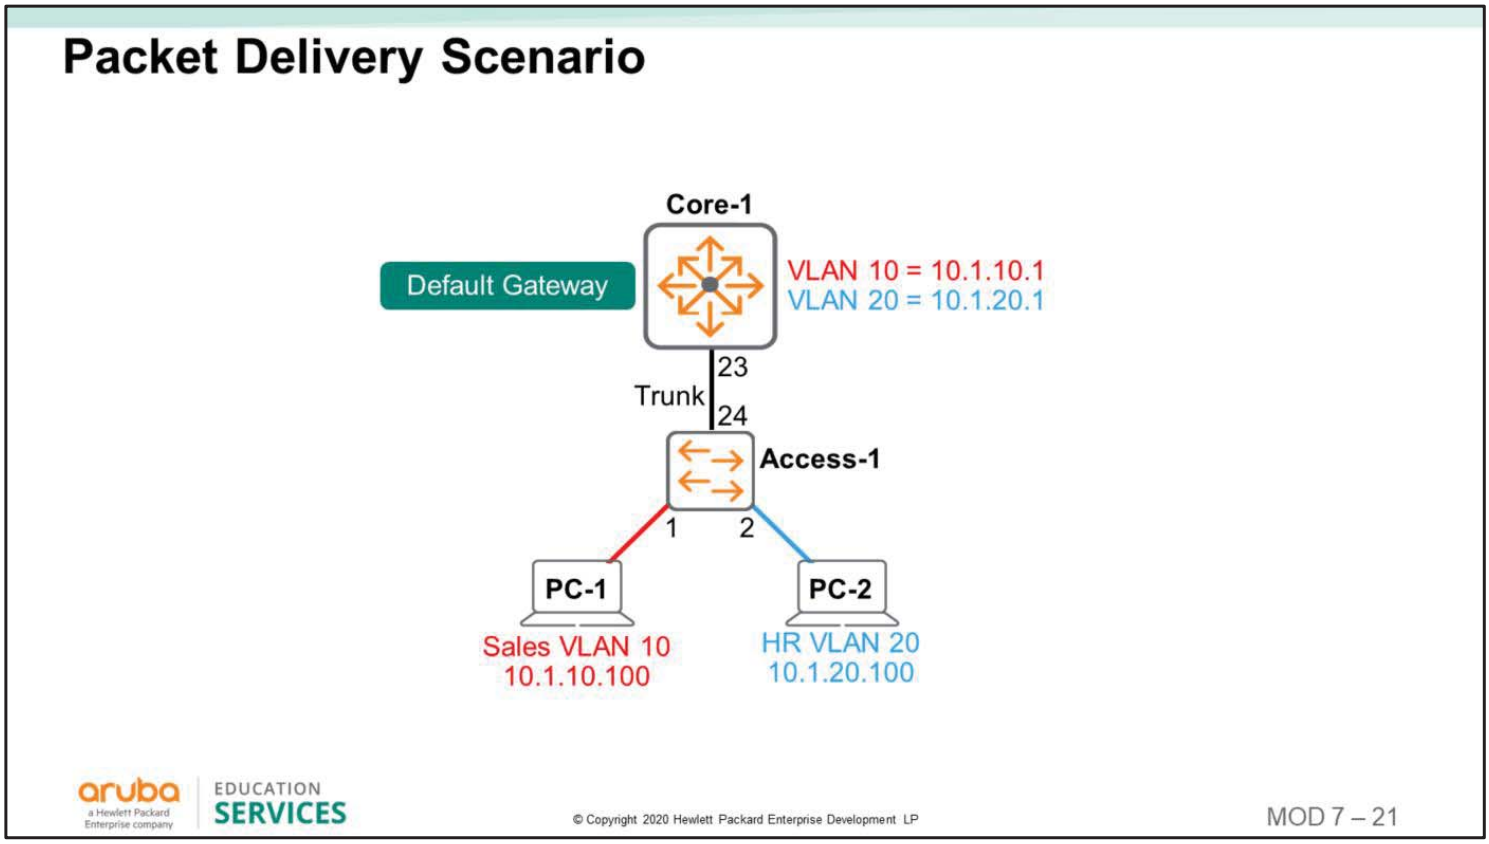

In this scenario, PC-1 must communicate with PC-2. Both endpoints connect to the same Layer 2 switch, but they are mapped to different VLANs. The Core-1 multilayer switch is there to do inter-VLAN routing.

PC-1 has IP address 10.1.10.100 on VLAN 10, and its DG is 10.1.10.1. It connects to port 1 of Access-1.
PC-2 has IP address 10.1.20.100 on VLAN 20, with DG = 10.1.20.1. It connects to port 2 of Access-1
Access-1 port 24 connects to Core-1 port 23.

Note. You learned that collapsing Layer 2 access services and Layer 3 routing services into a single multilayer switch can improve efficiency. However, for larger networks, there is typically a separate layer of pure Layer 2 switches for endpoint access, connected to a smaller set of L2/L3 multilayer switches. This design improves scalability.

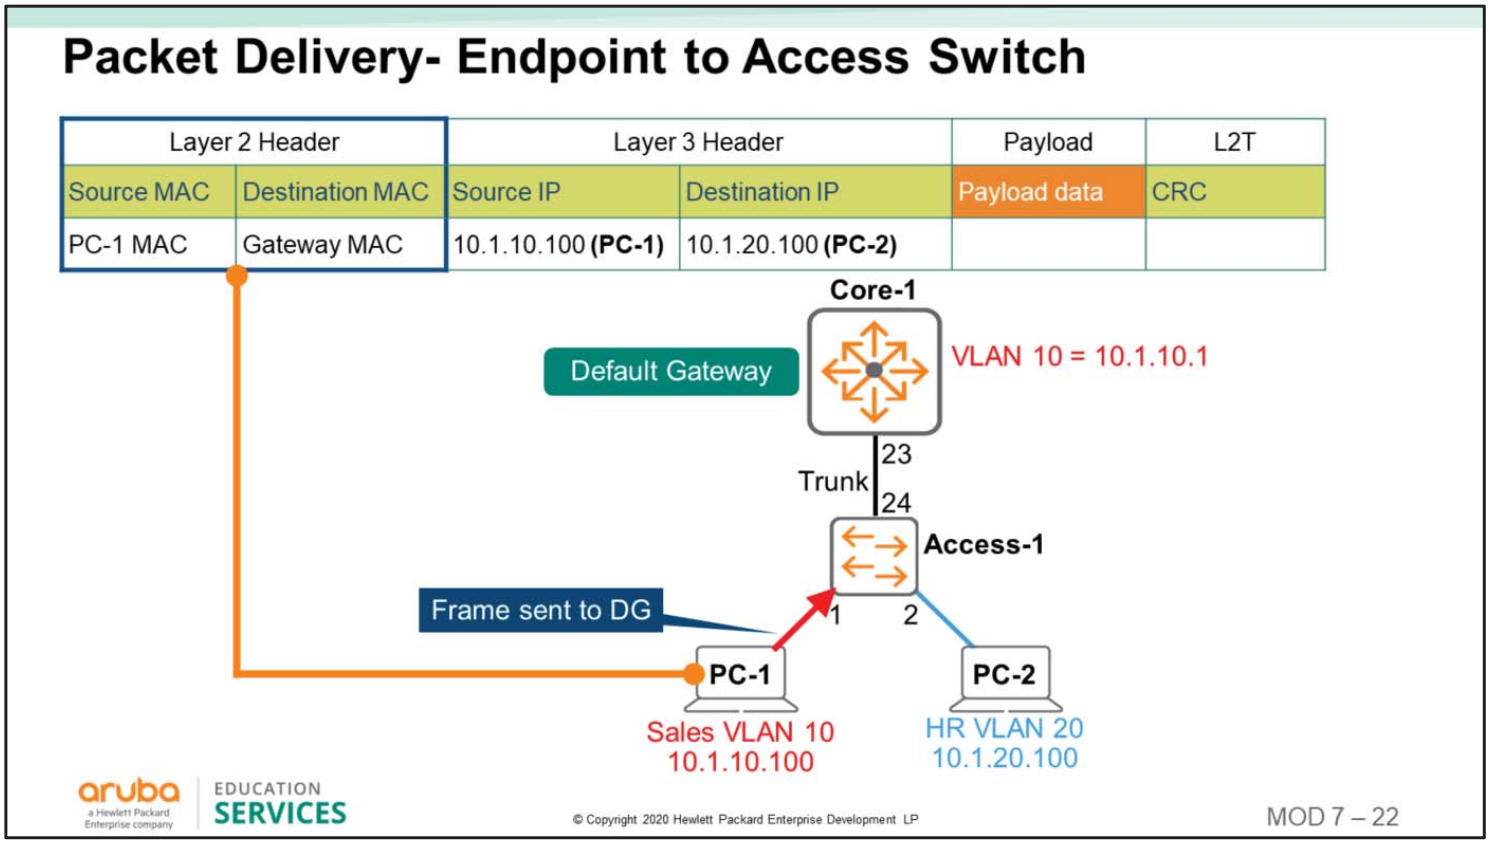

PC-1 generates a message that contains the following information:

Layer 3 header: Source IP address is PC-1's 10.1.10.100 and the destination IP is PC-2's 10.1.20.100.
Layer 2 header: Source MAC address is PC-1's MAC address and the destination MAC address is the default gateway (the MAC address associated with 10.1.10.1).

Remember, if PC-1 does not know the MAC address for 10.1.10.1 it performs an ARP process to get this information.

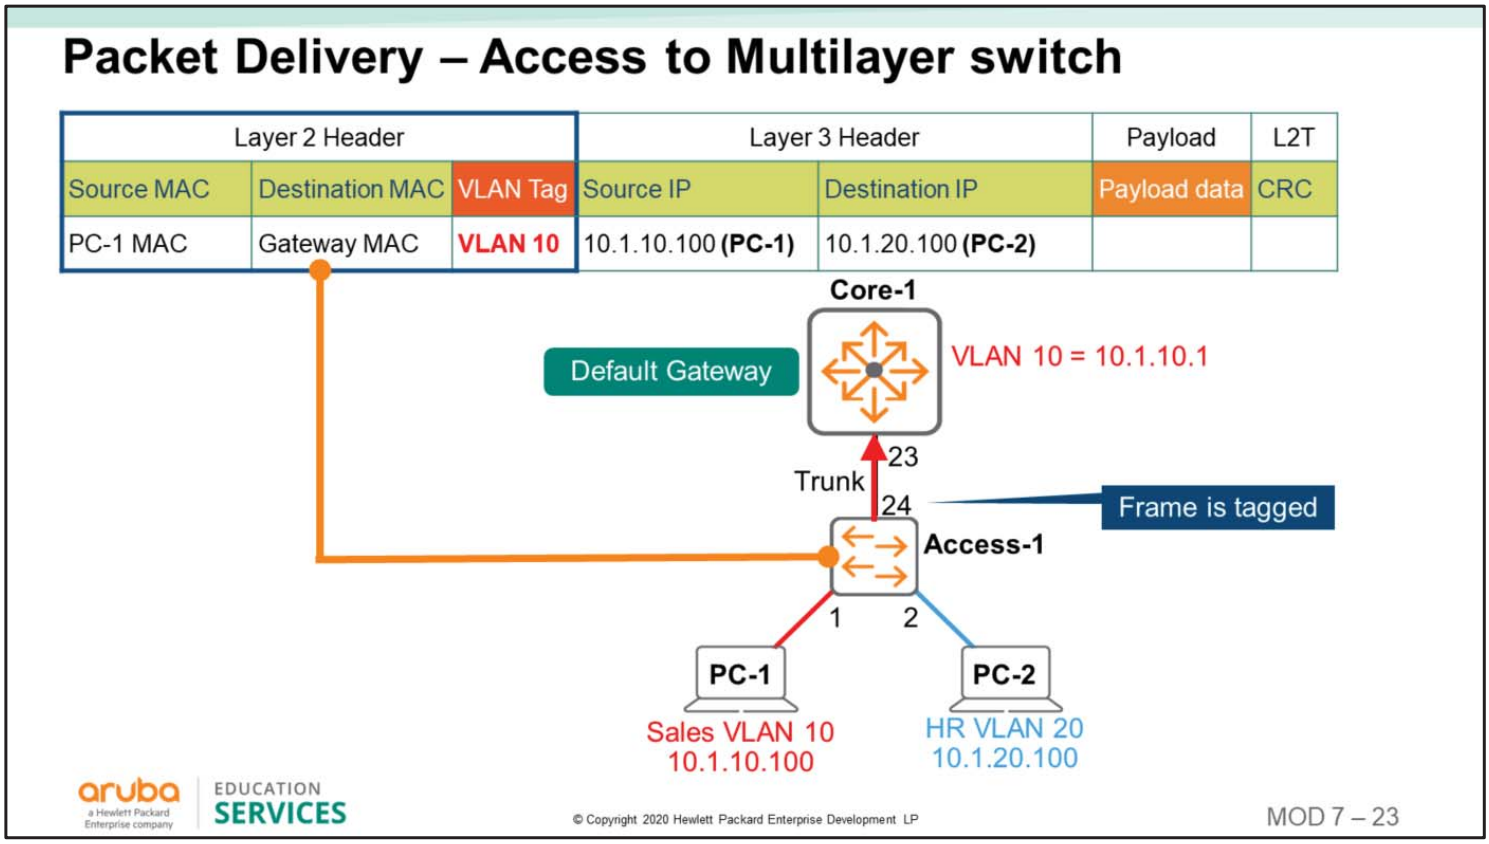

Access-1 receives the frame and analyzes the Layer 2 destination MAC address. It finds a match in its MAC address table and knows that it must forward this frame out its trunk link, port 24, to Core-1. It adds and 802.1q tag, VLAN = 10, and forwards the frame..

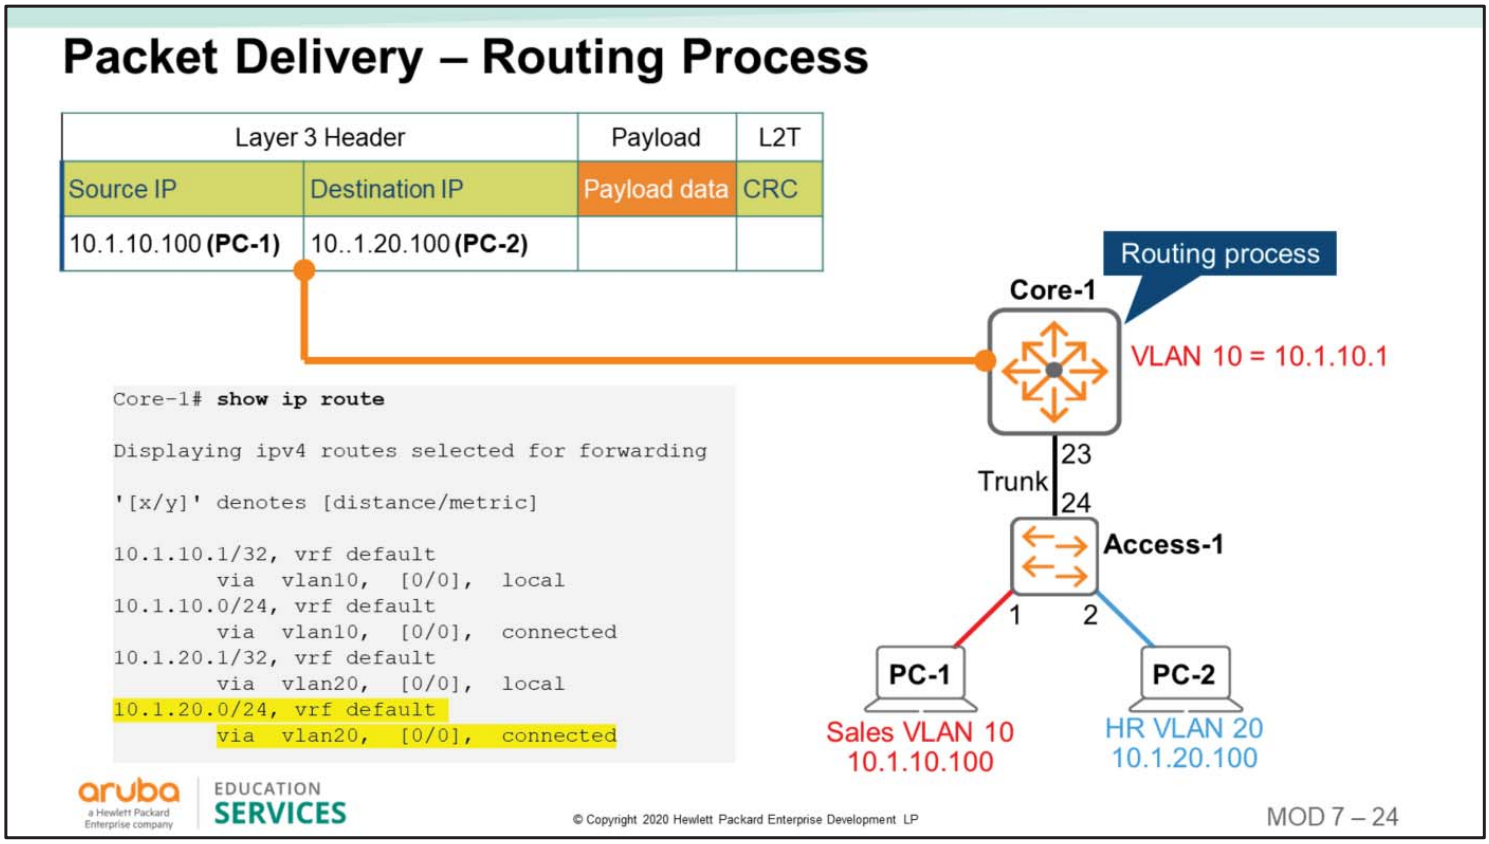

Multilayer switch Core-1 is the Layer 2 destination of this frame. It accepts the frame, strips off the Layer 2 header, and begins to perform its routing function, to analyze layer 3 header information.

It compares the Layer 3 destination IP address to its routing table entries. The figure shows Core-1's route table – the output of show ip route. Core-1 knows that destination network 10.1.20.0/24 is directly connected on its Switch Virtual Interface (SVI) VLAN20. Thus, Core-1 knows that it must forward the packet out its VLAN20 interface.

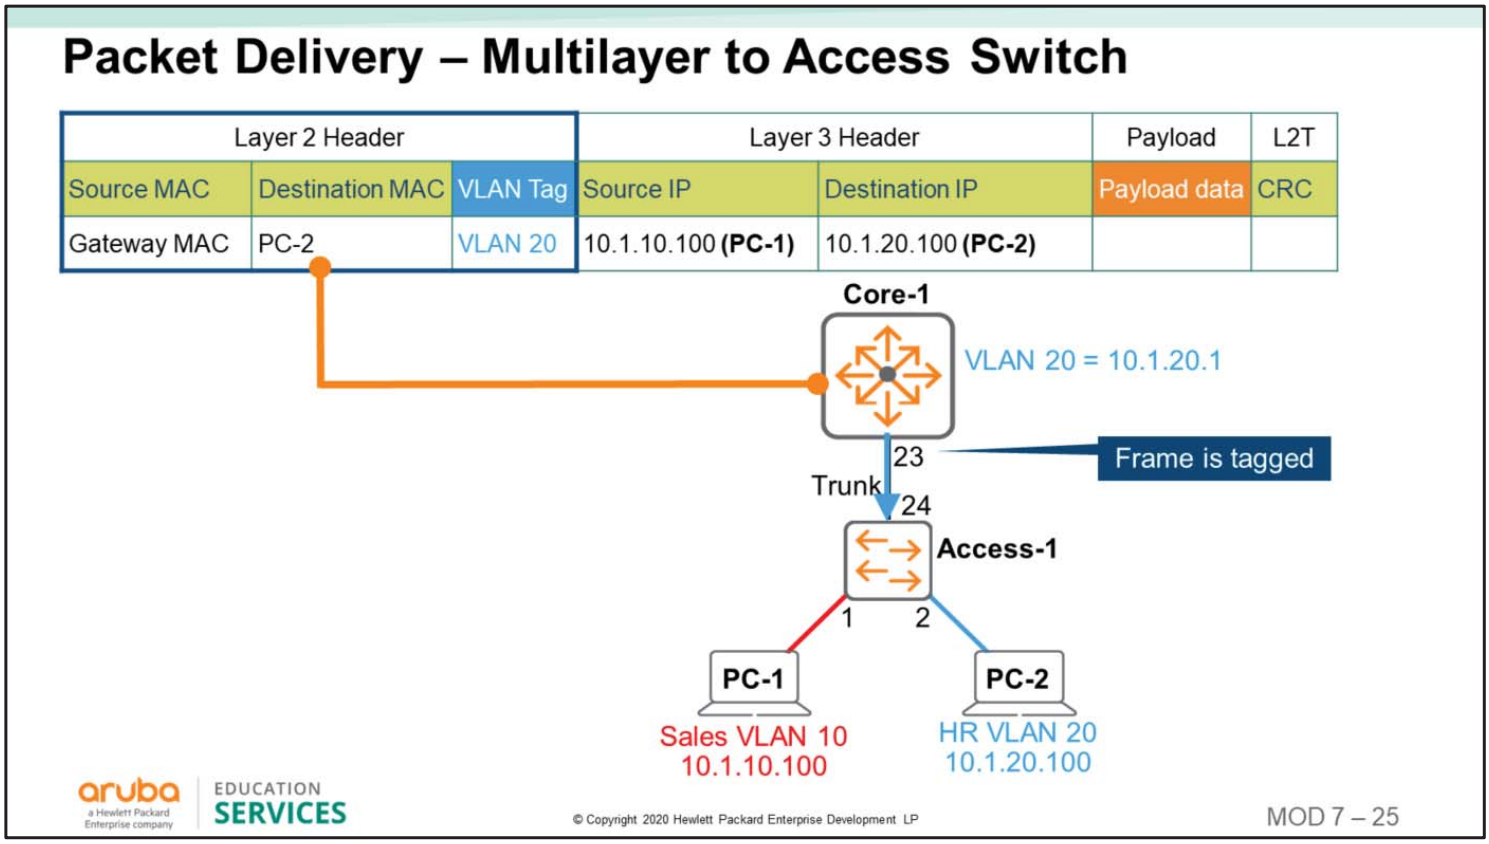

Core-1 builds a new frame to wrap around the IP packet. This frame includes an 802.1q tag – VLAN = 20. This frame is sent to L2 switch Access-1.

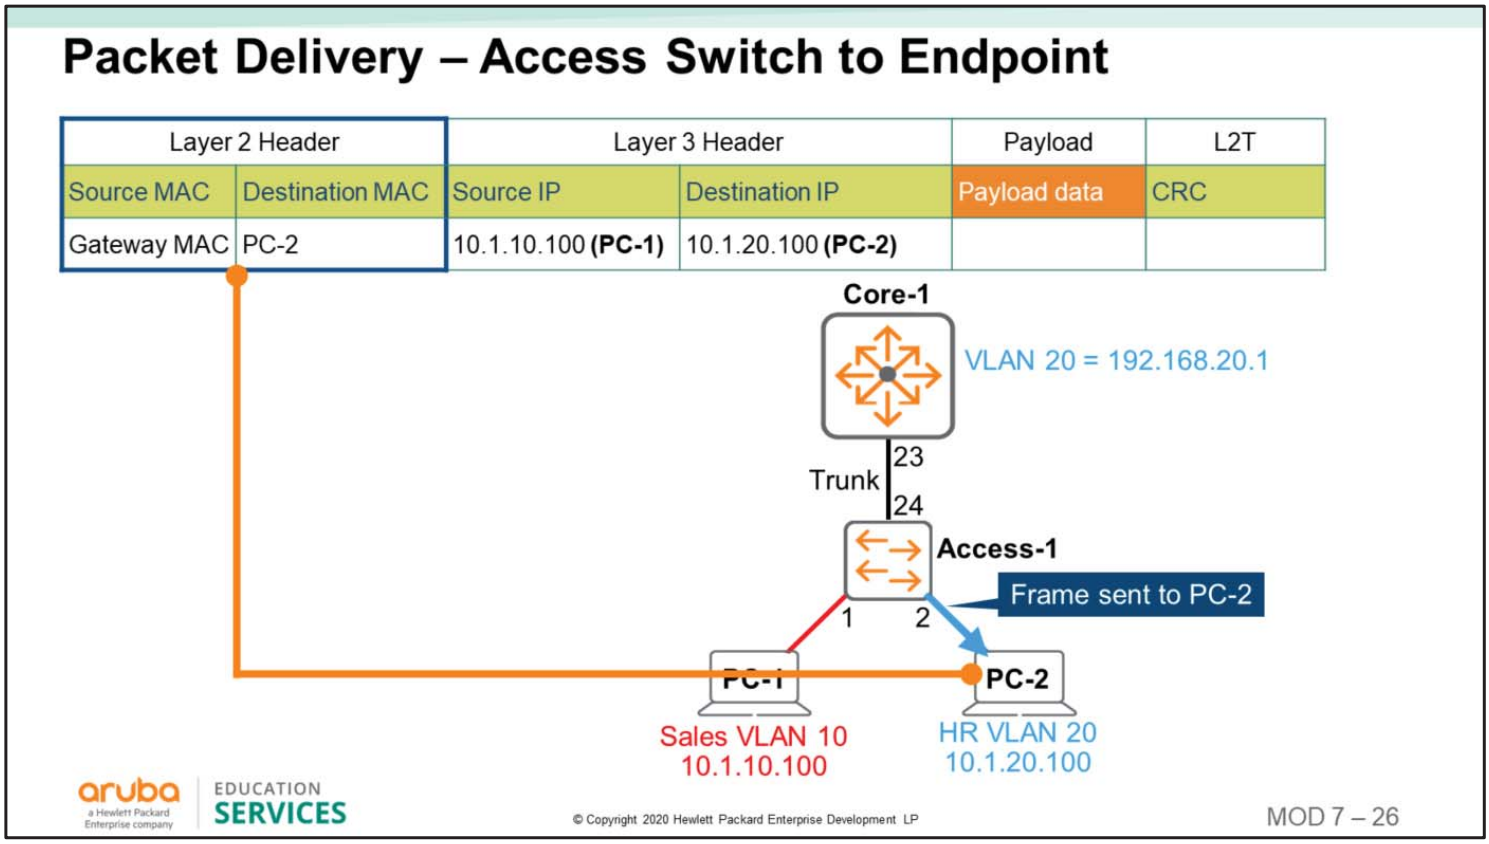

Access-1 receives the frame and learns from the tag that it is for VLAN 20. The 802.1q has served its purpose, and so it is removed. Access-1 compares the destination MAC address to its MAC address table, finds a match, and so forward it out port 2, toward PC-2.

Question #3

Packet Delivery

Which of the options below accurately describe a typical packet delivery process?

a. Access switches used their MAC address table to forward frames.

b. Multilayer switches use their MAC address table to forward frames.

c. Multilayer switches add 802.1q tags before sending frames to other switches.

d. Multilayer switches use an IP routing table to forward packets.

e. Access switches add 802.1q tags before sending frames to endpoints.

f. The source and destination IP addresses remain consistent throughout the packet delivery process, but the MAC addresses change.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 7 – 27

Correct answers:
a.
b.
c.
d.
f.

Virtual Routing and Forwarding

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

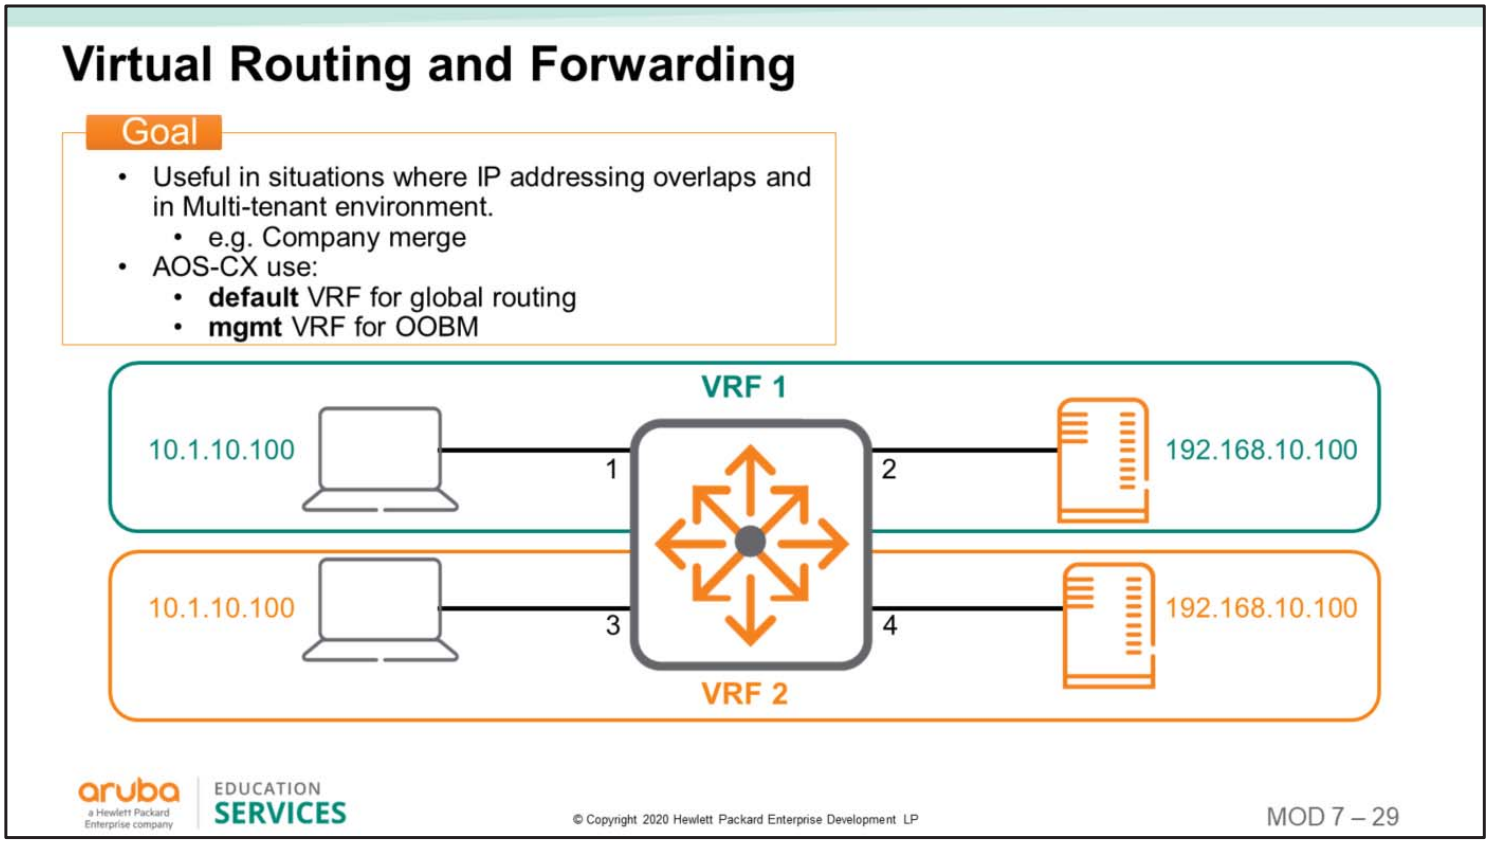

You learned that you can define several VLANs on a single physical switch. It is as if you have created multiple virtual switches inside the physical switch – one for each VLAN. Similarly, you can create separate virtual routers inside a single physical router, with Virtual Routing and Forwarding (VRF).

VRFs are useful in situations where the IP addressing overlaps in different places of the network. This could happen when two companies merge, for example. Also, another great use of VRF in a multi-tenant environment, customers A, B, and C, could be physically connected to the same L3 switch, but there routing tables are virtually kept separate and without using any ACL's to prevent communication between the customers.

The figure shows a single multilayer switch split into two separate VRFs. Interfaces 1 and 2 participate in VRF 1, and only interfaces 3 and 4 participate in VRF 2. These two VRFs do not interact. It is as if they are separate physical routers, with no connectivity between them. This is why the addressing can be the same in both VRFs without conflict.

In AOS-CX all interfaces (enabled with routing) by default are mapped to the Global VRF called default. In other words all interfaces are part of the same VRF – the physical router and the global VRF are essentially the same thing. Then you decide to create VRF 1 and 2, to split it up as shown in the figure. You know the two VRF's do not interact by default. However, you can configure the solution to route between the two VRFs if needed.

AOS-CX also includes a specific VRF for management purposes and that can only be used in the Out-of-Band Management (OOBM) port. To separate the data and control plane from the management plane.

Lab Activity

7.1 Creating a VRF

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

In previous modules, you learned quite a bit about Layer 2 processes and Layer 2 frames – communications within a single VLAN.
In this module, you learned how to connect those VLANs together and route between them. You saw how routing devices use Layer 3 packet analysis to forward Layer 3 IP packets, and explored IP addressing and masking – vital skills to design, deploy, and diagnose scalable routed networks.
You saw how IP routes and Default Gateways (DG) benefit end systems and explored Inter-VLAN routing and DHCP helper addresses. Armed with this information, you explored a Layer 3 packet delivery scenario.

You learned how a single physical router can be divided into multiple virtual routers using Virtual Routing and Forwarding (VRF). Then you applied this knowledge in a lab activity.

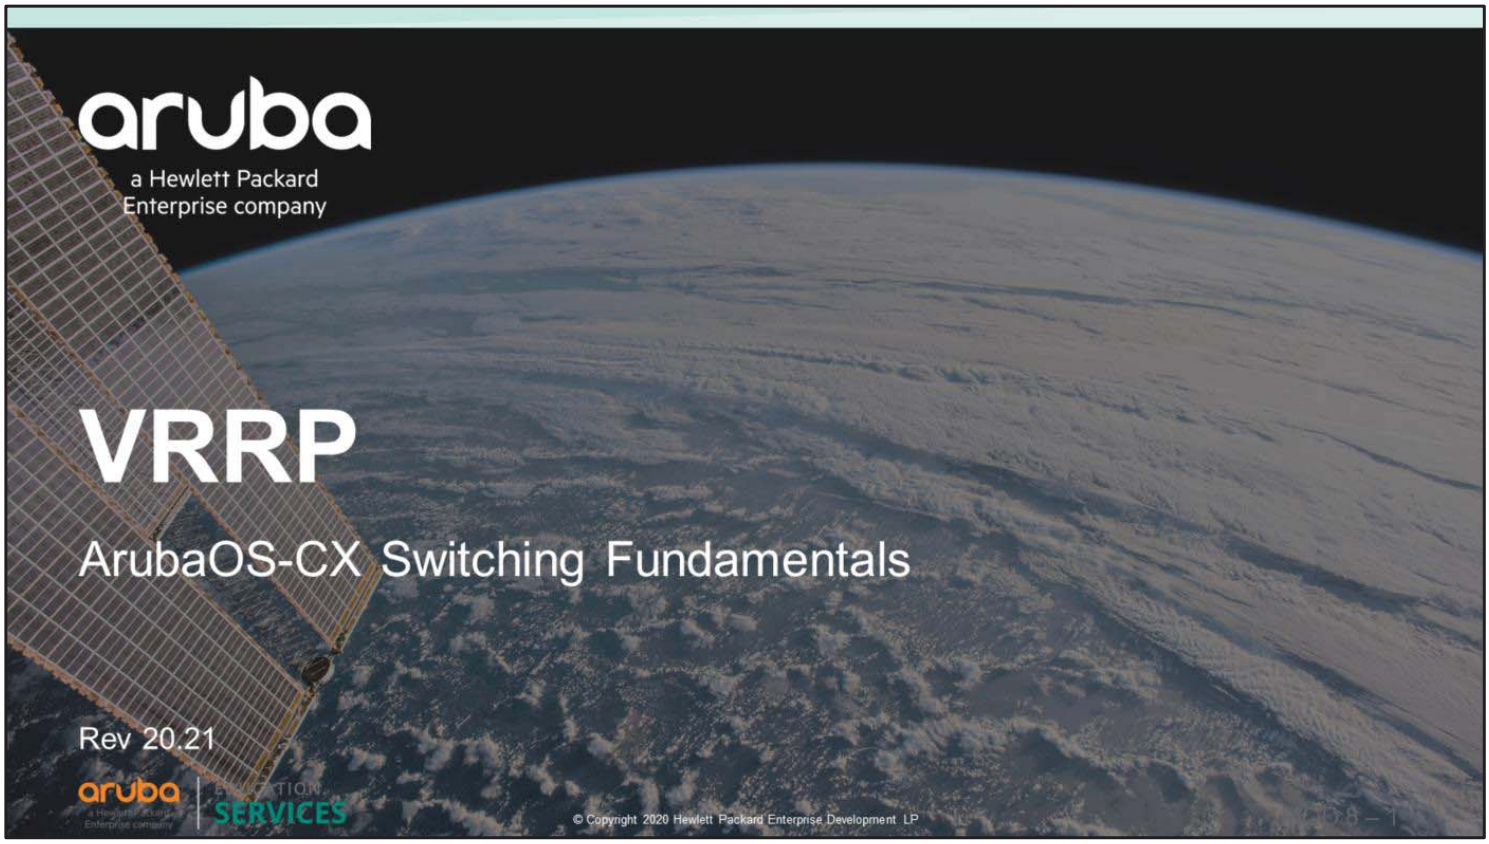
aruba
a Hewlett Packard Enterprise company
VRRP
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP

Objectives

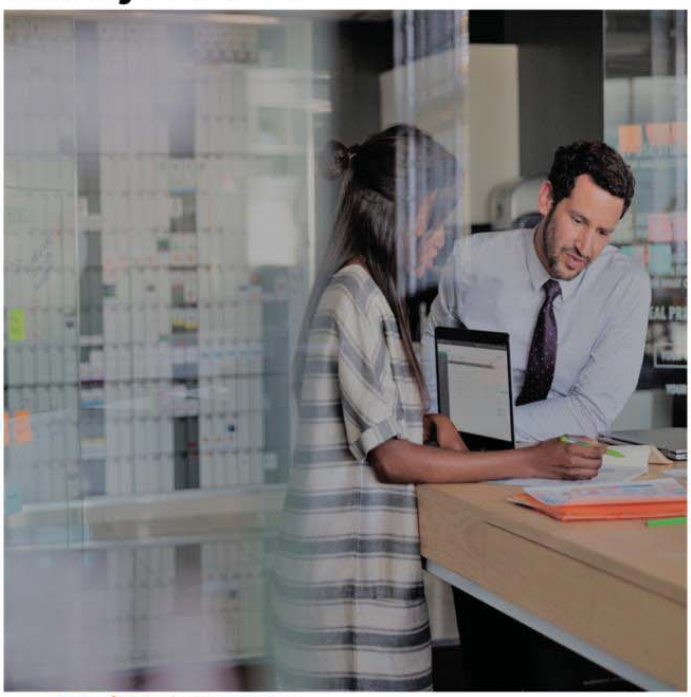

- Learning Objectives
  - Explain the need for L3 redundancy and First Hop Redundancy protocols
  - Describe VRRP and VRRP instances
  - Describe master election
  - Describe failover operation
  - Explain VRRP preemption
  - Describe VRRP and MSTP coordination

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 8 – 2

After completing this module, you will be able to explain the need for Layer 3 redundancy as relates to forwarding packets from endpoints, and First Hop Redundancy Protocols (FHRP). You will be able to describe a common standard implementation of an FHRP – The Virtual Router Redundancy Protocol (VRRP).
This includes VRRP instances, master election, failover operation, and VRRP preemption. Finally, you will be able to describe why you must ensure that Layer 3 VRRP redundancy and Layer 2 MSTP redundancy are coordinated.

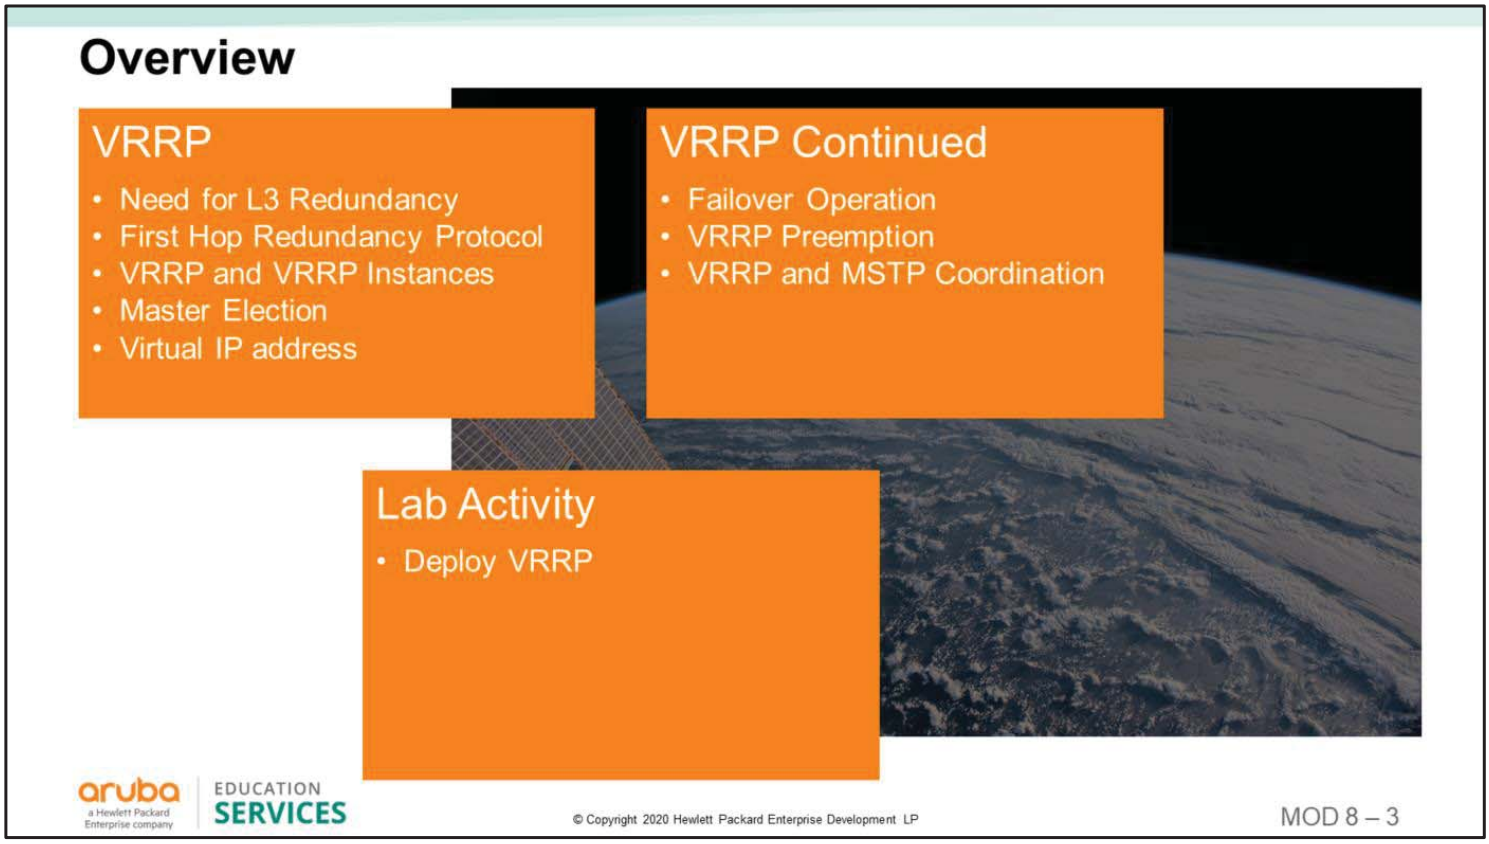

In this module, you will learn how to eliminate a potentially costly single point of failure – the endpoint's Default Gateway (DG). You will come to understand how a single DG can cause major outages, and why you cannot simply add a second DG.

You need to understand the Virtual Router Redundancy Protocol (VRRP), and how it automatically mitigates downtime upon DG failures. You will learn about VRRP operation, including how to load balance endpoint traffic by configuring multiple VRRP instances on a single set of routers.

You learn about the Master election process and how to control it by configuring priority values. You also learn about preemption options, to control what happens when a failed Master comes back online.

Finally you learn about the importance of coordinating Layer 3 VRRP redundancy with Layer 2 MSTP redundancy, to avoid strange network outages and unpredictable behavior. Then you will apply this knowledge with a lab activity

VRRP

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

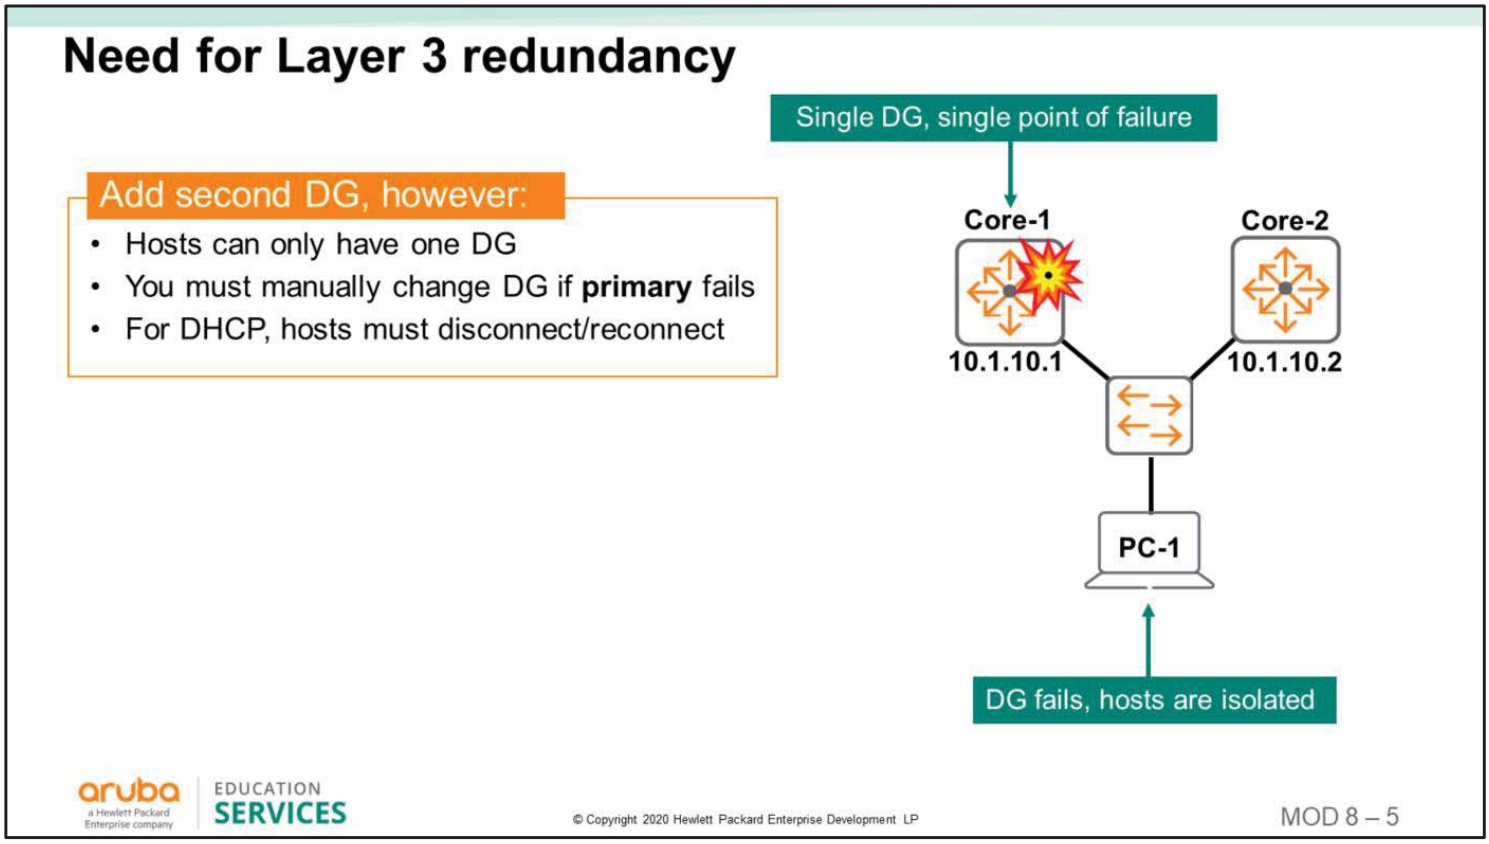

In the previous module you learned about the benefits a Default Gateway (DG) for endpoints. Recall that an endpoint may only have one DG, and a single DG means a single point of failure. In the example, if Core-1 fails, PC-1 and all other hosts using it for the DG are now isolated. What can you do?

You could add another DG for redundancy. This seems like an easy solution...but maybe not. Each endpoint's DG is either configured manually or obtained via DHCP. When Core-1 fails, you must either manually reconfigure each host with a new DG, or reconfigure the DHCP scope, and then ask all end users to disconnect from the network and reconnect, or power cycle their PC, or teach them how to trigger a DHCP release and renew action. (In a Windows command prompt for example, use ipconfig /release, then ipconfig /renew). However you do it, it is not very elegant, or scalable, and disruptive for end users.

There must be a better way, right?

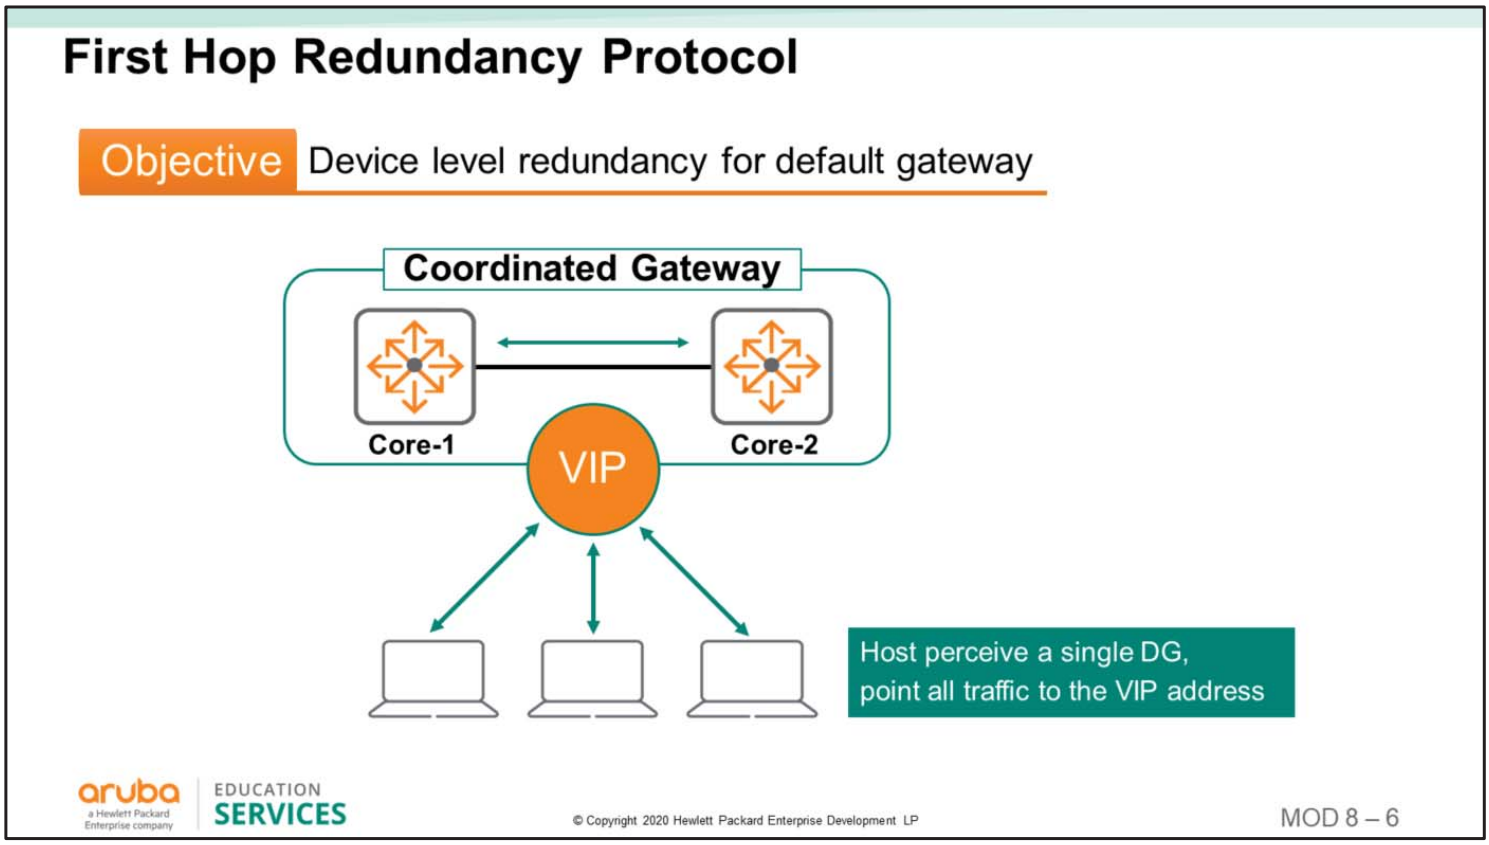

The real solution for this challenge can be resolved from the network side instead from the endpoint side, using some kind of First Hop Routing Protocol (FHRP). This adds resiliency for endpoints by using a coordinated gateway solution – no change to endpoint IP configuration, no DHCP modifications, and no end user disruption. It is automatic!

Note. The term resiliency refers to the ability of the network to adapt to changes and failures.

An FHRP solution creates a single coordinated gateway from two or more physical routers. The two physical routers present themselves to endpoints as a single device, with a single Virtual IP (VIP) address. This VIP acts as the endpoint's DG.

Normally, the Primary routing device serves the DG role, forwarding traffic for endpoints. The Secondary unit monitors the Primary device state. If the Primary fails, the Secondary device takes over. It takes on the Primary role and VIP and forwards endpoint traffic.

From the endpoint perspective the Virtual IP address is always available. There is no DG address change, and there is no disruption for end users.

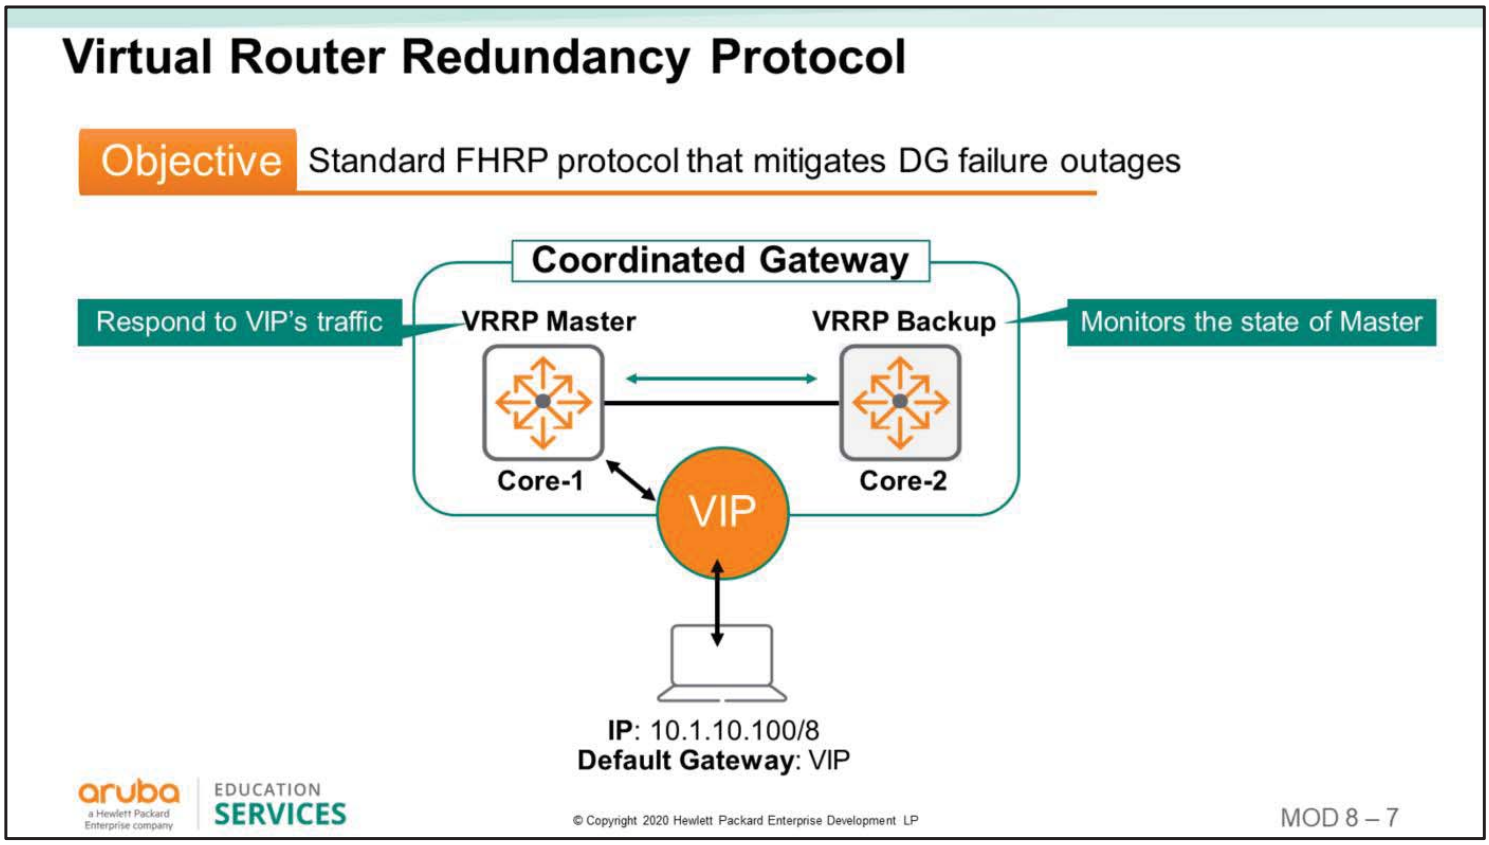

RFC 5798 defines the Virtual Router Redundancy Protocol (VRRP) - a standard FHRP that enables two or more routing devices to provide gateway redundancy for endpoints.

VRRP uses a Master-Standby architecture. Only one gateway actively forwards traffic sent to the VIP address. This primary forwarding device is called the Master, while the non-forwarding device is the Backup.

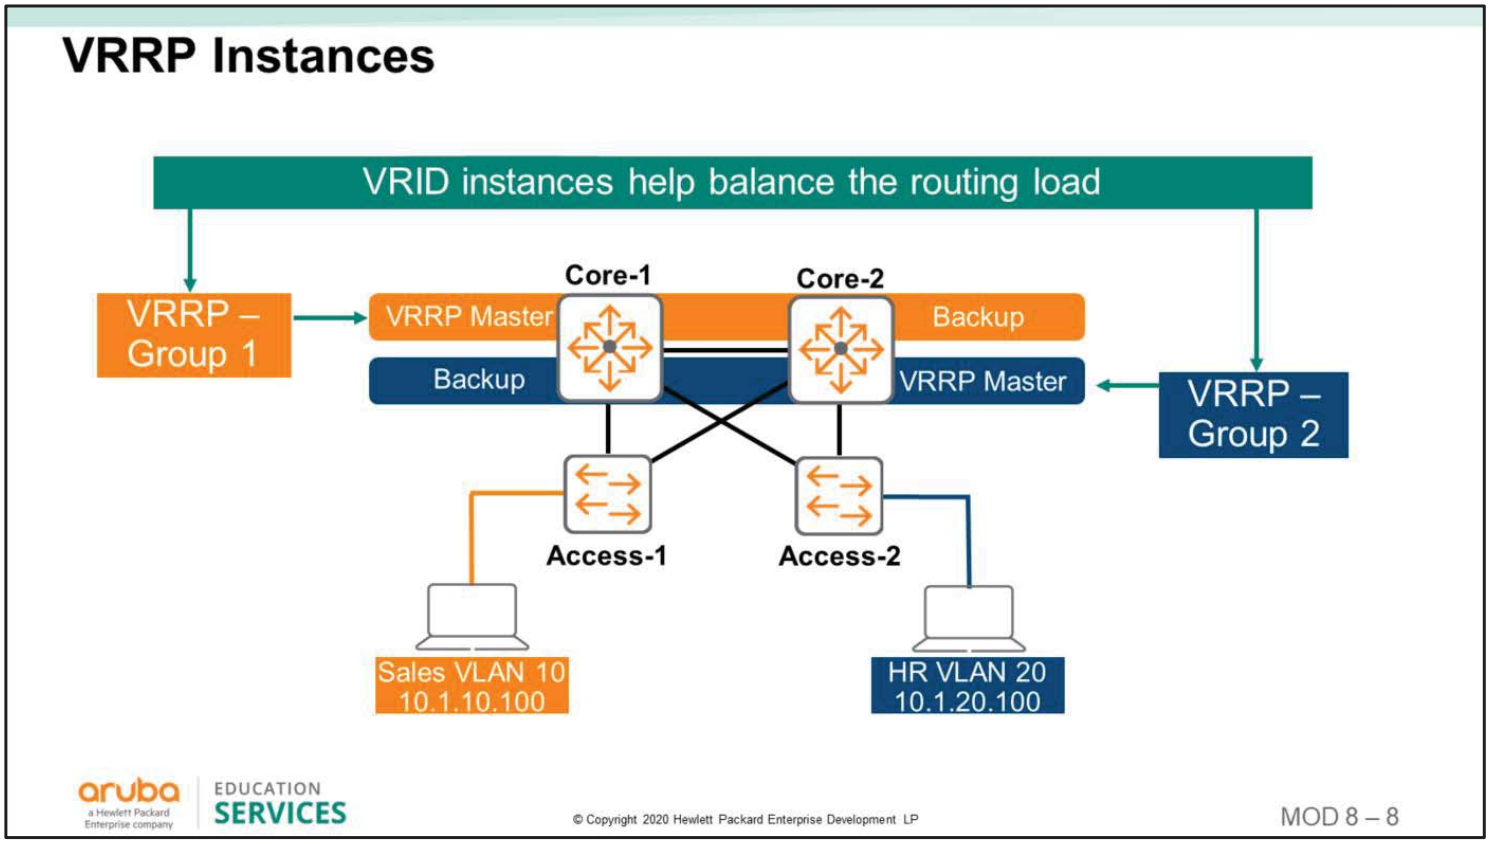

ArubaOS-CX allows you to deploy multiple instances of VRRP, often to balance the load for VLANs, as shown in the figure. Each instance has a unique Virtual Router ID (VRID) number, ArubaOS-CX refers as Group ID as shown in the figure. VRRP instance 1 serves VLAN 10, while instance 2 serves VLAN 20.

Switch Core-1 is the Master or active forwarder for VLAN 10, with Core-2 as the Standby. Meanwhile, Core-2 is the Master for VLAN 20, with Core-1 as the Standby. This gives you a nice load balancing capability.

VRRP Instances – Validate Capacity

```
6300# show capacities vrrp

System Capacities: Filter VRRP
Capacities Name                                                                    Value
-----------------------------------------------------------------------------------------
Maximum number of unique IPv4 VRRP VRIDs configurable between 1 to 255               8
Maximum number of VRRP IPv4 addresses supported                                   1024
Maximum number of VRRP IPv4 addresses supported per virtual router                  16
Maximum number of VRRP IPv4 virtual routers supported per port                       8
Maximum number of VRRP virtual routers supported                                   256
<<<<Omitted output>>>>
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 8 – 9

Instances are also known in VRRP as Virtual Router IDs (VRIDs), the number of available VRIDs that a switch supports depends on the switch capacity. In ArubaOS-CX you can verify the capacities of your switch using the show capacities vrrp command.

Master Election

Master election based on priority

- 1- 255, highest priority wins
- ArubaOS-CX default priority = **100**
- Same priority - highest IP address is the winner

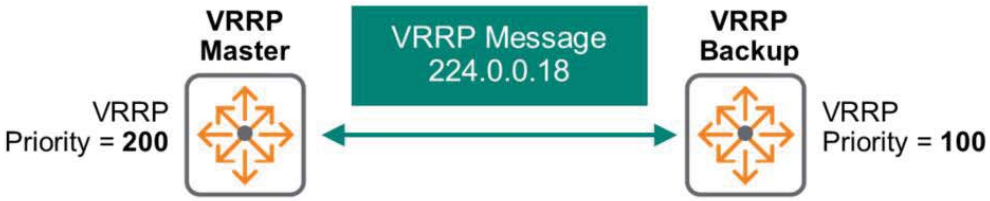

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 8 – 10

VRRP members exchange multicast advertisement messages to elect the Master gateway, using address 224.0.0.18, IP protocol number 112.

To control the Master election, you set a priority value from 1 – 255 – highest priority wins. If both devices have the same priority, then the gateway with the highest IP address wins the election.

In ArubaOS-CX the default priority value is set to 100.

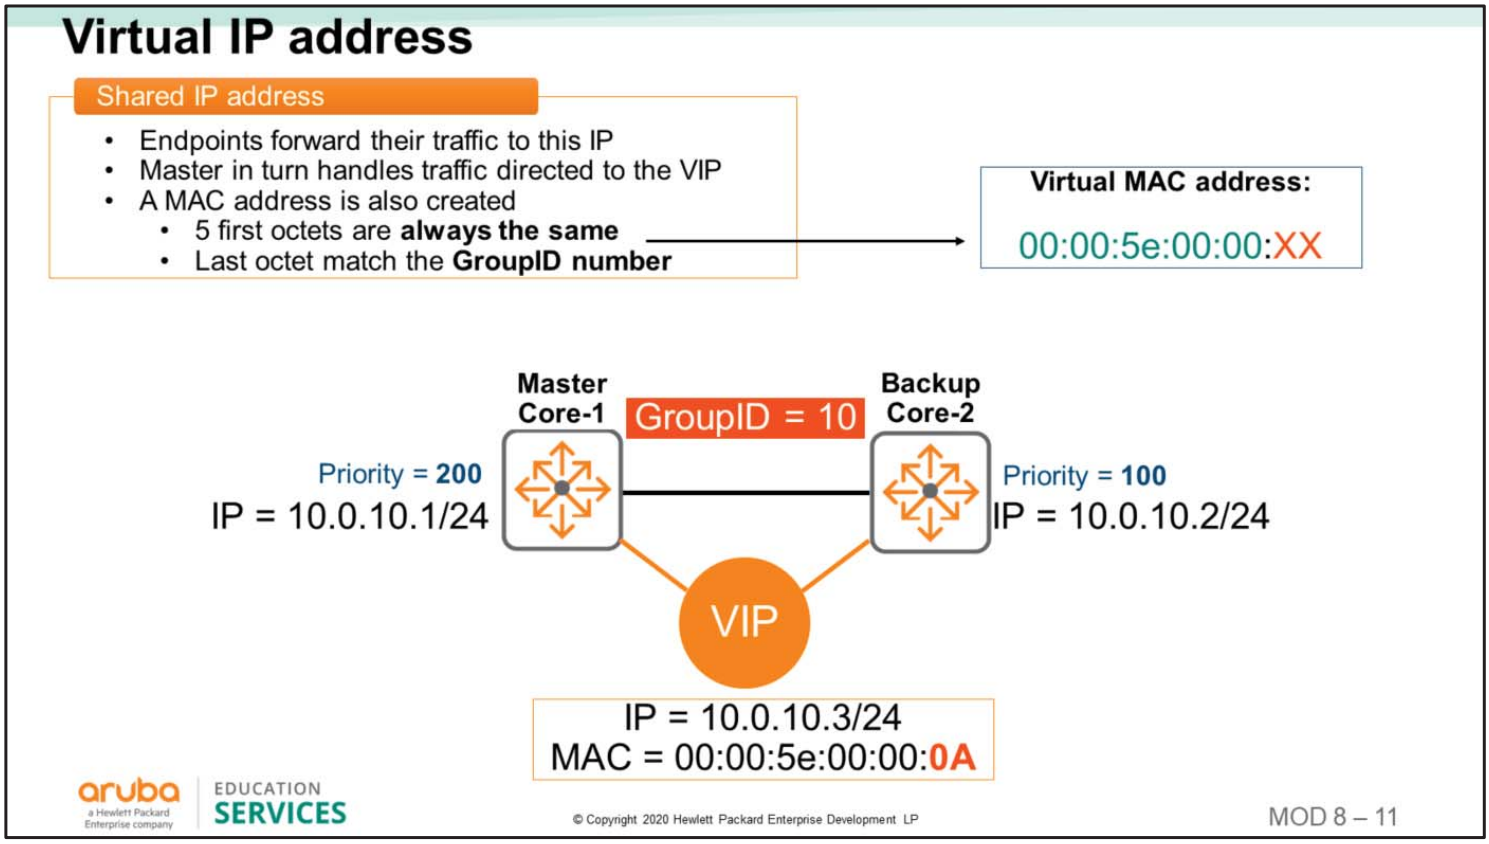

The Virtual IP (VIP) address is the result of the Gateway coordination. You assign a unique, “real” IP address to each individual physical gateway, as normal. The VIP address must also be unique. In the figure, you assigned 10.0.10.1 to the Master, 10.0.10.2 to the Backup, and 10.0.10.3 is used as the VIP.

Each VRRP instance has a unique VRRP address. 10.0.10.3 might be the VIP address for Instance 10, used for VLAN 10. meanwhile, 10.0.20.3 might be the VIP address for Instance 20, used for VLAN 20.

You must ensure that this VIP address is configured as the endpoint Default Gateway. Thus, endpoints forward their traffic to the VIP. The VRRP Master receives these packets and routes them. If the Master fails, the Standby unit takes over. The devices do not learn about the router’s physical IP addresses - 10.0.10.1 and 10.0.10.2 in the example.

A virtual MAC address (vMAC) is automatically assigned to the VIP. As defined in the standard, this address shall be 00:00:5e:00:00:XX, where XX = the VRID. In the figure, the VRID = 10, and so the vMAC is 00:00:5e:00:00:0A (Hex A = 10 in decimal). Thus, when endpoints ARP for their DG address of 10.0.10.3, they will learn the vMAC address, and add it to their ARP table.

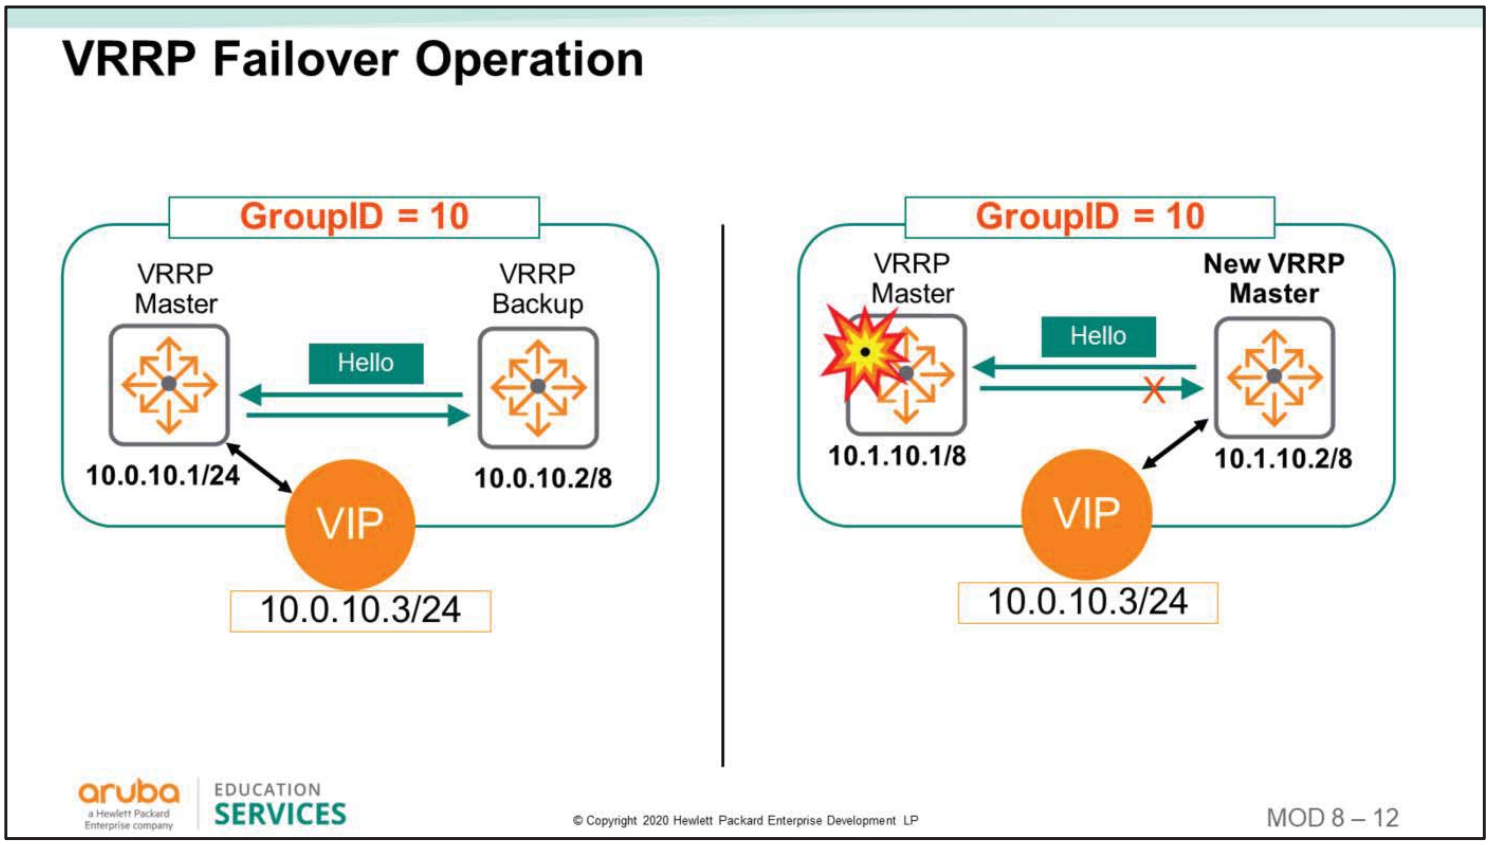

The image shows a scenario where two Core switches run VRRP. Core-1 has a higher priority, and so is the VRRP Master. Core-2 is the VRRP Standby - constantly monitoring Core-1 status via a keepalive mechanism. Endpoints forward their traffic to 10.0.10.3, which is services by the VRRP Master.

Then the Master fails. The standby stops receiving keepalive messages, and so knows that the Master is down. The former Standby takes over as the new VRRP master and begins to forward traffic for VIP 10.0.10.3.

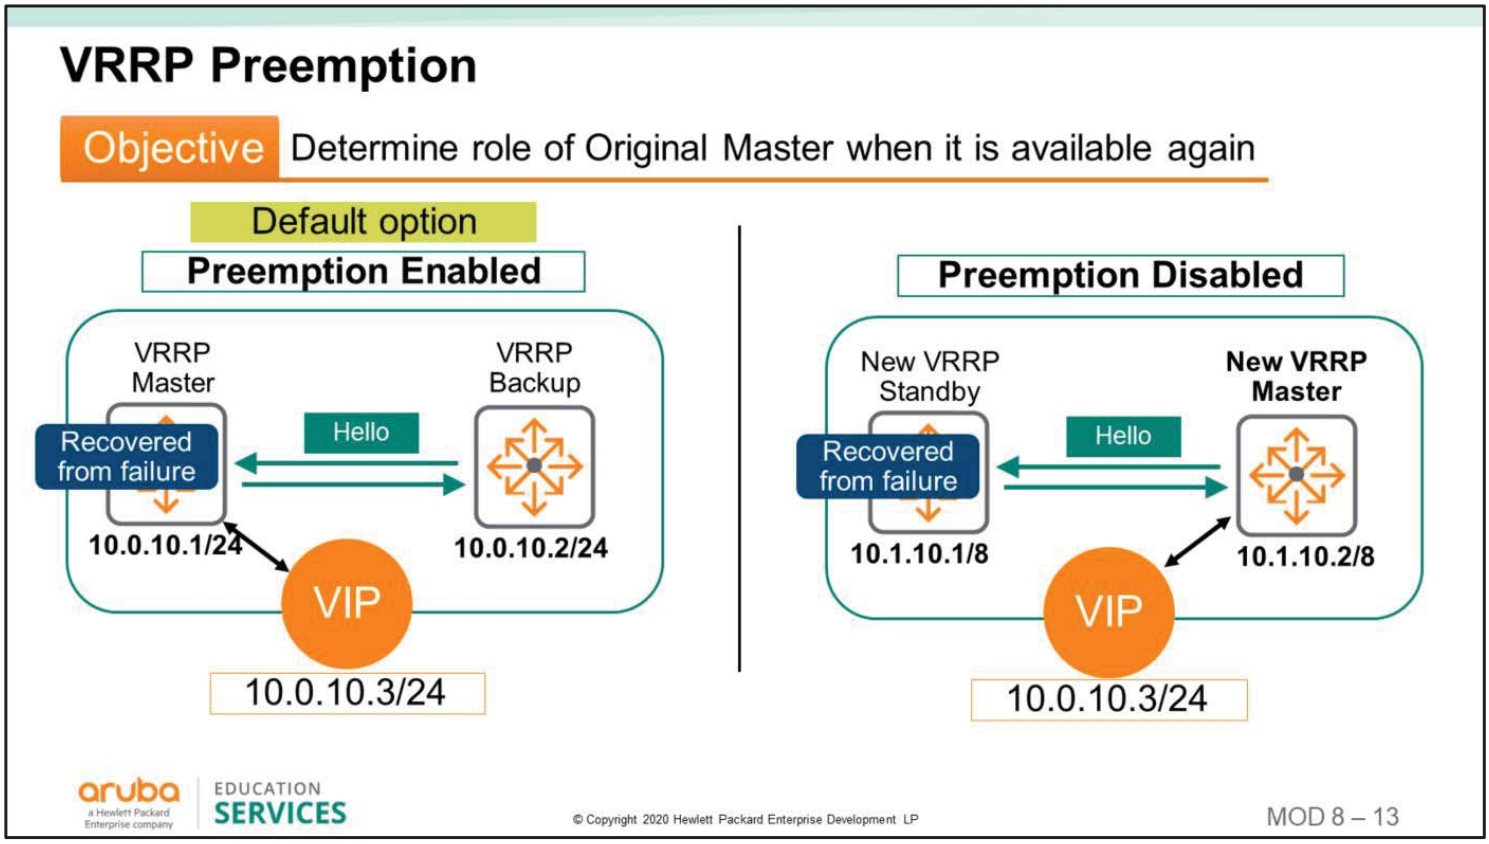

If you continue to review the previous figure and failover operation, you can note that that Core-1 failed, and so Core-2 took over as the new Master. What happens when Core-1 comes back online? This depends on how you configure VRRP preemption.

If Preemption is enabled, then Core-1 will reassume its original Master role and Core-2 reverts back to its Standby role. This is the ArubaOS-CX default setting. It is nice because you know that under normal operating conditions, when all devices are up, the same router always acts as the Master. This can be especially important if you are using multiple VRRP instances. If a Master fails, the remaining device(s) will carry the load for all endpoints, while the former Master, once again operational, remains unused.

If Preemption is disabled, then Core-1 will not resume its original Master role. Core-2 remains the Master, and Core—1 takes on the Standby role. You must manually disable preemption with the command no prempt. With preemption disabled you lose the benefits described above. Some administrators might choose to disable preemption if they are worried about the (very brief) time lag that might occur during the preemption process – while the routers switch back over to new operational states, potentially during the middle of a busy day. Due to the high-performance nature of ArubaOS-CX devices, this is rarely a concern.

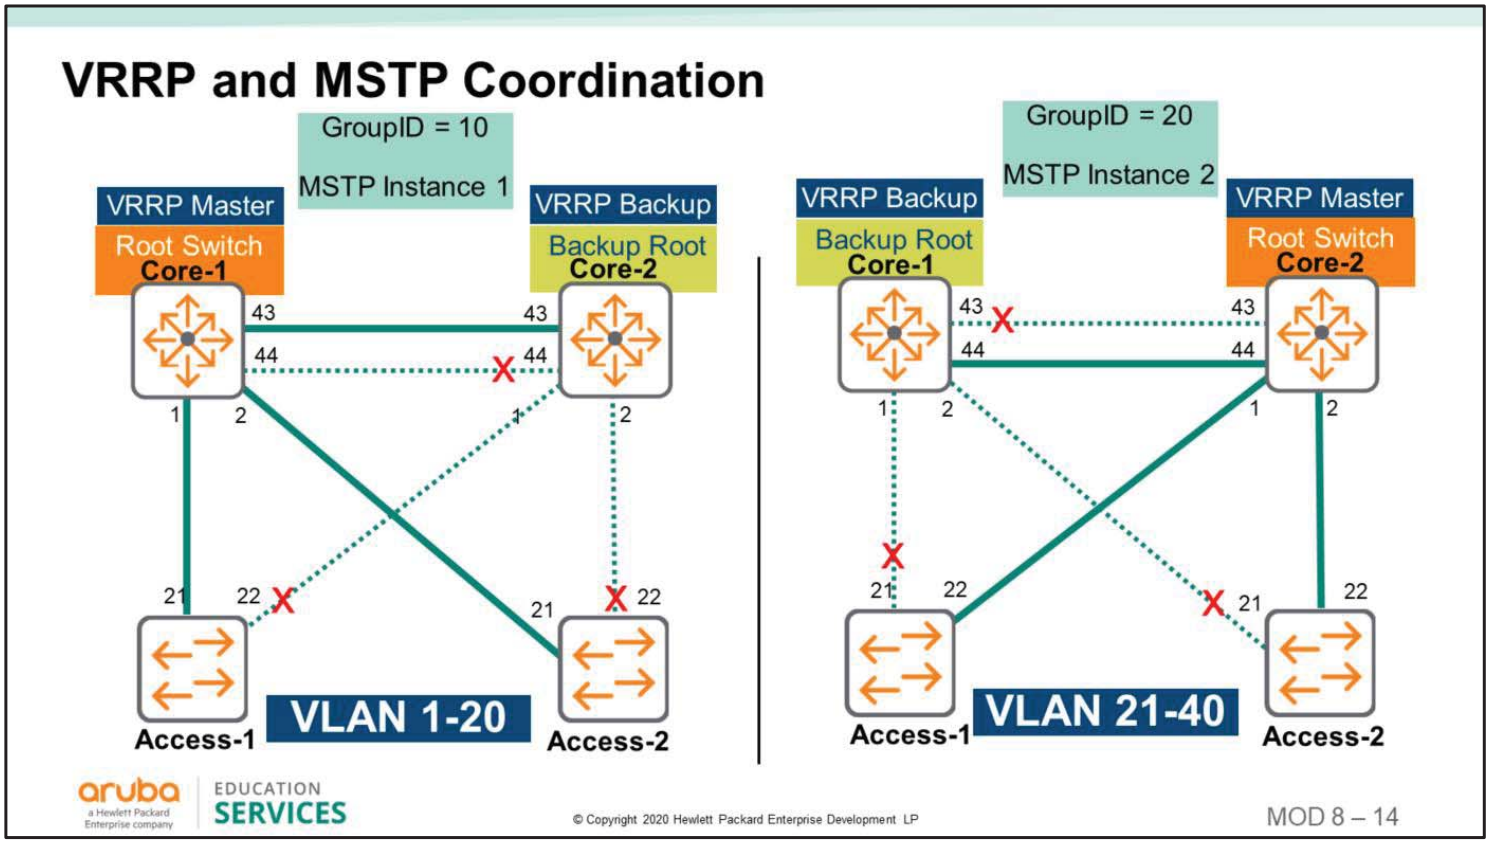

You may have noticed in the previous scenario that a Layer 2 Loop can be created between Access-1 and the Core switches. You know how to avoid this – enable Spanning Tree.

In these scenarios there must be coordination between the MSTP Root Bridge and the VRRP Master. This ensures proper traffic forwarding - Layer 2 and Layer 3 protocols converge with the same devices.

For example, Core-1 is configured as the Root Bridge for MSTP Instance 1 which supports VLANs 100 to 199. Meanwhile, Core-1 is also configured as the VRRP Master for this same VLAN range. If a failure occurs on Core-1, Core-2 becomes the new VRRP master and the new Root Bridge for Instance 1. Thus, both Layer 2 and Layer 3 protocols are coordinated – L2 STP uses the same forwarding path as L3 routing. This is vital to avoid unexpected behaviors.

Question #1

VRRP, Master Election,

Which of the statements below accurately describe VRRP concepts and operation?

a. A unique VRRP instance can run in an Aruba Switch.

b. To control VRRP master election between two routers, you must configure a priority on both devices.

c. If two VRRP routers have the same priority, the router with the highest IP address becomes the VRRP Master.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 8 – 15

Correct Answer:
c.

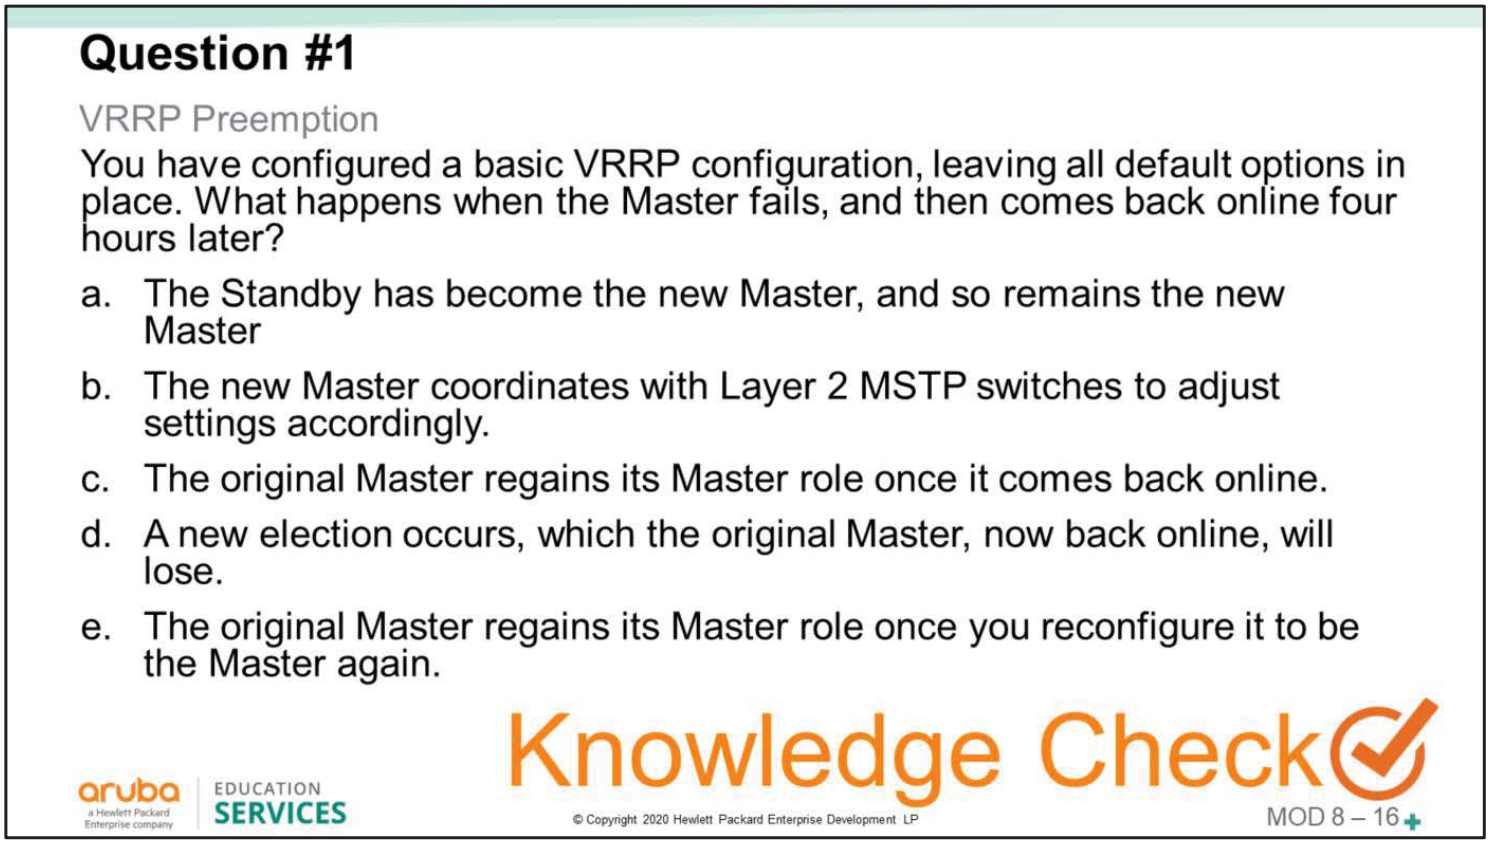

Correct answer:
c.

Lab Activity

8.1 Deploy VRRP

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

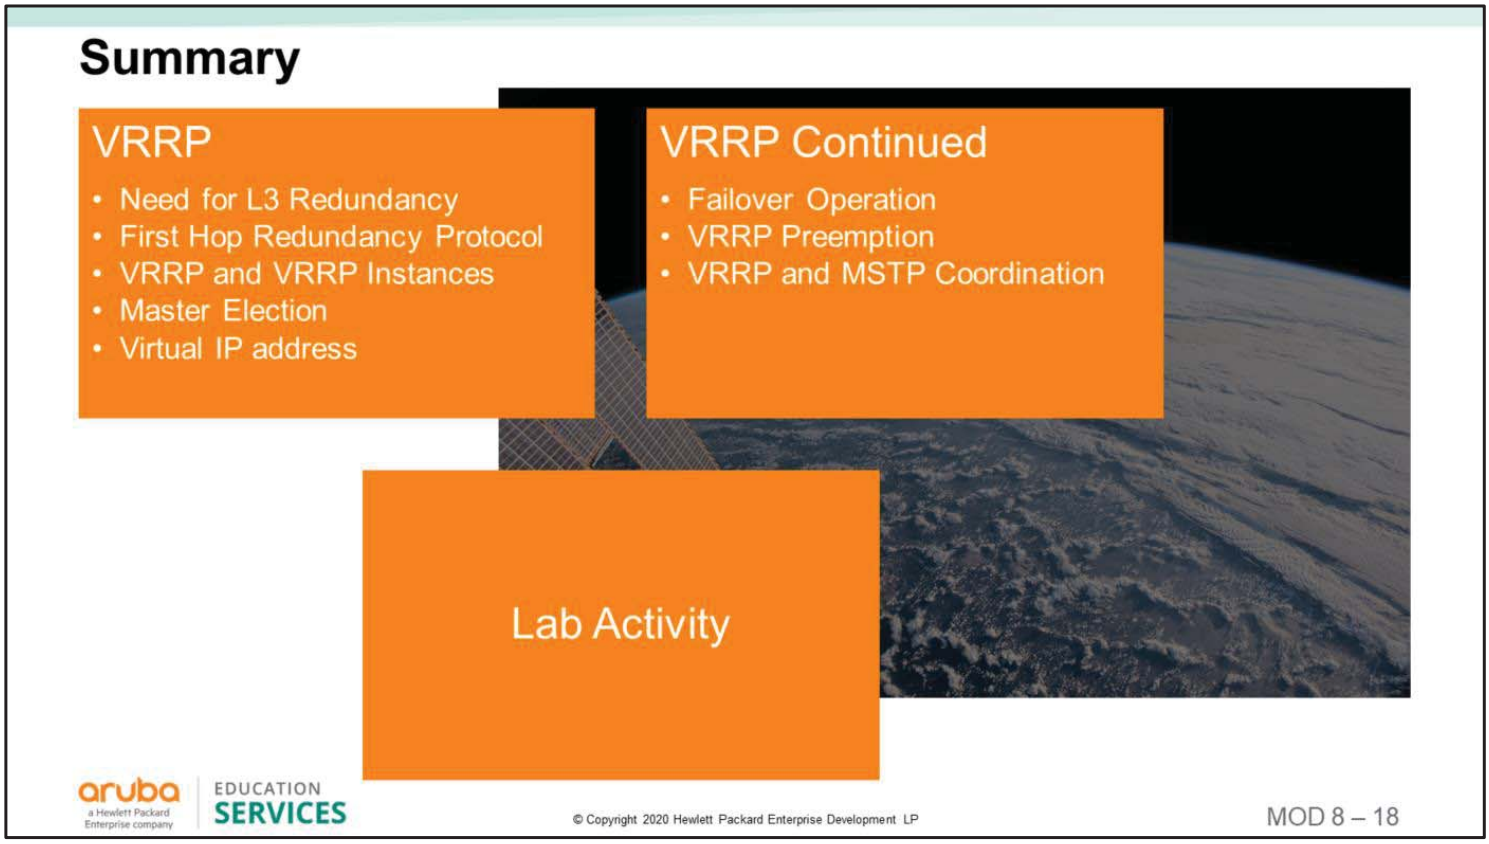

In this module, you learned how to eliminate a potentially costly single point of failure – the endpoint's Default Gateway (DG). You came to understand how a single DG can cause major outages, and why you cannot simply add a second DG.

You explored the Virtual Router Redundancy Protocol (VRRP), and how it automatically mitigates downtime upon DG failures. You learned about VRRP operation, including how to load balance endpoint traffic by configuring multiple VRRP instances on a single set of routers.

You learned about the Master election process and how to control it by configuring priority values. You also explored preemption options, to control what happens when a failed Master comes back online.

Finally you learned about the importance of coordinating Layer 3 VRRP redundancy with Layer 2 MSTP redundancy, to avoid strange network outages and unpredictable behavior.

Then you applied this knowledge with a lab activity.

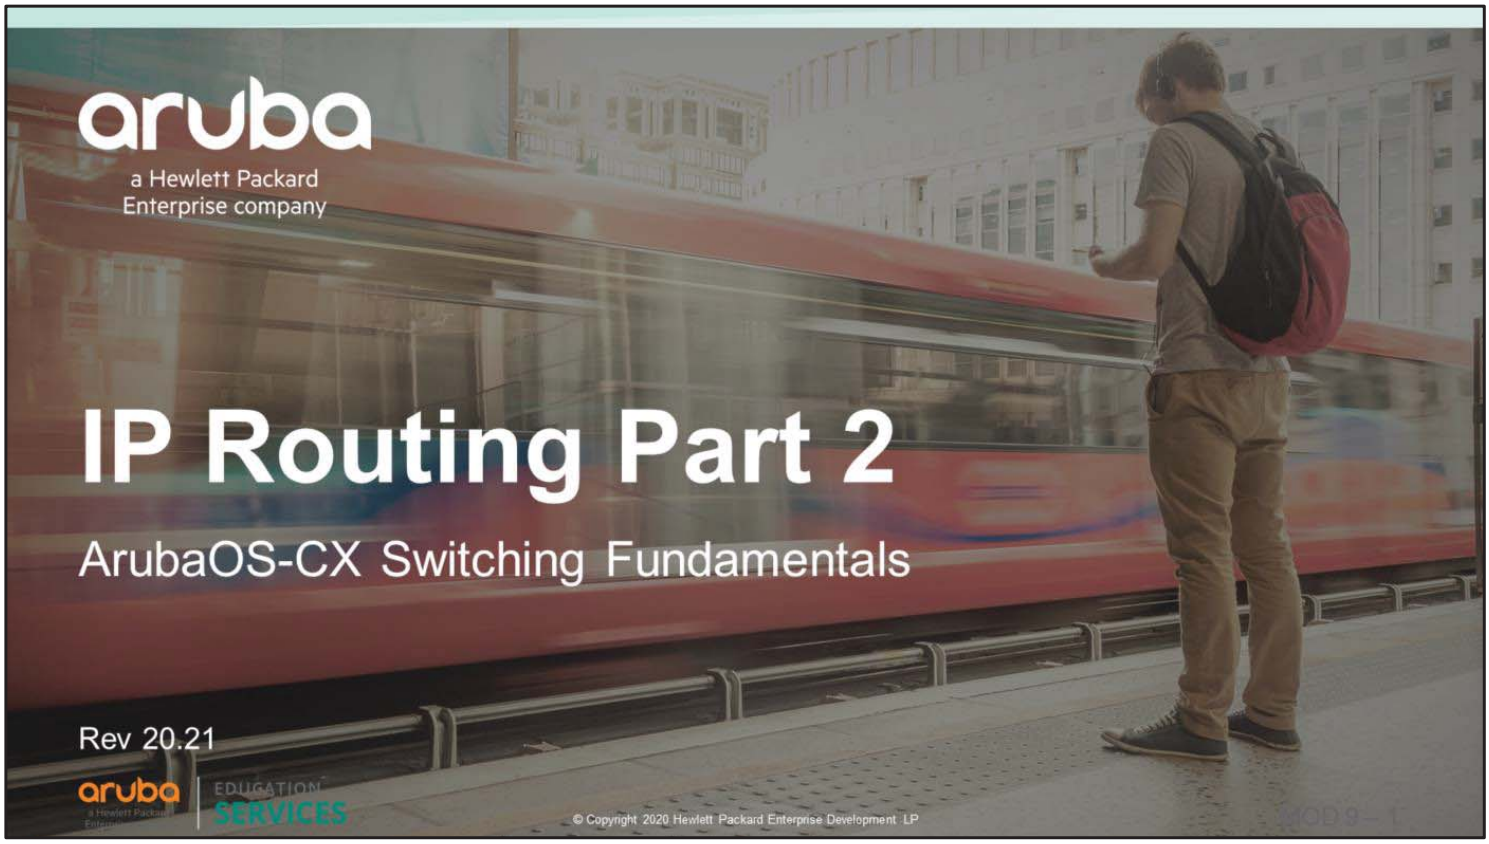
aruba
a Hewlett Packard Enterprise company
IP Routing Part 2
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP

Objectives

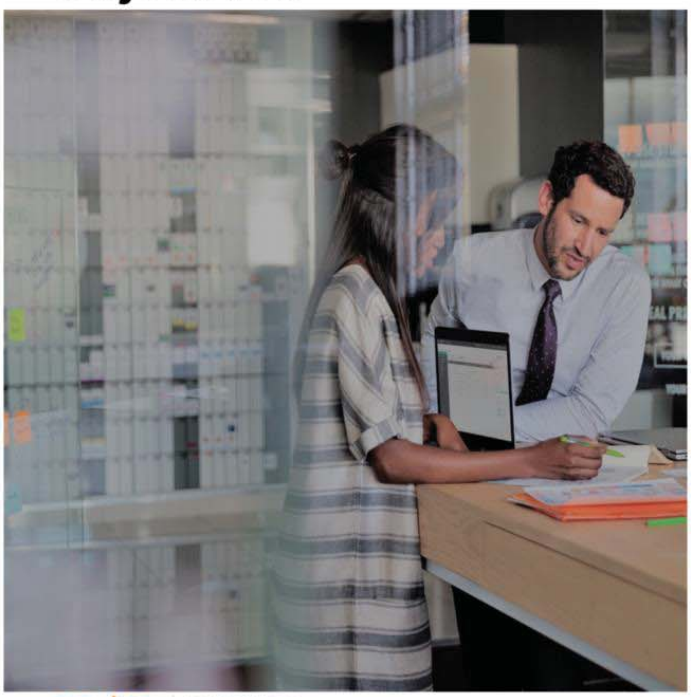

–Learning Objectives

- Describe and perform subnetting
- Explain the advantages of VLSM and CIDR

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 2

After completing this module, you will be able to leverage your knowledge of binary and decimal number systems to describe and perform subnetting.
You will also be able to explain the advantages of Variable Length Subnet Masking (VLSM) and Classless Interdomain Routing (CIDR).

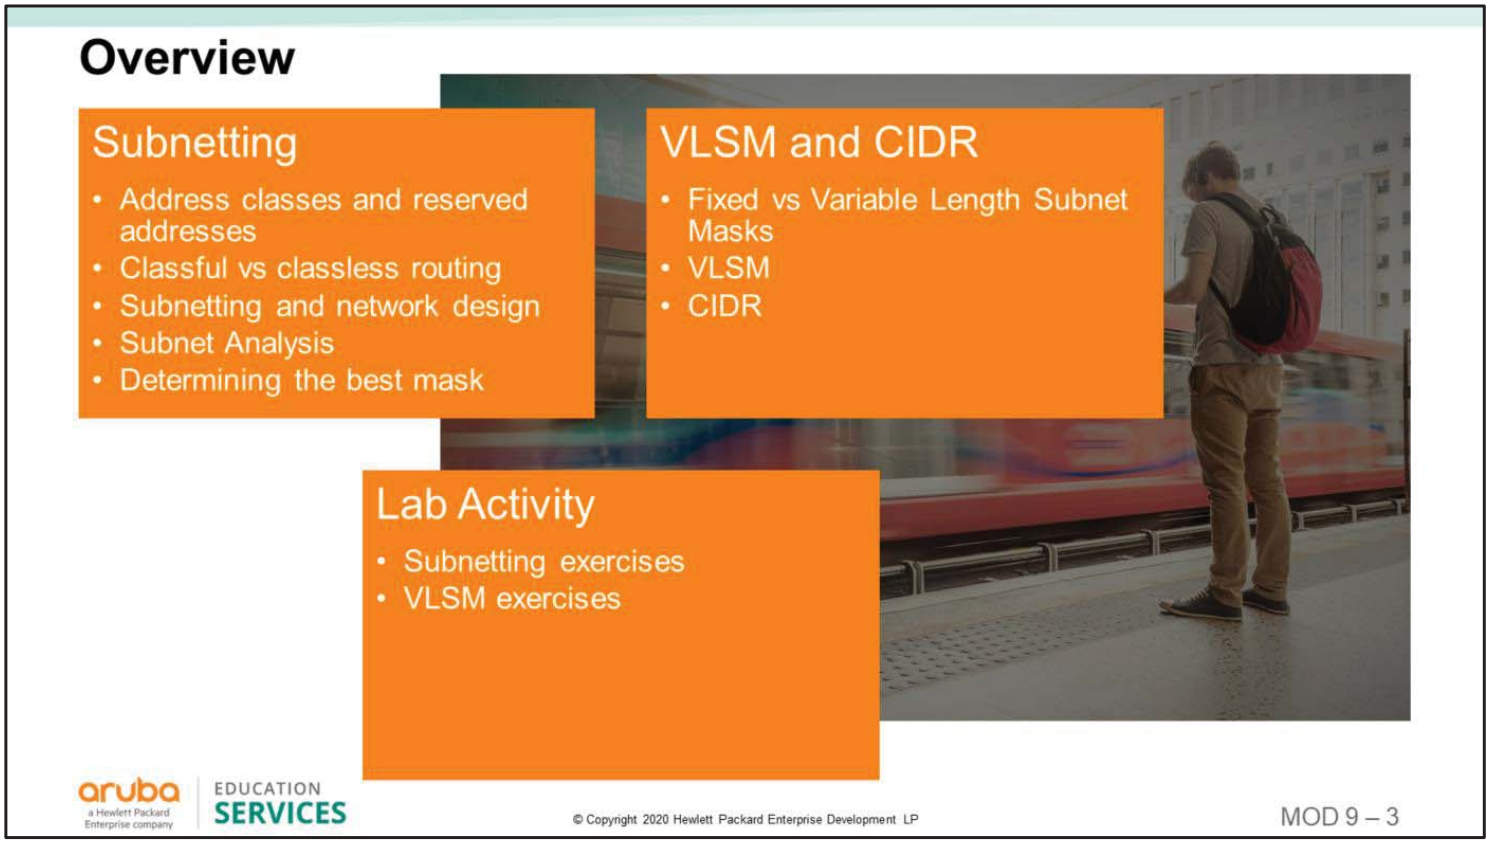

This module should serve to significantly elevate your ability to design and deploy more complex, scalable IP networks. You will leverage your knowledge of binary and decimal number systems as you explore subnetting – how to more efficiently use available IP address. Related concepts include the various address classes and reserved address space, and why classful routing can seem quite rigid when compared to classless routing. Now you can apply your subnetting knowledge to network design and analysis scenarios.

Then you will explore Variable Length Subnet Masking (VLSM) and Classless Interdomain Routing (CIDR). VLSM help you to assign address space more efficiently, while CIDR helps the routers to advertise and work with that address space more efficiently.

Finally, a lab activity will give you more experience with these concepts.

Subnetting

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

IPv4 Address Classes

Network Class	Address scope	Leading bits	Description
Class A	0.0.0.0. to 127.255.255.255	0000	Unicast address
Class B	128.0.0.0 to 191.255.255.255	1000	
Class C	192.0.0.0 to 223.255.255.255	1100	
Class D	224.0.0.0 to 239 255.255.255	1110	Multicast address
Class E	240.0.0.0 to 255.255.255.255	1111	Reserved for future or experimental purposes

Note.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 5

The available IPv4 address space is divided into five classes - Class A to E, which each class designed for a particular purpose. This is summarized in the table, which shows each class and its address range.

It also shows the first few most-significant bits of the address, if the range was converted to binary. For example, Class A addresses always have the most significant bits set to 0000, Class B addresses always begin with 1000, and so on.

Note. This training is focused on Classes A through C, which are use for Unicast addressing. These are by far the most common addresses that you will work with in your career.

The table also shows that the first octet of any address reveals its Network class. This is important for you to know, as it will help you with upcoming concepts and activities, and prepare you for real-world design, administration, and troubleshooting tasks. The main difference between address classes A- C is first octet's range, and the distribution of Network ID and Host ID within the address. Here is how it works:

Class A: The first octet is between 0 and 127. The first octet is reserved for the Network ID, which means that 3 octets (24 bits) are available to assign to hosts on that network.
Class B: The first octet is between 128 and 191. The first two octets (16 bits) are reserved for the Network, which leaves 16 bits available for host assignment on that network.
Class C: The first octet is between 192 and 223. The first three octets (24 bits) are reserved to specify the subnet, leaving only 8 bits available to specify hosts on that subnet.

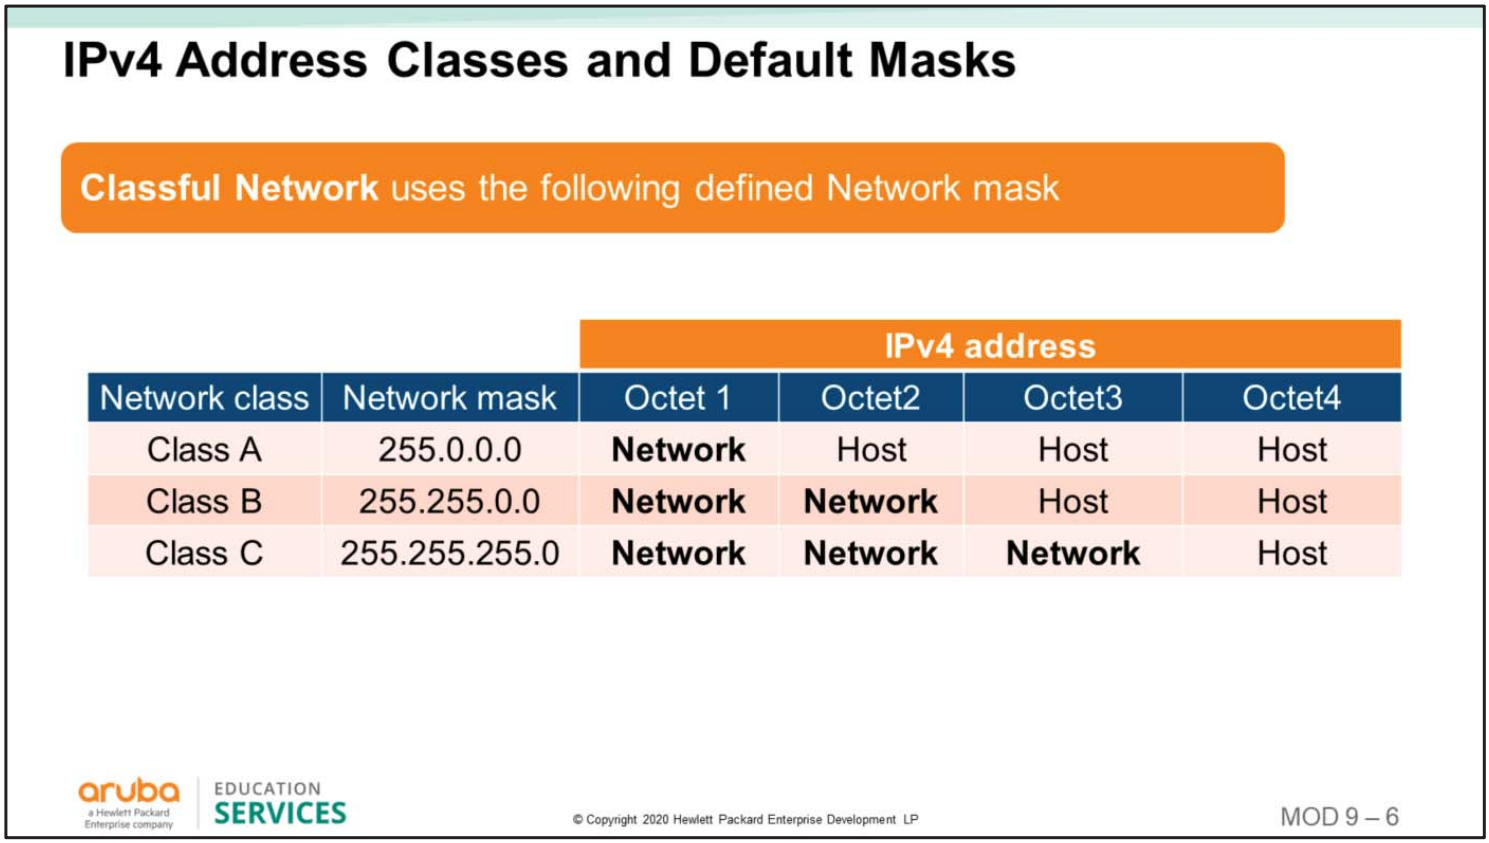

From the previous discussion, you might realize that there is a fixed Network Mask for each class. Think of a **classful network** as a network that uses these default Network Masks for their corresponding Network Class.

As shown in the figure:

- Class A uses 255.0.0.0 Network Mask.
- Class B uses 255.255.0.0 Network Mask.
- Class C uses 255.255.255.0 Network Mask.

For example, a Class A network might be 10.0.0/8. With 24 bits to specify hosts, you can have a lot of hosts! 10.0.0.1 – 10.255.255.254

Look at the Class B example, network 172.16.0.0/16. With 16 bits to specify hosts, you have 172.16.0.1 – 172.16.255.254.

The Class C example is network 192.168.1.0/24, leaving only 8 bits available to specify hosts on that network. You can have 192.168.1.1 – 192.168.1.254.

Notice the numbers in the example. Why doesn't the example assign 192.168.1.0 to a host? Why didn't the example assign 192.168.1.255 to a host? Because these addresses are reserved. They have special meaning.

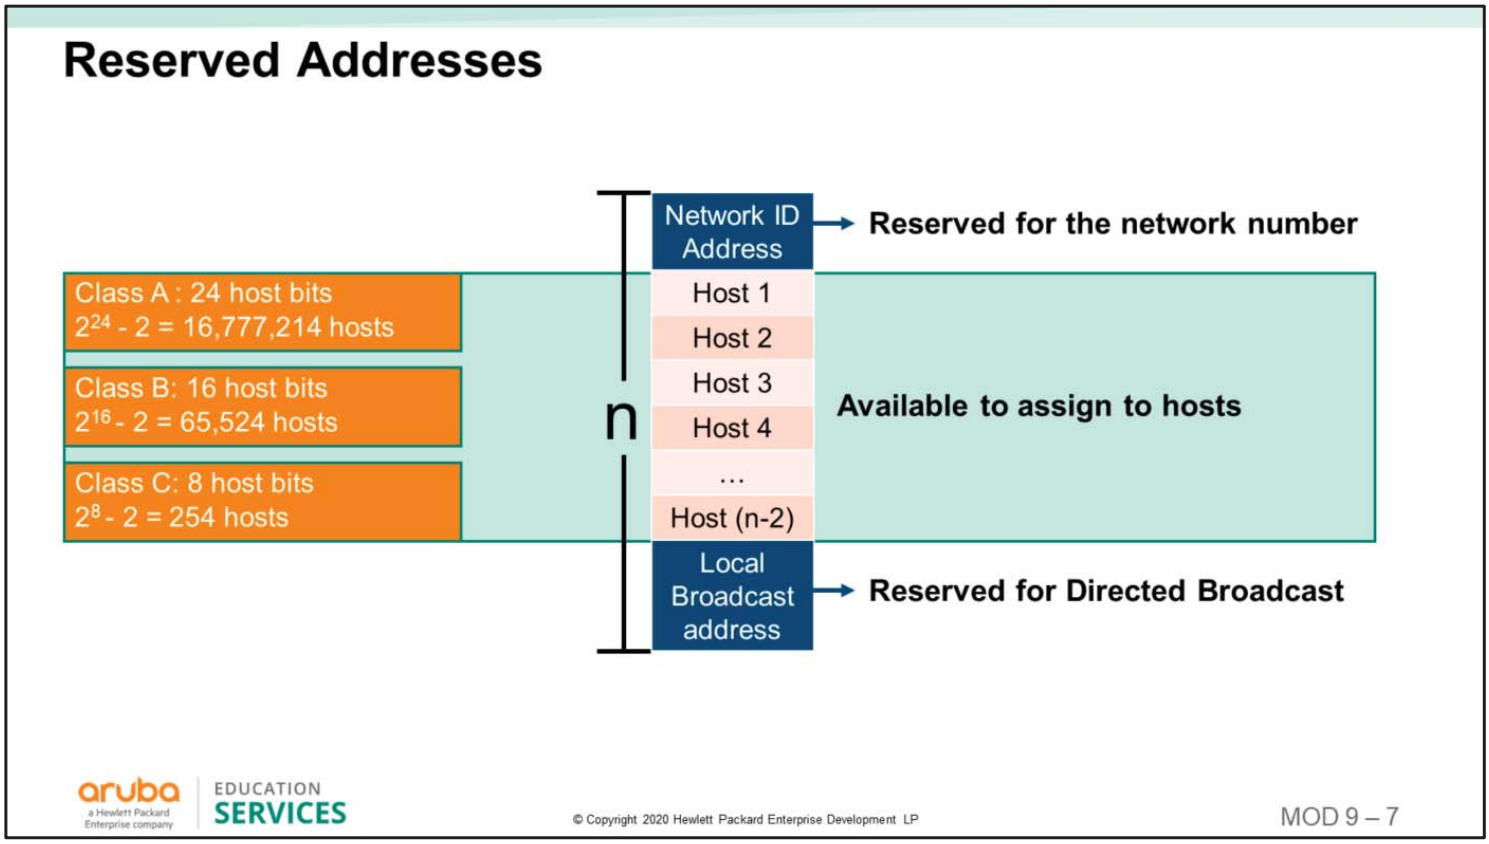

In any network there are two reserved addresses that cannot be assigned to the hosts. The first one is the **Network ID** or **Network Number** and the second one is the **Local Broadcast Address**.
As in the previous Class C example, IP address 192.168.1.0 cannot be assigned to a host – this number is reserved to mean “the network itself”. Assigning this address to a host would be like telling your mail delivery person that your address is “Main Street”. You can’t just live in the middle of the street...you must have a street address, “123 Main Street”. Another way of saying this is that the network number is indicated when all host bits are binary 0s. This reserved network number is used by routers to find the best path to a subnet.

Similarly, you cannot assign the address 192.168.1.255 to a host. This is reserved for the broadcast address. It means, “Attention everyone on this subnet”. Assigning this address to a host would be like having a child and naming her “Everyone”.

So, the first available number is always the network ID (10.0.0.0, 172.16.0.0, 192.168.1.0), and the last available number is always reserved for a directed broadcast (10.255.255.255, 172.16.255.255, 192.168.1.255).

The figure also indicates how many hosts you can have on a particular network, with the variable n. “n” represents how many unique numbers can be created, given the number of available hosts bits. To do this, use the formula 2^n (2 raised to the nth power).
Class A networks only use 8 bits for the network number, so 24 bits are available to specify hosts. 2^24 = 16,777,216 numbers, minus the 2 reserved addresses = 16,777,214 addresses – that is how many hosts you can have on a Class A subnet.

Class B networks use 16 bits for the network number, leaving 16 bits available to specify

hosts. 2^16 = 65, 536 – 2 = 65,534 hosts on any Class B network.

Class C networks use up 24 bits to specify the network number, leaving only 8 bits available to specify hosts. 2^8 = 256 – 2 = 254 hosts on a Class C network.

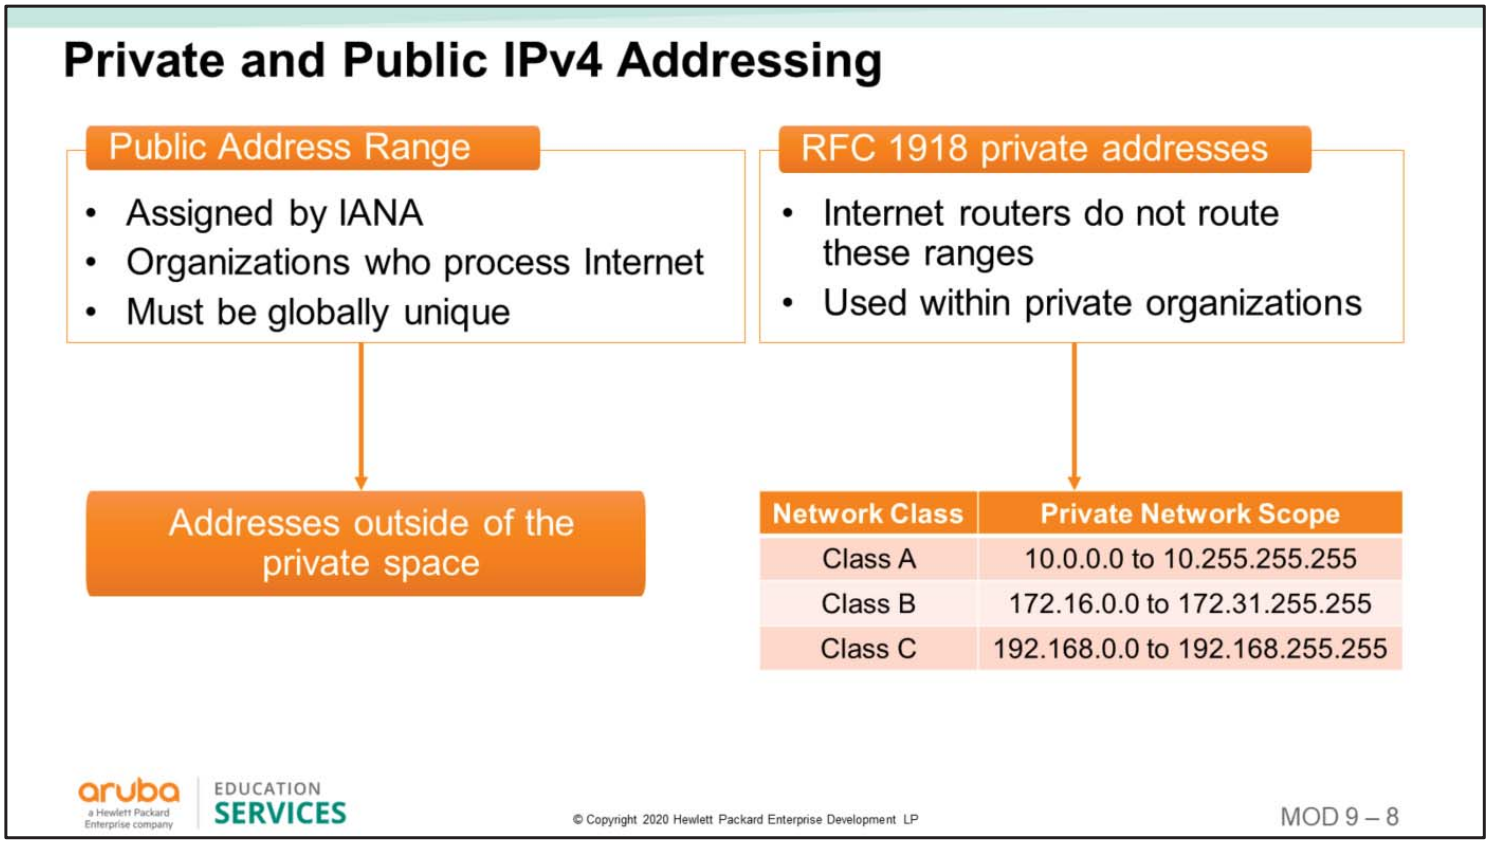

The IPv4 addressing scope is divided in private and public addresses. Typically the public address space is assigned by the Internet Assigned Numbers Authority (IANA), for organizations that need to process packets from the Public Internet (typically Internet Service Providers). This is globally unique address space.

The private address space is defined in the RFC 1918 document published by the IETF organization (https://tools.ietf.org/html/rfc1918). This address space is reserved for private organizations to use within the confines of their organization. Traffic with these addresses should never be sent to the Internet, nor ever be seen by the Internet. If an Internet router within your ISP sees these as destination IP addresses, they will not route that traffic. It will be discarded, because it is not globally unique.

RFC 1918 defines three private address ranges, as shown in the figure.

Note: If this private traffic cannot be routed on the connected Internet, how to private organizations who use this space connect to the internet? Later you will learn about Network Address Translation (NAT), which solves this issue. NAT converts private addresses into public addresses before forwarding packets to the connected Internet.

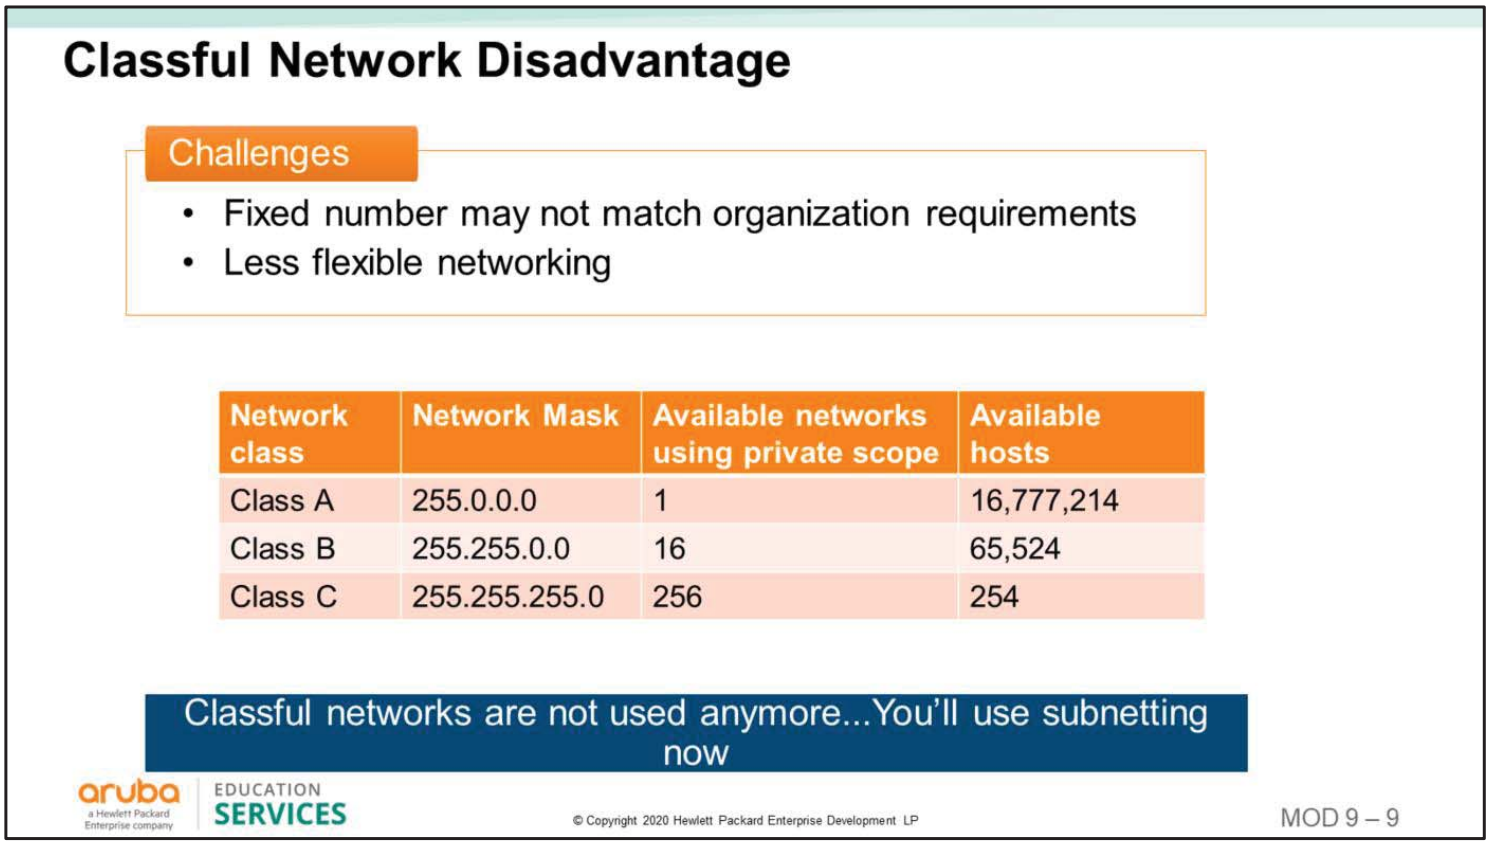

When you deploy a Classful network, you can allocate a fixed number of hosts for each network type in this case, as previously calculated, and shown in the figure. This approach has some disadvantages:

- The fixed host number for each class will unlikely accommodate organizational needs. Suppose that you need a network to allocate 2,000 endpoints. Of course, a Class B network will do the job, but what would you do with the remaining 63 thousand addresses? They are wasted.
- Classful networks restrict the use of the number of networks and hosts that can be used. Imagine that you must set up 50 networks with 2000 on each network. This is simply not possible, since none of the Network classes accomplish this task.

Thankfully, this is only an issue if you use the default network masks shown in the figure. You can resolve many of these issues by changing the default masks, using a process called subnetting.

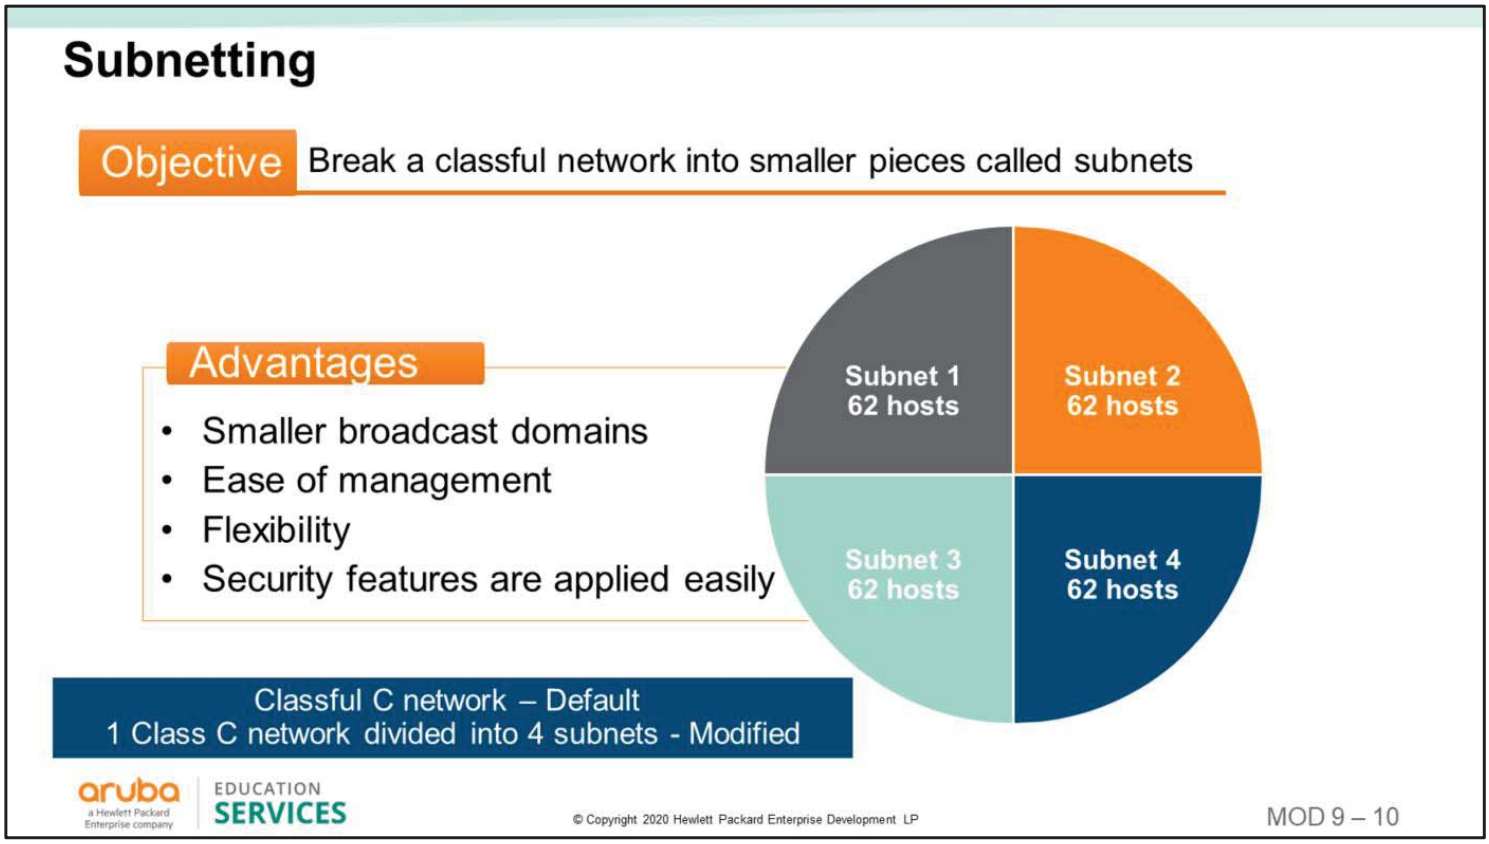

Subnetting enables you to break up a single classful network into smaller pieces called subnetworks – or simply, subnets. Thus, you can create logical network numbers to accommodate the physical design of your network infrastructure. The advantages include

- Smaller broadcast domains: Reduces broadcast overhead, potentially increasing performance.
- Ease of management: You create the number of networks and hosts that you need.
- Flexible network: Addressing can change and grow along with your business.
- Security features are applied easily: You can segment traffic and apply unique security to each subnet.

The image shows how a single Class C network with 254 hosts is broken up into four subnets, each with 62 hosts. You might notice that 62 times 4 is 248 and you have six missing addresses, those address are assigned to the Network and Local Broadcast addresses.

OK, let’s explore how this works.

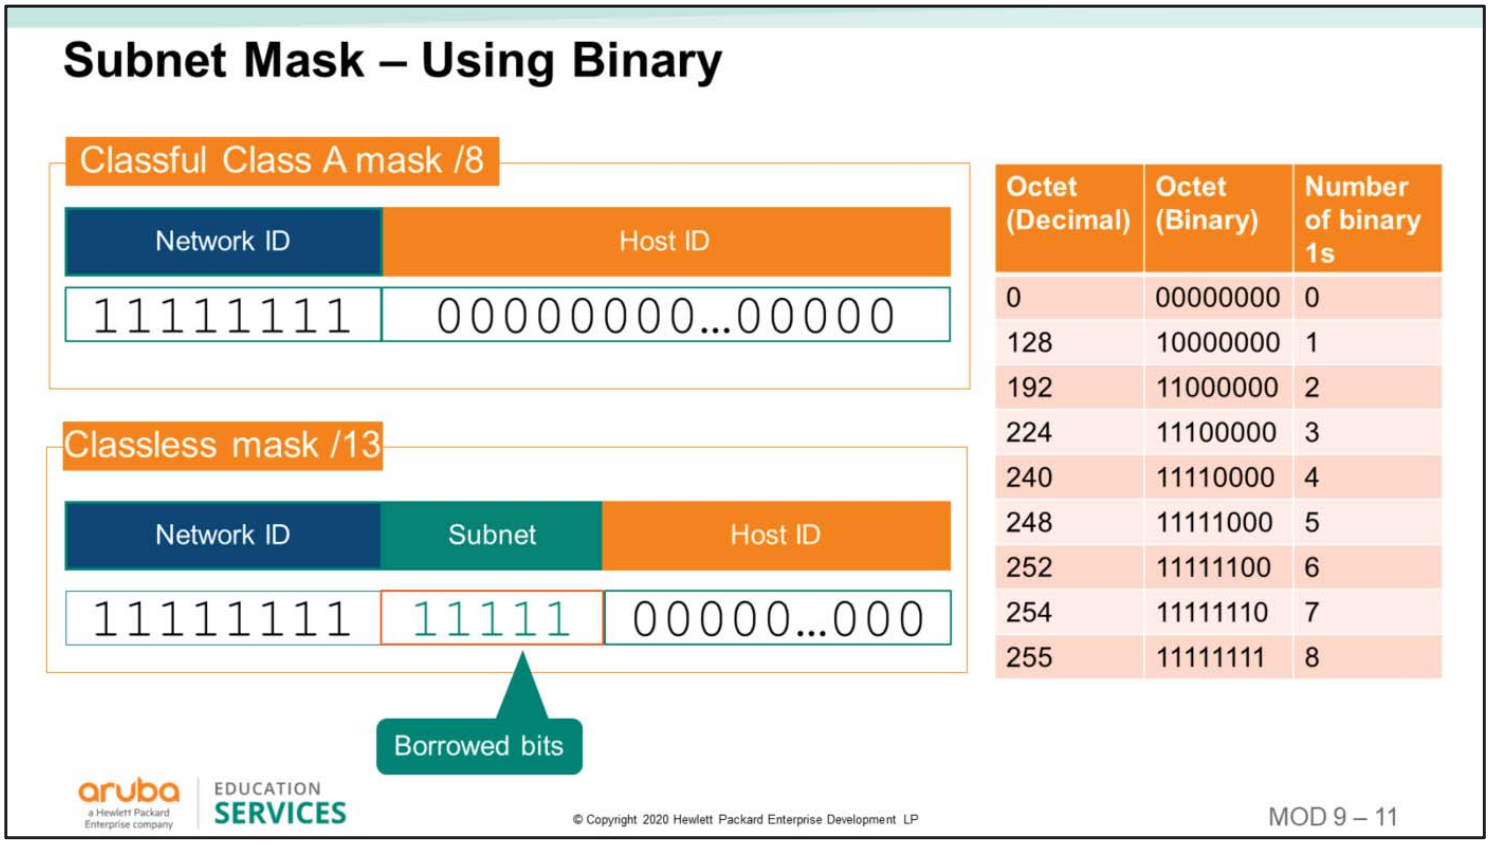

Octet (Decimal)	Octet (Binary)	Number of binary 1s
0	00000000	0
128	10000000	1
192	11000000	2
224	11100000	3
240	11110000	4
248	11111000	5
252	11111100	6
254	11111110	7
255	11111111	8

To accomplish the task of dividing a network, you must use a custom network mask. This mask borrows bits from the host portion of an IP address to create subnets. This mask is known as the subnet mask. This mask is known as a Classless mask.

Notice that each reference to the term “borrow” means that the bits must be set to one. Also, a subnet mask follows the same rule as the Network mask, its format is a sequence of **ones** follow by a block of **zeros**.

A Network mask can have two notations: dot-decimal and prefix. Understanding subnetting will require for you to quickly convert a mask from one notation to the other. The simplest way to do this is by memorizing the following table.

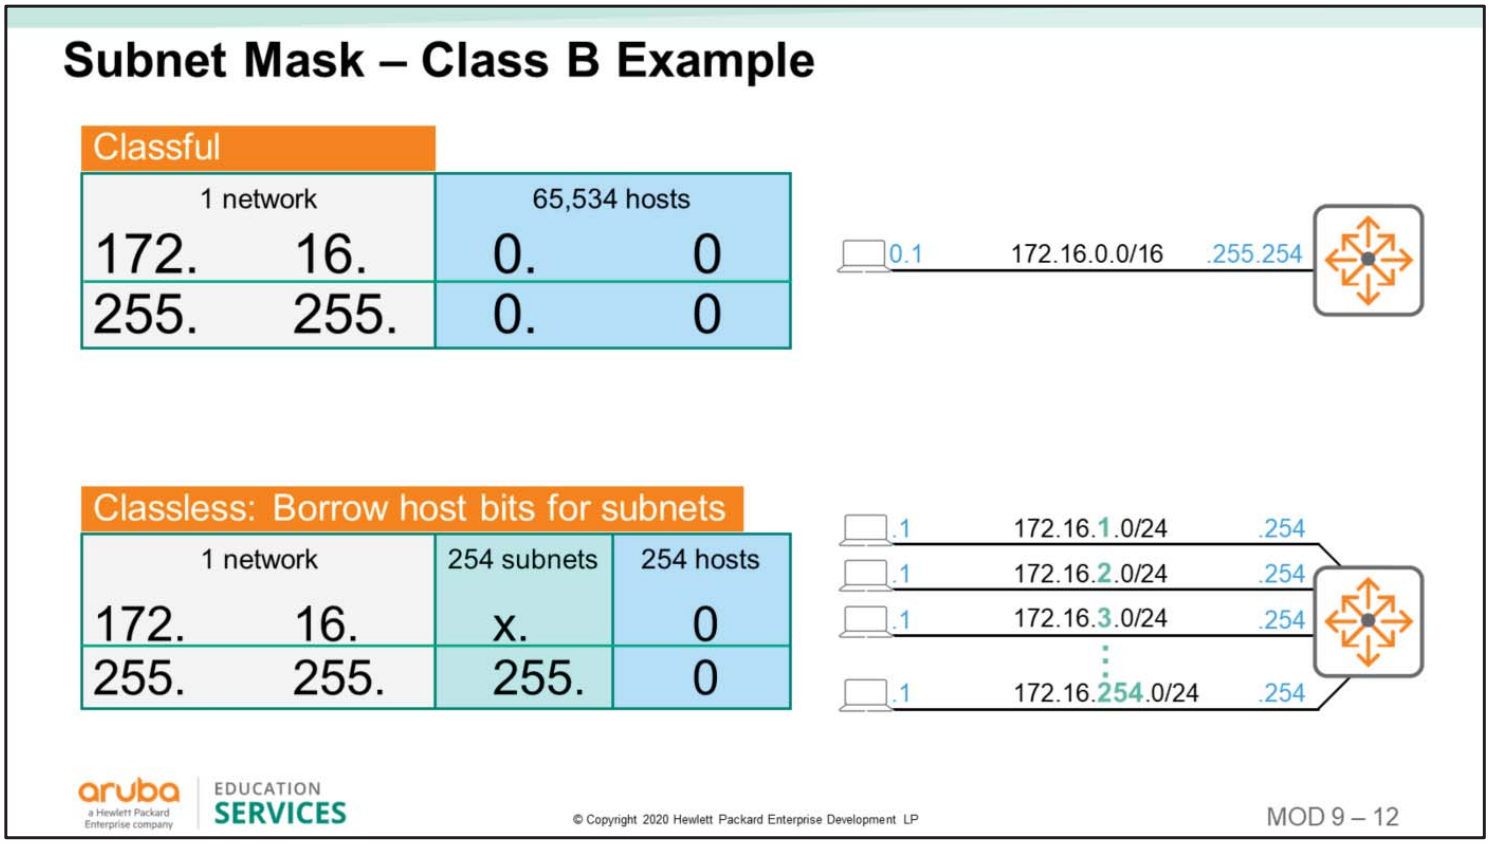

Look at the top example in the figure, which uses the Class B default mask. This means that you can have 1 network, and on that network you can have 65,534 hosts. The problem is that you actually have 150 networks, and you never have more than 200 hosts on each network. You might think that you can use 172.17.0.0, 172.18.0.0, up to 172.250.0.0. However, you cannot do this because your department has been assigned the 172.16.0.0 address space, while your co-workers have been assigned 172.17.0.0, and 172.18.0.0.

You must therefore customize the network mask. You borrow host bits to create subnetworks, using a subnet mask, as shown in the bottom example.

Now you have Class B network 172.16.0.0, split into 254 subnetworks – 172.16.1.0/24, 172.16.2.0/24, and so on – up to 172.16.254.0/24 . You still have 8 bits left over to specify hosts on each subnet. So each subnet can have 254 hosts, as shown in the figure.

Note. In this document, the term “Network Mask” specifically refers to a Classful Network. The term “Subnet Mask” refers to a Classless Network (which is a Network using the non-default masks).

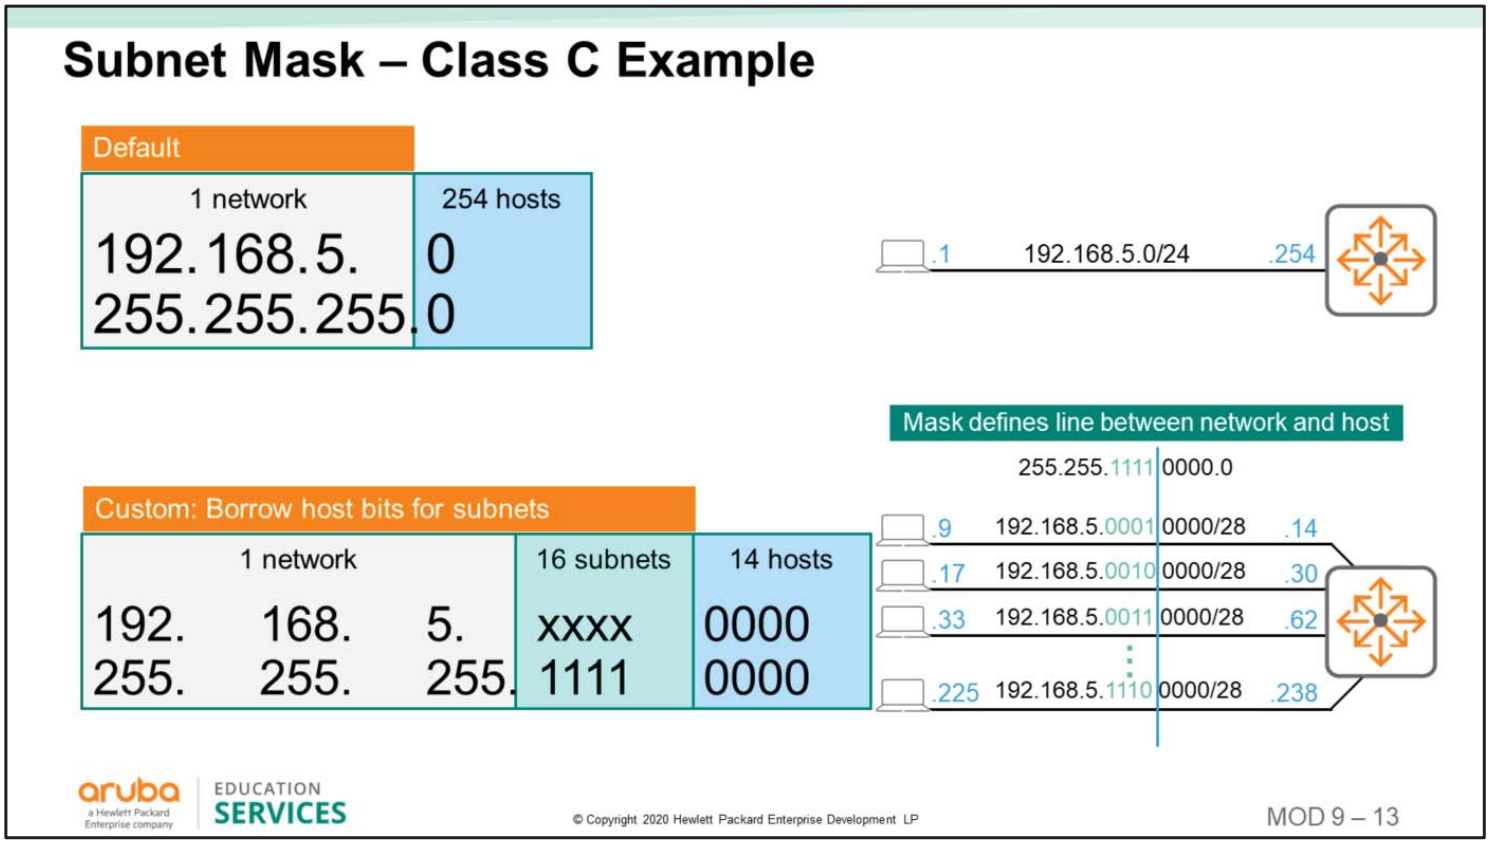

The figure shows an identical example to the Class B example, only using a Class C address. To subnet this you must think in binary, at least for the octet that is “split in two” by the subnet mask. The 3rd octet is split in two by the mask in this case, so only that octet is shown in binary – the others are shown in decimal.

The top example uses the default /24 mask for a Class C address – 255.255 .255.0. This means that you can have a single subnet with 254 hosts. However, you have 10 subnets, and each subnet has 12 hosts. You can do the same exact thing you did in the previous example – take host bits to use for subnetting. Thus, you can look at the last octet in it binary format.
You’ll stake the first 4 bits of the host portion to use for subnetting, leaving the last 4 bits available for hosts. 2^4 = 16, and so 16 unique numbers can be created with 4 bits, minus the reserved 0000 for the network number, and the reserved 1111 for directed broadcast means you can have 14 hosts.

You also have 16 network numbers. Most modern implementations allow you to use all 16 of them. In some rare cases, some legacy or specialty equipment may have issues with the “all zeros’ subnet” or the “all one’s subnet”, but they are usually OK to use.

Look at top subnet in the example – 192.168.5.0001 0000/28. that forth octet – 0001000 = 8 in binary. So, the reserved subnetwork number is 192.168.5.8/28. The reserved broadcast is 192.168.5.15. Therefore, assignable host addresses on this subnet are .9 - .14.

The remaining subnets are similarly addressed. Let’s look at this in a bit more detail.

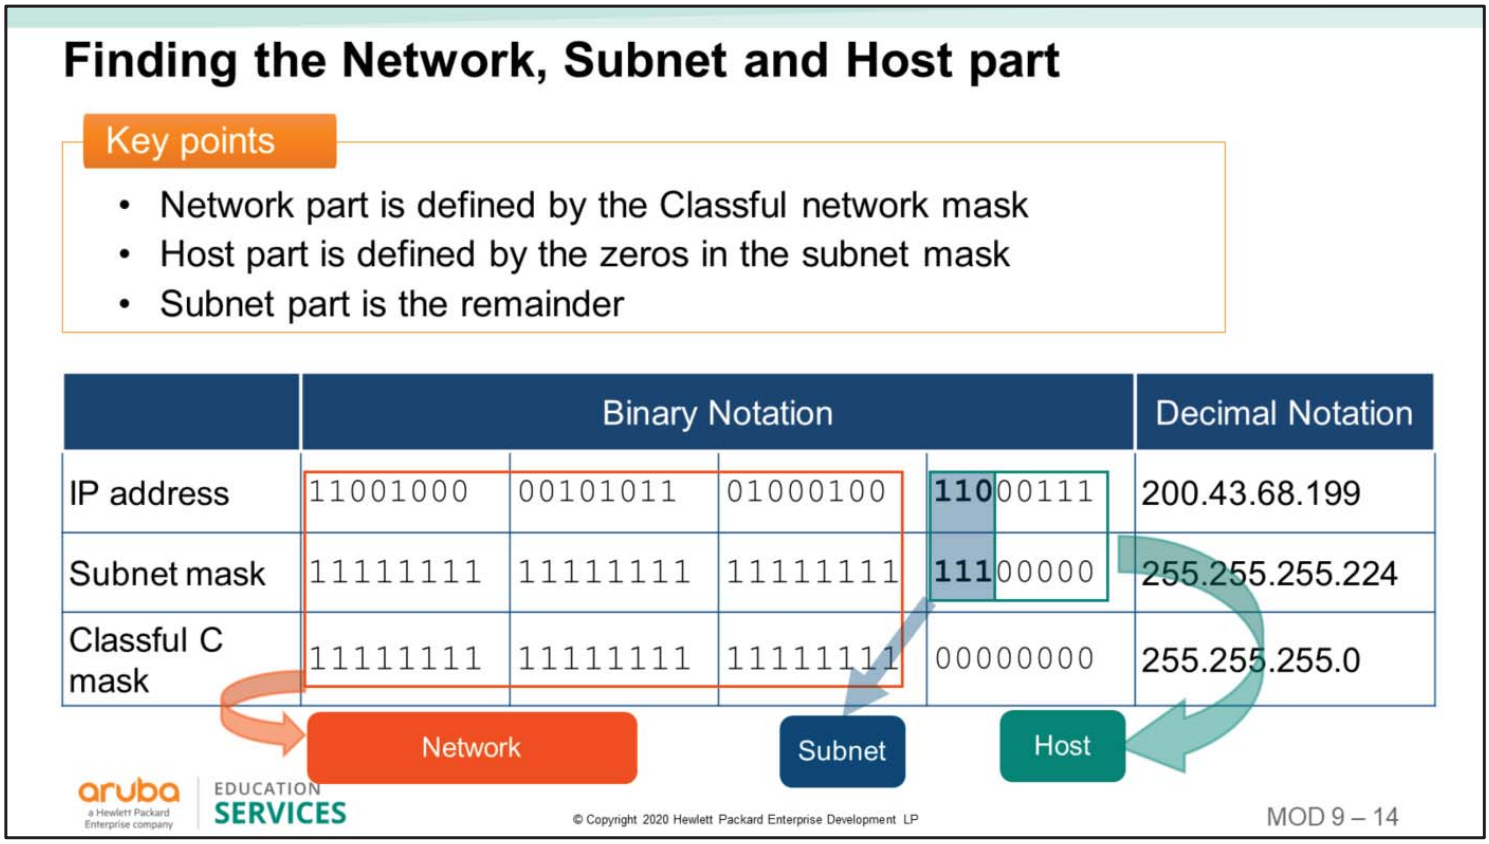

Here are some key points to remember:

- The Network part is always defined by the Classful Network mask (not the subnet mask).
- The Host part is always defined by the zeros in the subnet mask.

As an example consider the following IP and Subnet mask: 200.43.68.100/25. To understand the process this explanation will use the Binary notation.

- The first step is to convert the IP address and the subnet mask into its binary form.
- The network side is always obtained from the Classful Mask. In this case since the IP address is a Class C network, the first 24 bit will represent the Network side.
- To obtain the Host side this is simply the bits in the subnet mask that are set to zero. In this example those bits are the last seven.
- Finally the Subnet side is what is not Network or Host, this is simply the bit 25.

The image shows the previous process.

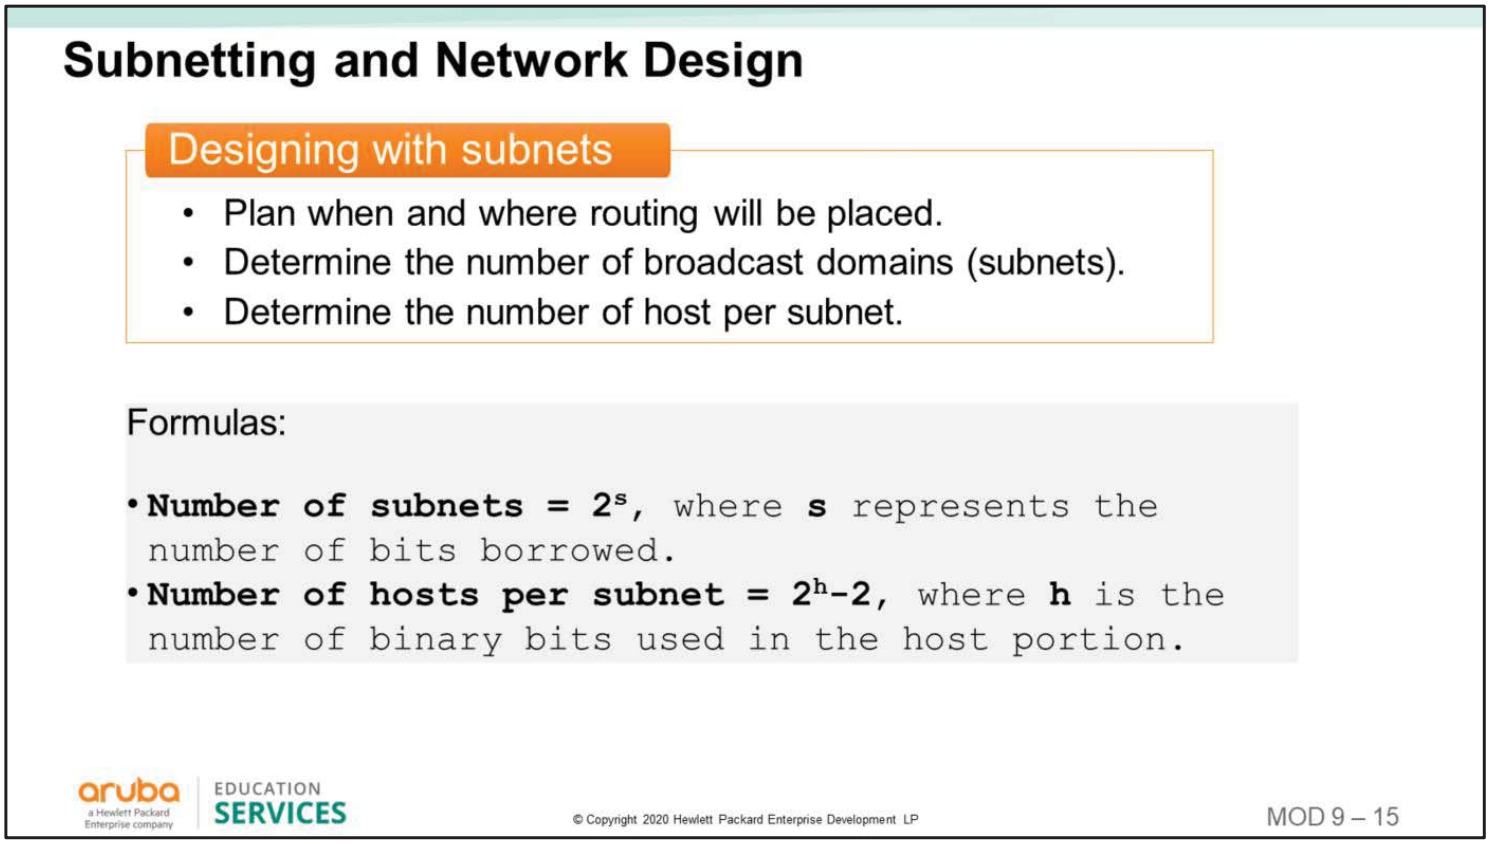

Subnetting has an important implication for network administrators because it helps to design broadcast domains. This dictates where and when routing must be done.

Subnetting is typically based on the number of Subnets needed and the number of Hosts per subnet that are required. To determine this information, two formulas can be used:

- Number of subnets = 2s, where s represents the number of bits borrowed.
- Number of hosts per subnet = 2h-2, where h is the number of binary bits used in the host portion.

Note. Remember that you need to subtract the 2 addresses because those are reserved to the Network ID and to the Local Broadcast address.

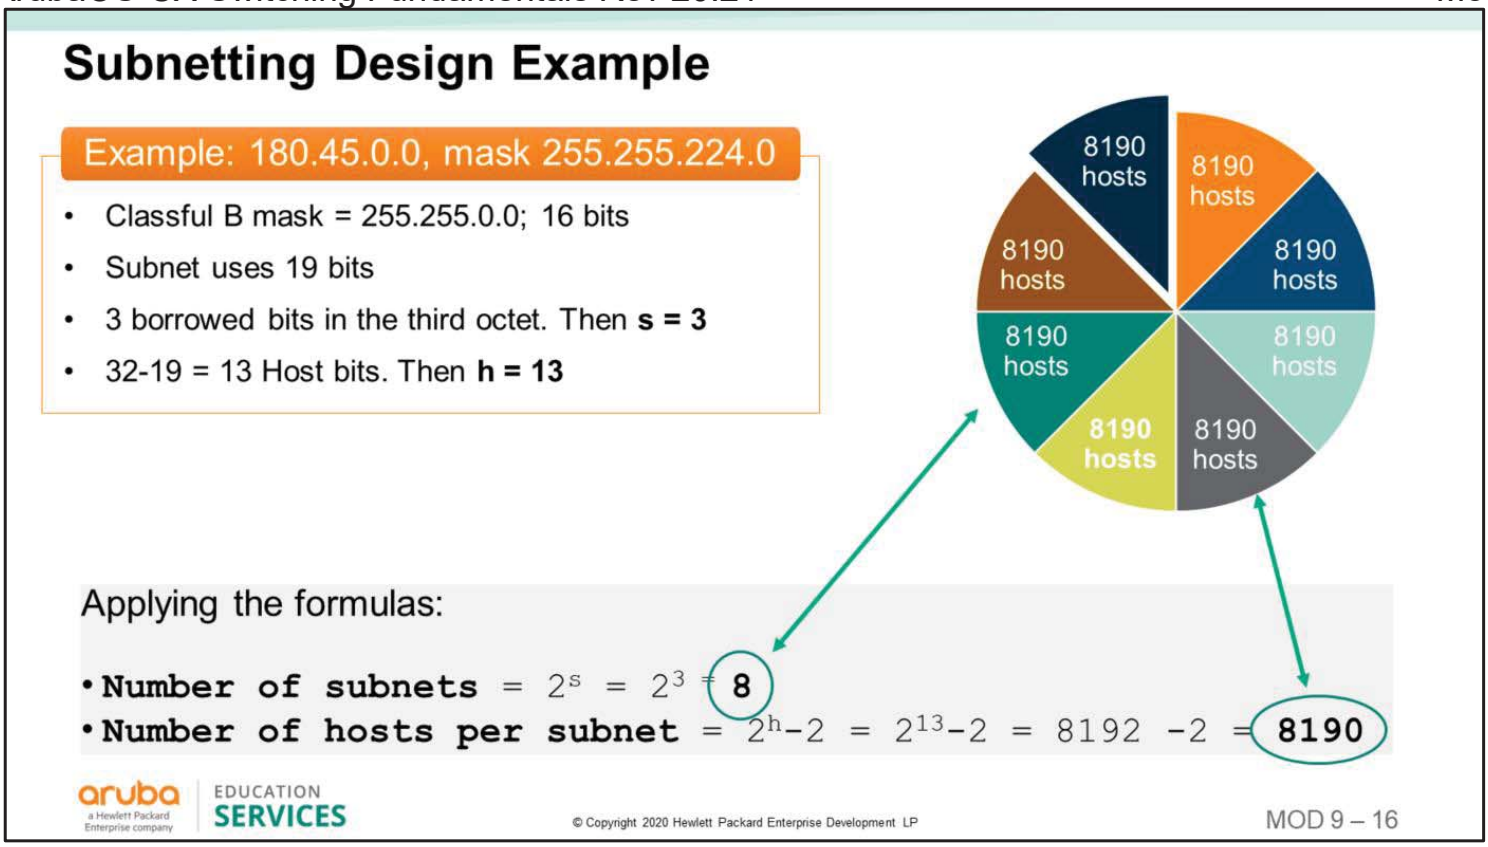

As an example, consider the Class B network – 180.45.0.0 with a subnet mask of 255.255.224.0 From this information you can determine the following:

- The classful mask for a class B is 16 bits.
- The subnet mask uses 19 bits.
- There are three borrow bits in the third octet. Then **s = 3**
- Host part is determined by a simple subtraction. Host = 32 – bits set to 1. In this case 13. Therefore **h = 13**.

You can now apply the two formulas:

- Number of subnets = 2^3 = **8**

The number of hosts that these subnets can use depends of the reminder zeros in the subnet mask, in this case you have 21 zeros.

- Number of hosts per subnet = 2^{13}-2 = 8192 – 2 = **8190**

The figure shows an example of the actual addressing that you might assign. The third octet is shown in binary, and the 1st, 2nd, and 4th octets are shown in decimal. This saves you time, since you only need to work with binary in the 3rd octet.
The top example is 180.45.001 00000.0/19. 001 00000 = 32 in decimal, and so the subnet is 180.45.32.0/19. This means the first assignable IP address is 180.45.32.1, and the last assignable IP address on that subnet is 180.45.63.254. the following slides will break this process down into distinct steps for you.

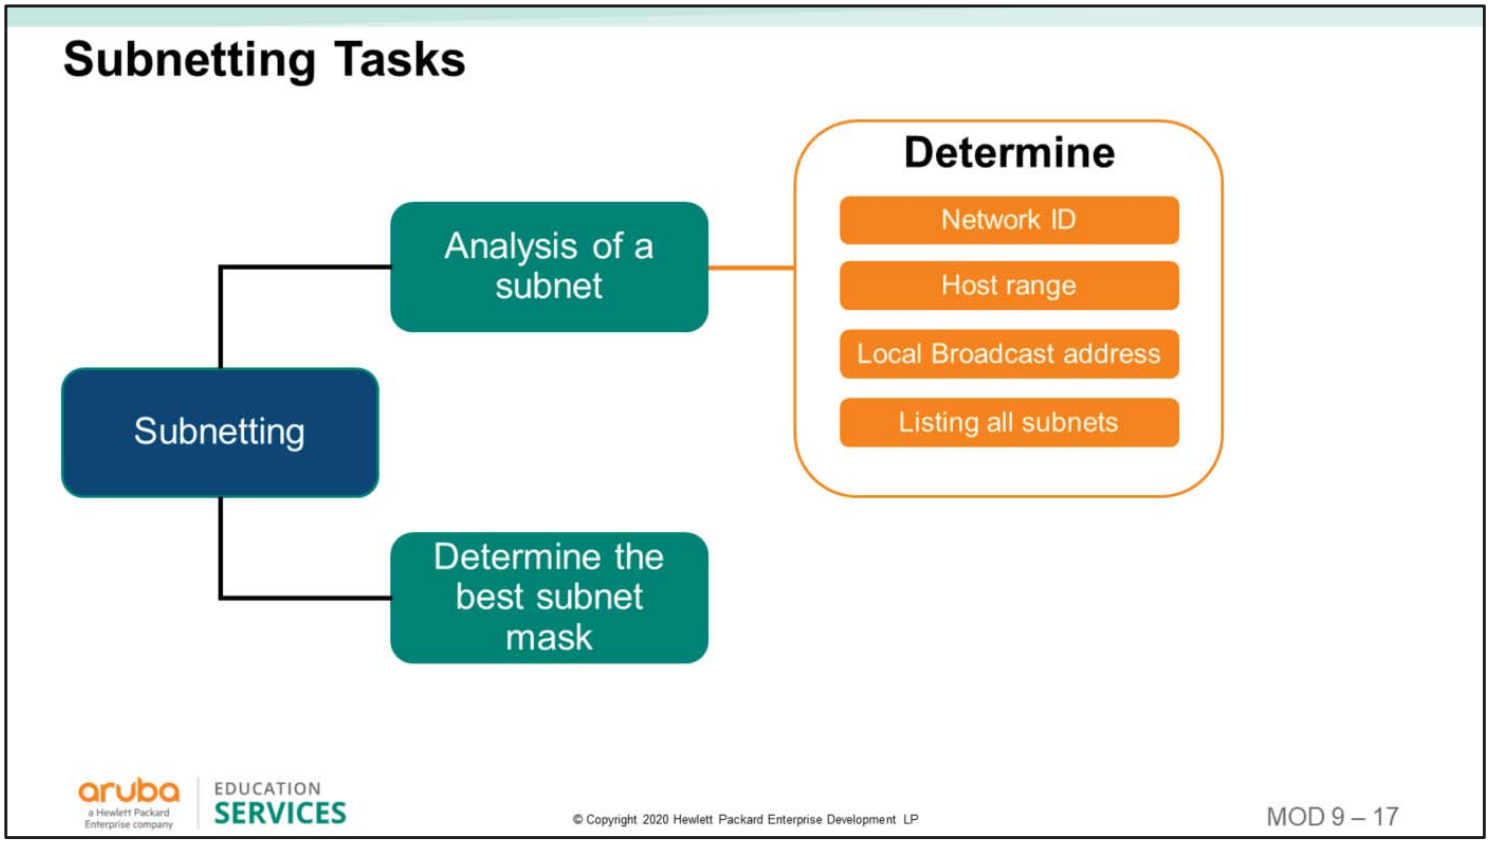

The analysis of the subnet requires that you determine the following:

- The Network address or Network ID (obtain to what network an IP belongs)
- The Local Broadcast address.
- Assignable range of addresses to hosts.
- Listing the subnets (optional)
-

Note. Decimal and Binary approaches can be used to determine the previous parameter, however in real life the Binary approach is not commonly used since it takes more time. This training then, considers the analysis using only the decimal notation.

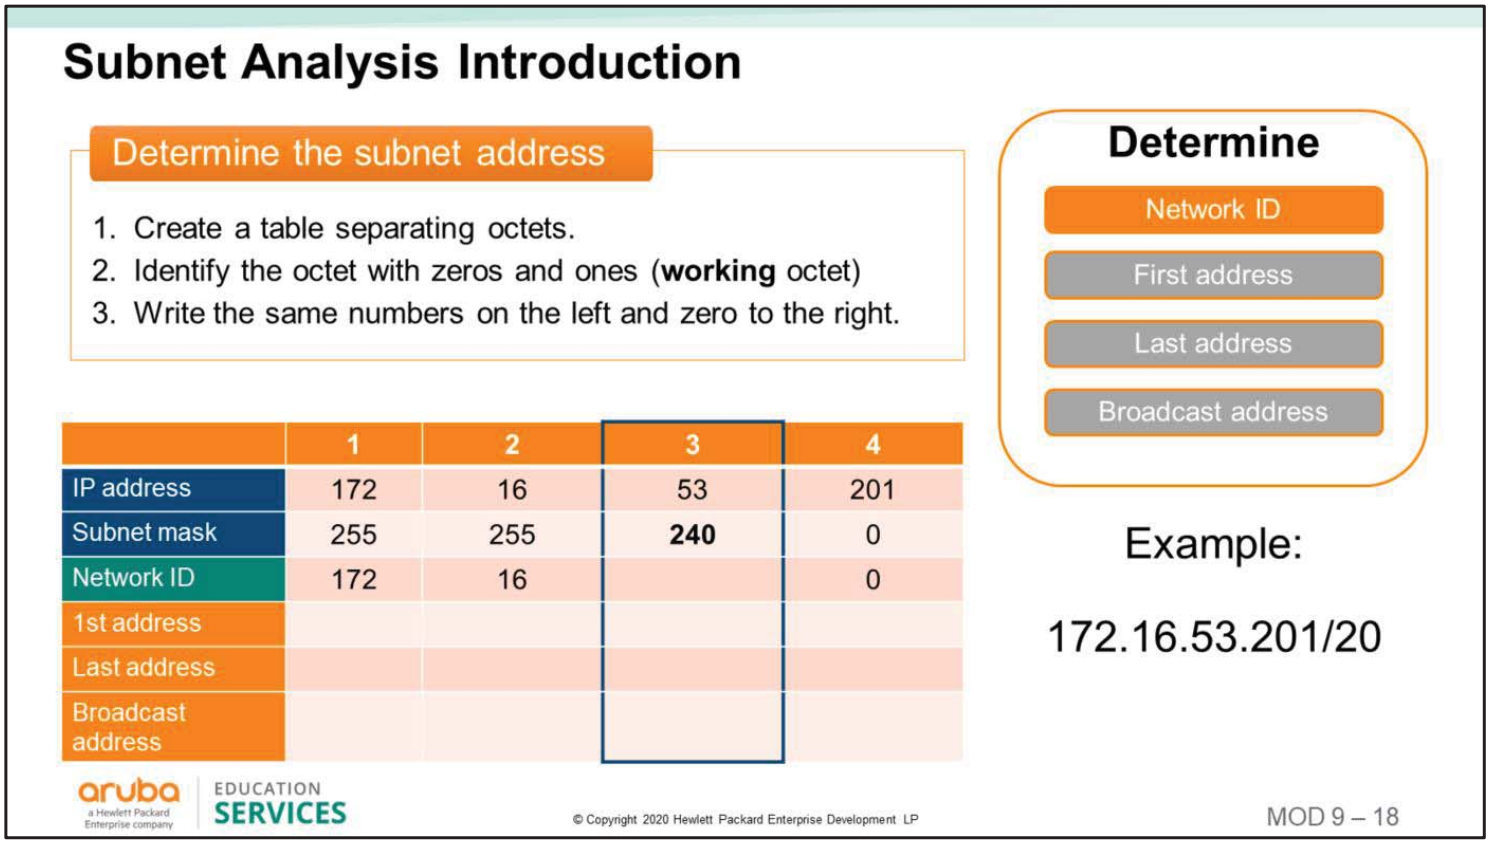

Determining the Subnet address

As an example, consider the following IP and mask: 172.16.53.201/20

1. First create a simple table where you can separate the different octets and write down the IP address and the network mask.
2. Identify the octet with zeros and ones. This octet will require all our attention. In this case a quick analysis of the subnet mask (the octet that is not 255 or 0) indicates that the third octet is the one that meets this criterion. This will be our working octet.
3. Write down the same numbers of the IP address on the left octets of the working octet. And write a zero on the right octets of the working octet.

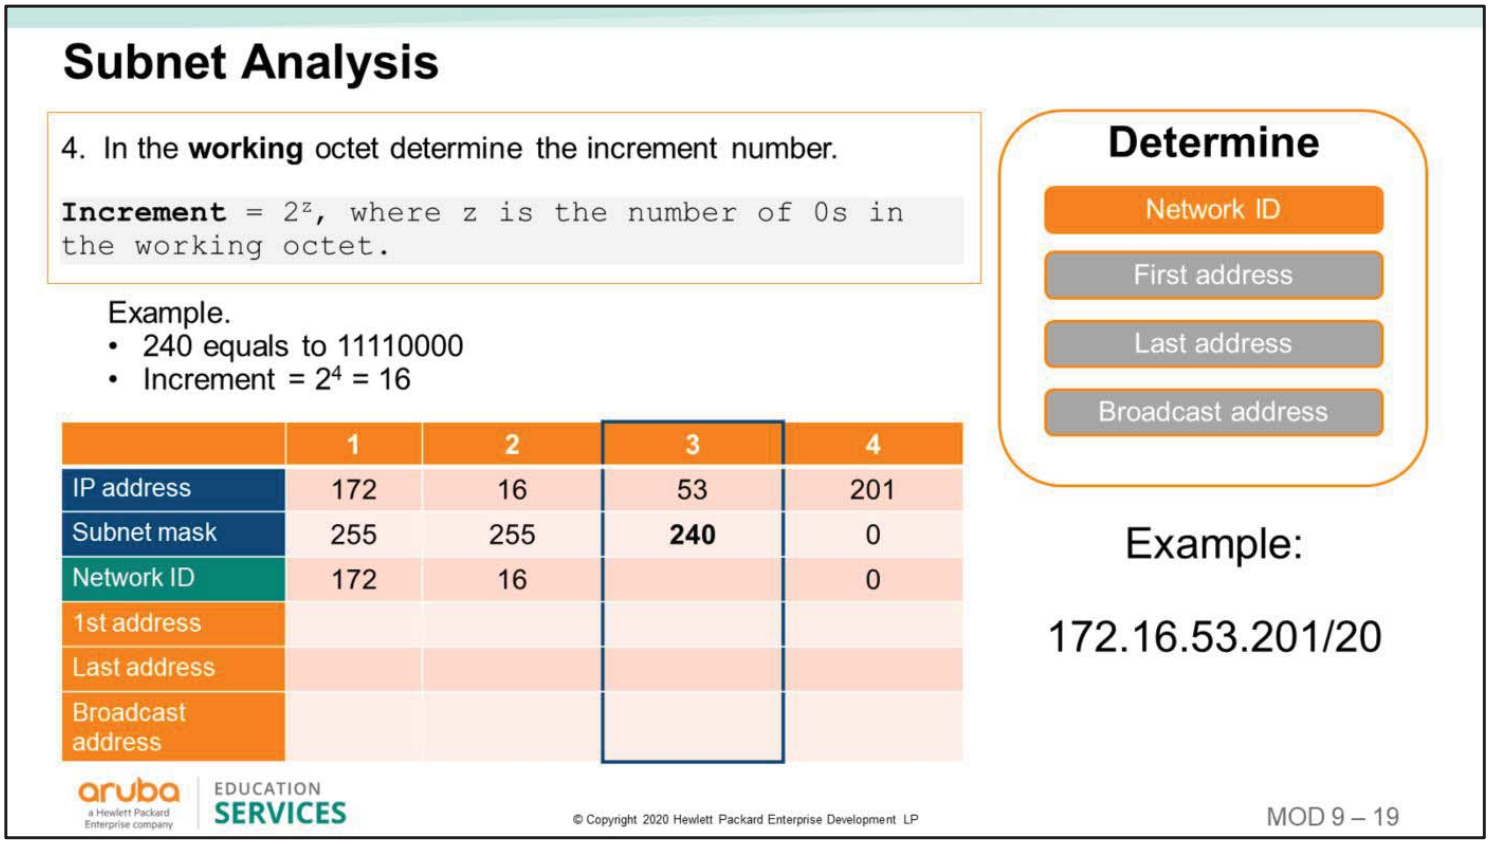

4. In the working octet, determine the increment number. This number is determined with the formula:

Increment = 2z, where z equals the number of zeros in the working octet.

In this case the third octet has four zeros on it which implies that the increment is 16 (24)

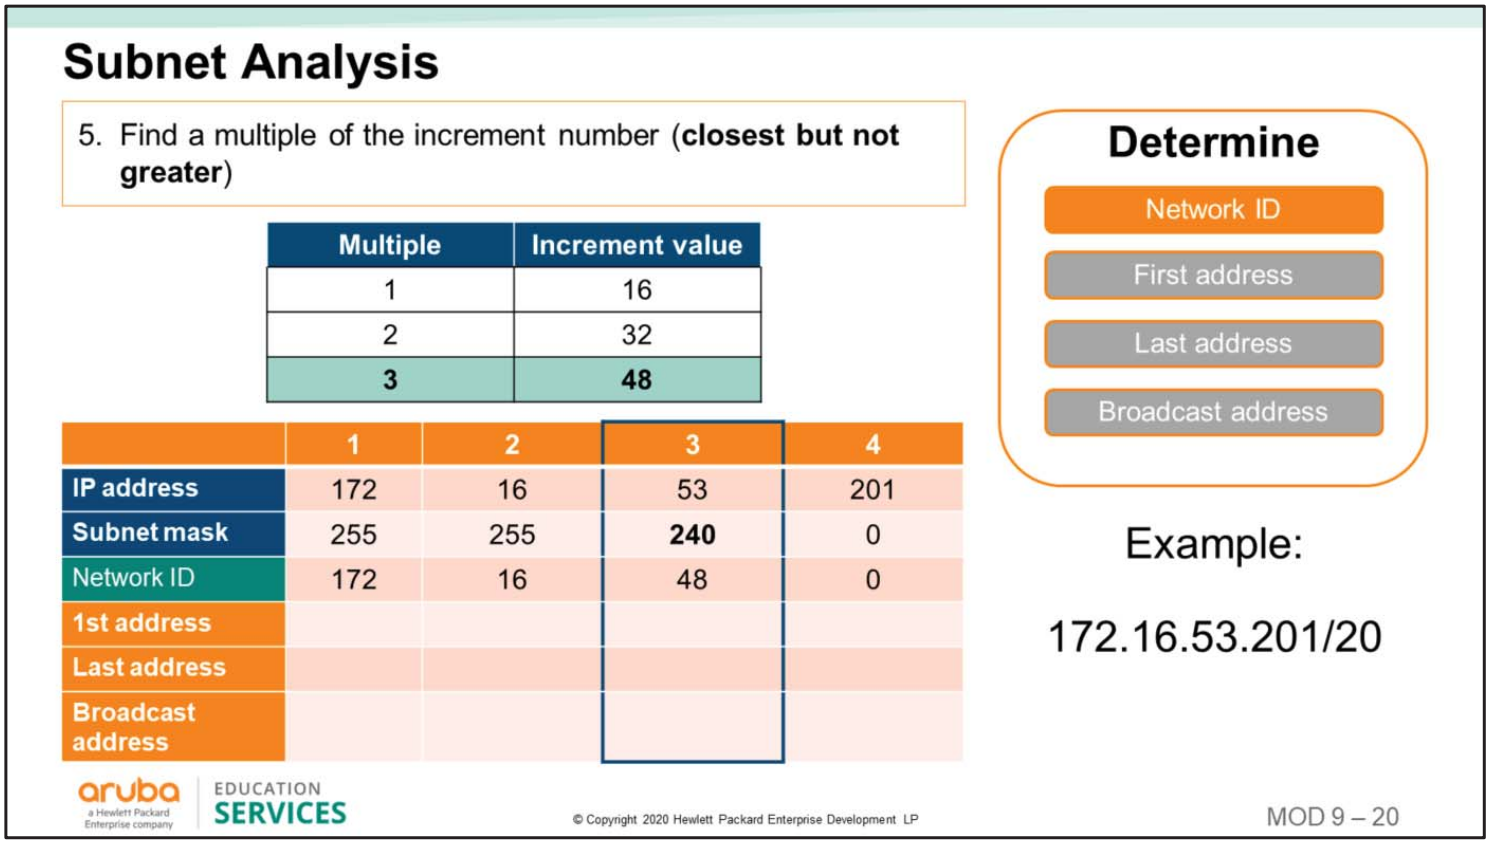

5. Find a multiple of the increment number that is closest to but not greater than the IP address's working octet value. And write down the value in the working octet.

In this example the multiple that is closest to the IP is 48.

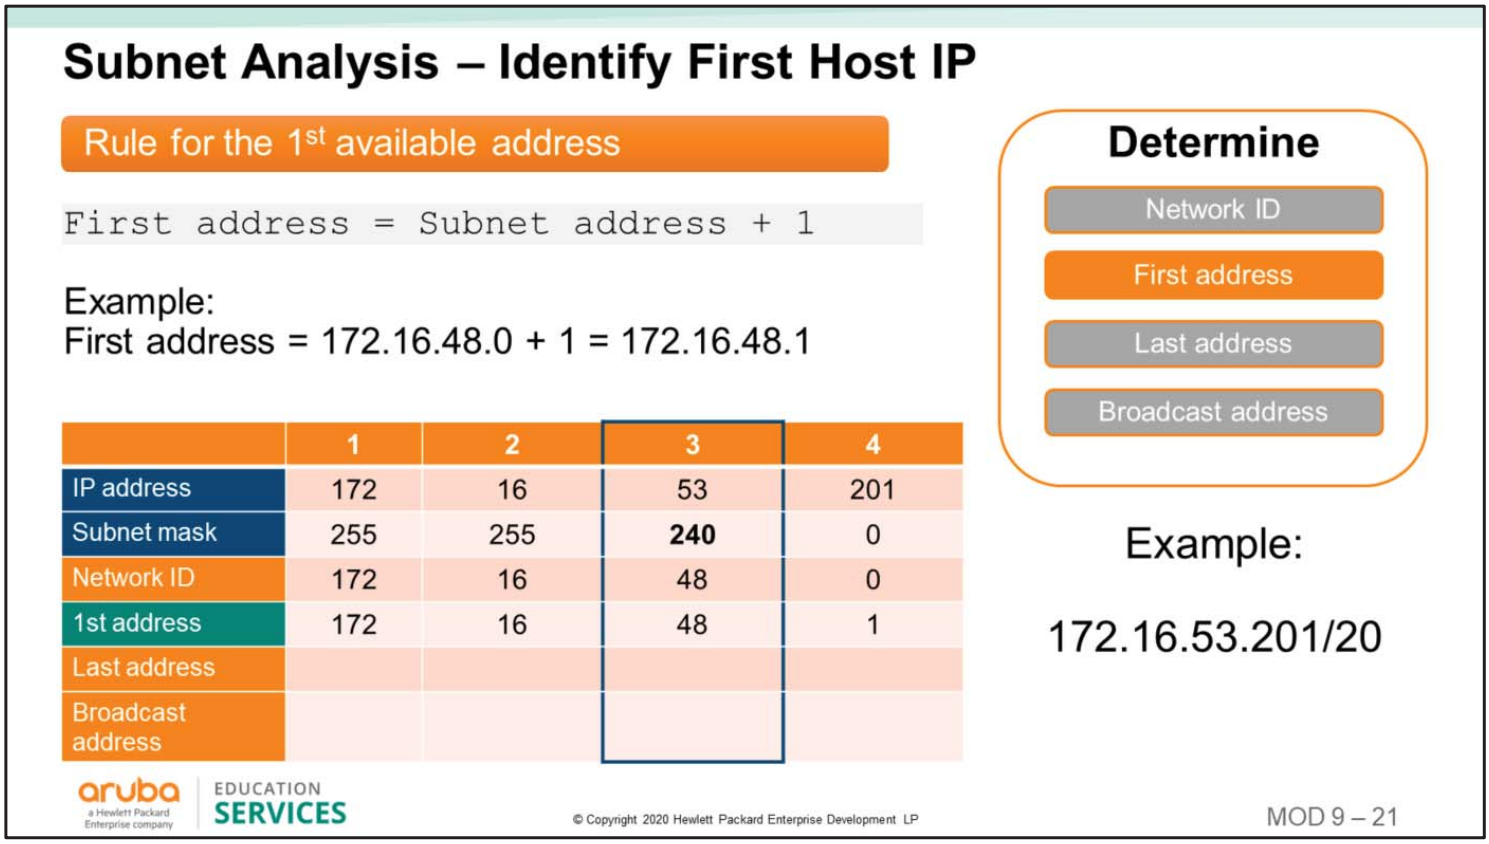

Determining the First address in the subnet

Use the following rule for the first address.

First address = Subnet address + 1

In the example: First address = 172.16.48.0 + 1 = 172.16.48.1

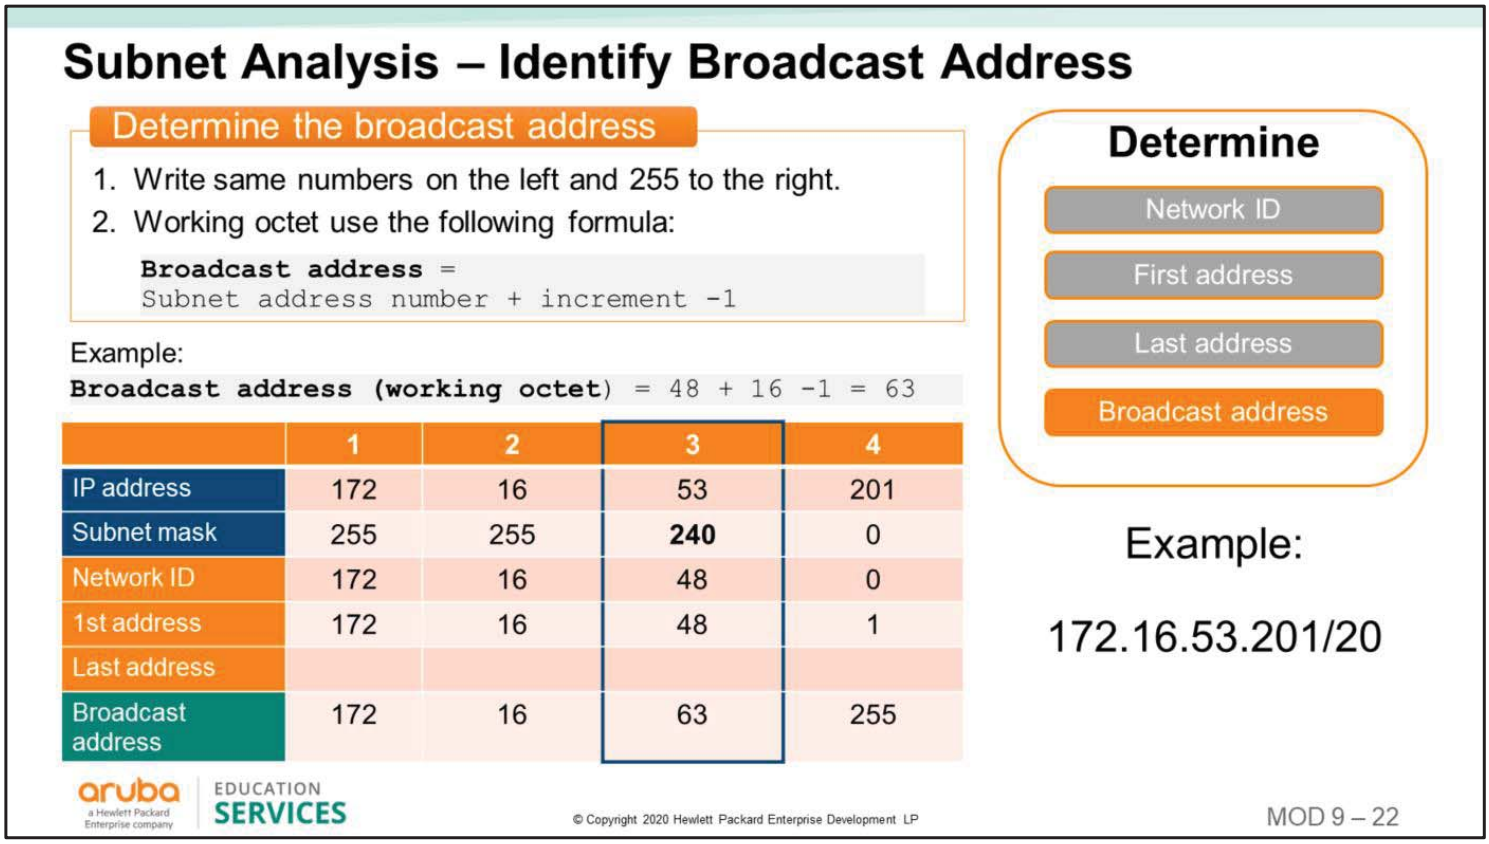

Determining the Broadcast address

1. Write down the same numbers of the IP address on the left octets of the working octet and write a 255 on the right octets of the working octet.
2. For the working octet follow the next formula:

Broadcast address (working octet) = Number of the Subnet address in the working octet + increment - 1

In this example:

Working octet in broadcast address = 48 + 16 -1 = 63

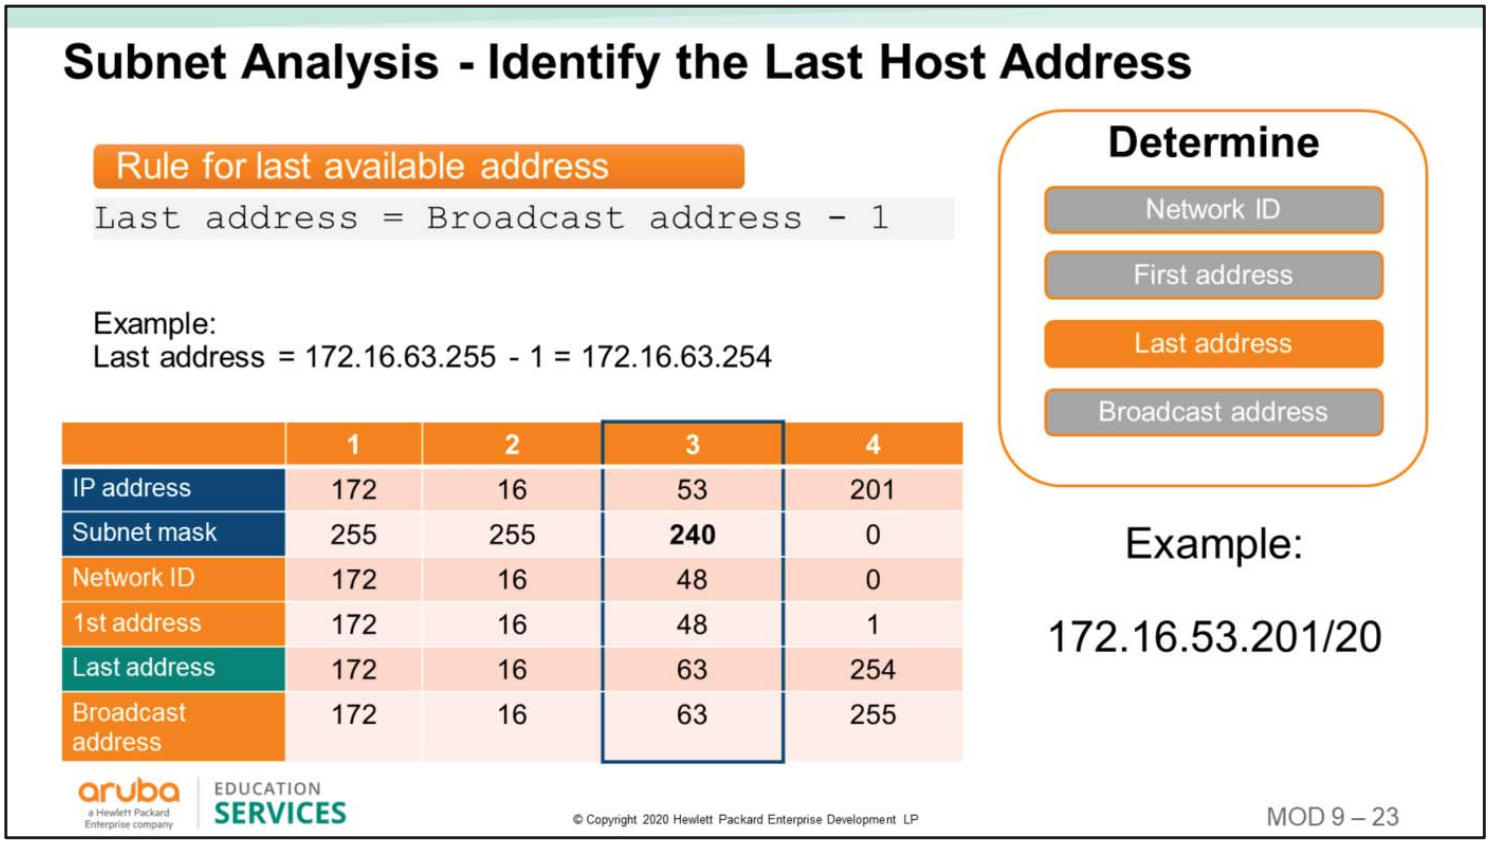

	1	2	3	4
IP address	172	16	53	201
Subnet mask	255	255	**240**	0
Network ID	172	16	48	0
1st address	172	16	48	1
Last address	172	16	63	254
Broadcast address	172	16	63	255

Determining the Last address in the subnet

1. Use the following formula:

Last address = Broadcast address – 1

In the example: Last address = 172.16.63.255 - 1 = 172.16.63.254

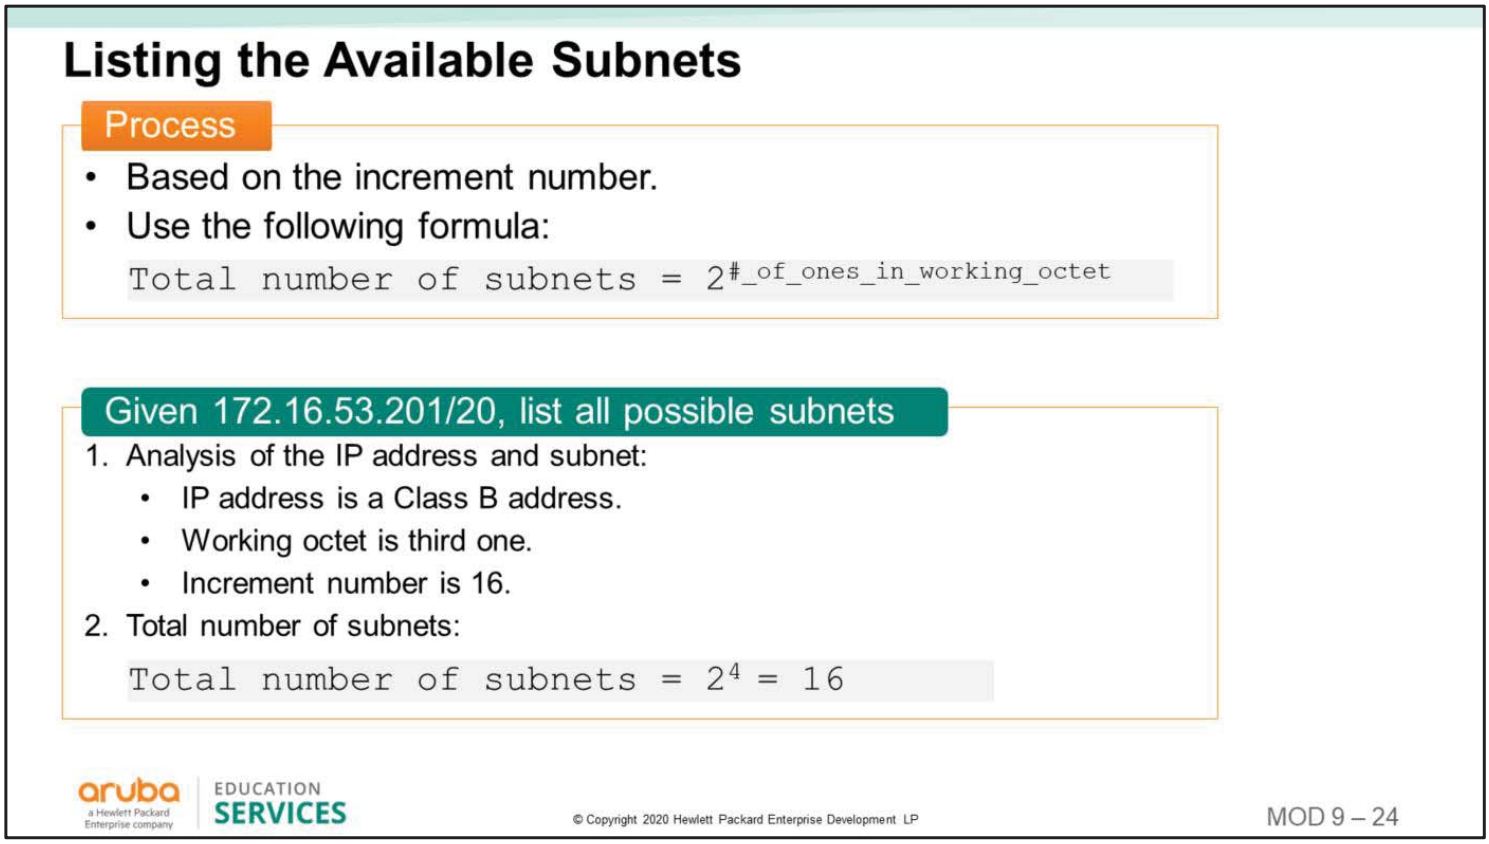

Listing all subnets that are available is based on the increment number. To know how many networks will be generated you can use the following formula:

Total number of subnets = 2#_of_ones

The process is explained in the following example:

From the following IP address and mask list all possible subnets 172.16.53.201/20

1. From the previous section you know the following information:
   - The IP address is a Class B address.
   - The Increment number is 16.
   - The Working octet is the third one.
2. Total number of subnets is 24 = 16

Listing the Available Subnets

3. Write classful network as the first subnet.
4. Add the **increment** number to the working octet.
5. Repeat the previous step to list all networks.

	1	2	3	4
IP address	172	16	53	201
Subnet mask	255	255	240	0
Classful network	172	16	0	0
Next subnet	172	16	**16**	0

Subnet	Network address
1	172.16.0.0
2	172.16.16.0
3	172.16.32.0
…	

Subnet	Network address
…	
14	172.16.208.0
15	172.16.224.0
16	172.16.240.0

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 25

3. To define the first network simply write down the Classful network.
4. To obtain the next Network address, add the increment number to the working octet of the first subnet.
5. Repeat the previous step to obtain the rest of the addresses.

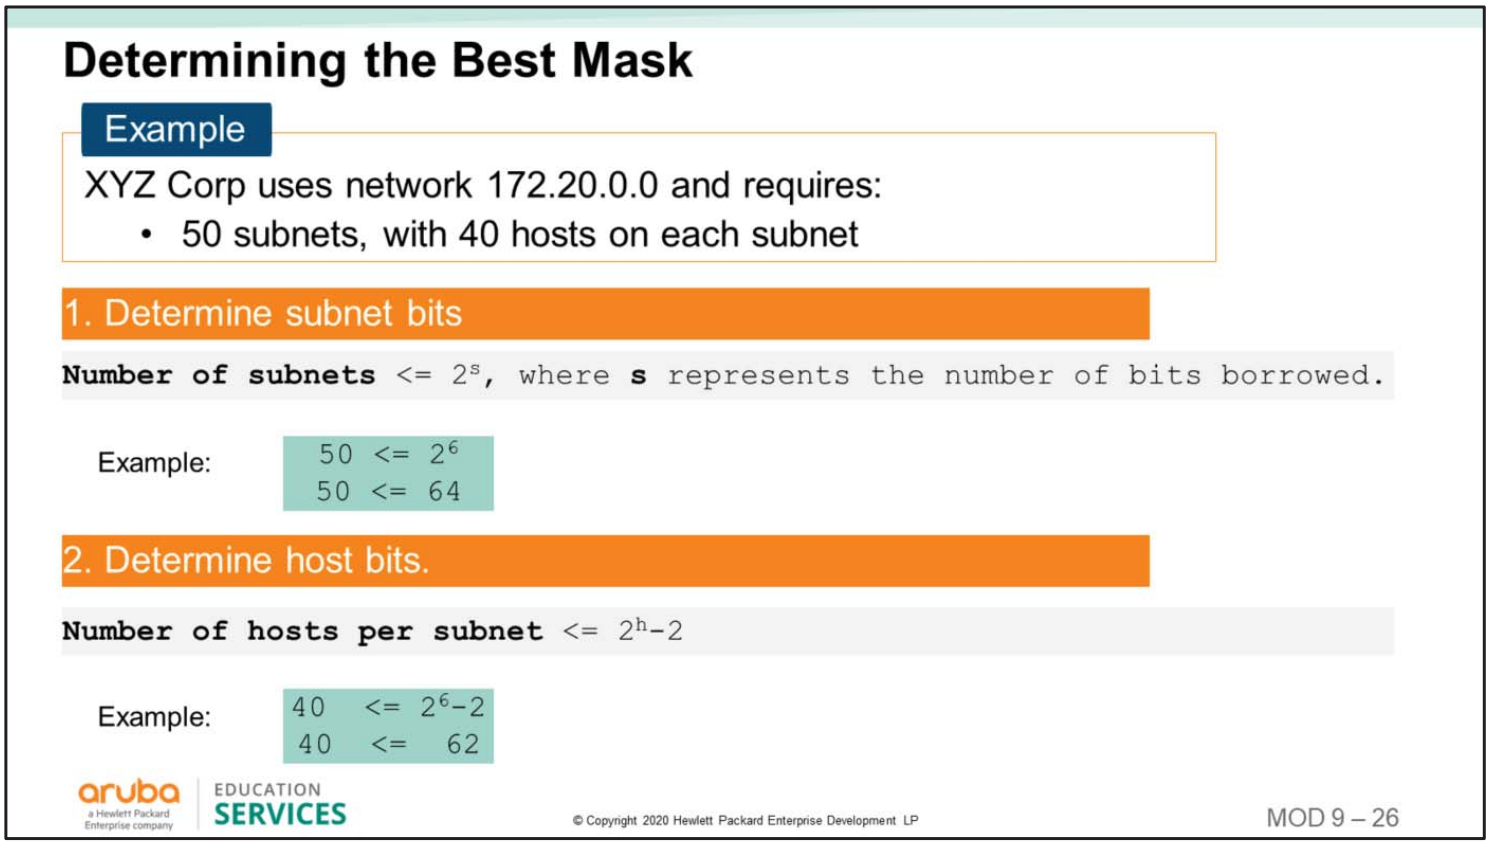

Now you'll learn how to determine the best mask that suits your organizational needs.

Suppose that an organization uses Class B network **172.20.0.0,** which must be divided into 50 subnets. Each subnet must accommodate at least 40 hosts. This number considers the future growth for the next 5 years. Here is the process:

1. Determine how many subnet bits permit 50 subnets. You can use the following formula:
Number of subnets <= 2s

In this case you are using the formula in the inverse way: 50 <= 2s. The best way to solve this equation is to replace s with different values and find which value is equal or greater than 50. In this case the result is 6, because 26 is 64.

In this case the number of bits borrowed from the host portion of the address is 6.

2. Determine how many host bits will be required to allocate the need. Use the formula:
Number of hosts per subnet <= 2h-2

In this case you are using the formula in the inverse way: 40 <= 2h-2. The best way to solve this equation is to replace **h** with different values. In this case the result is 6, because 26 is 64.

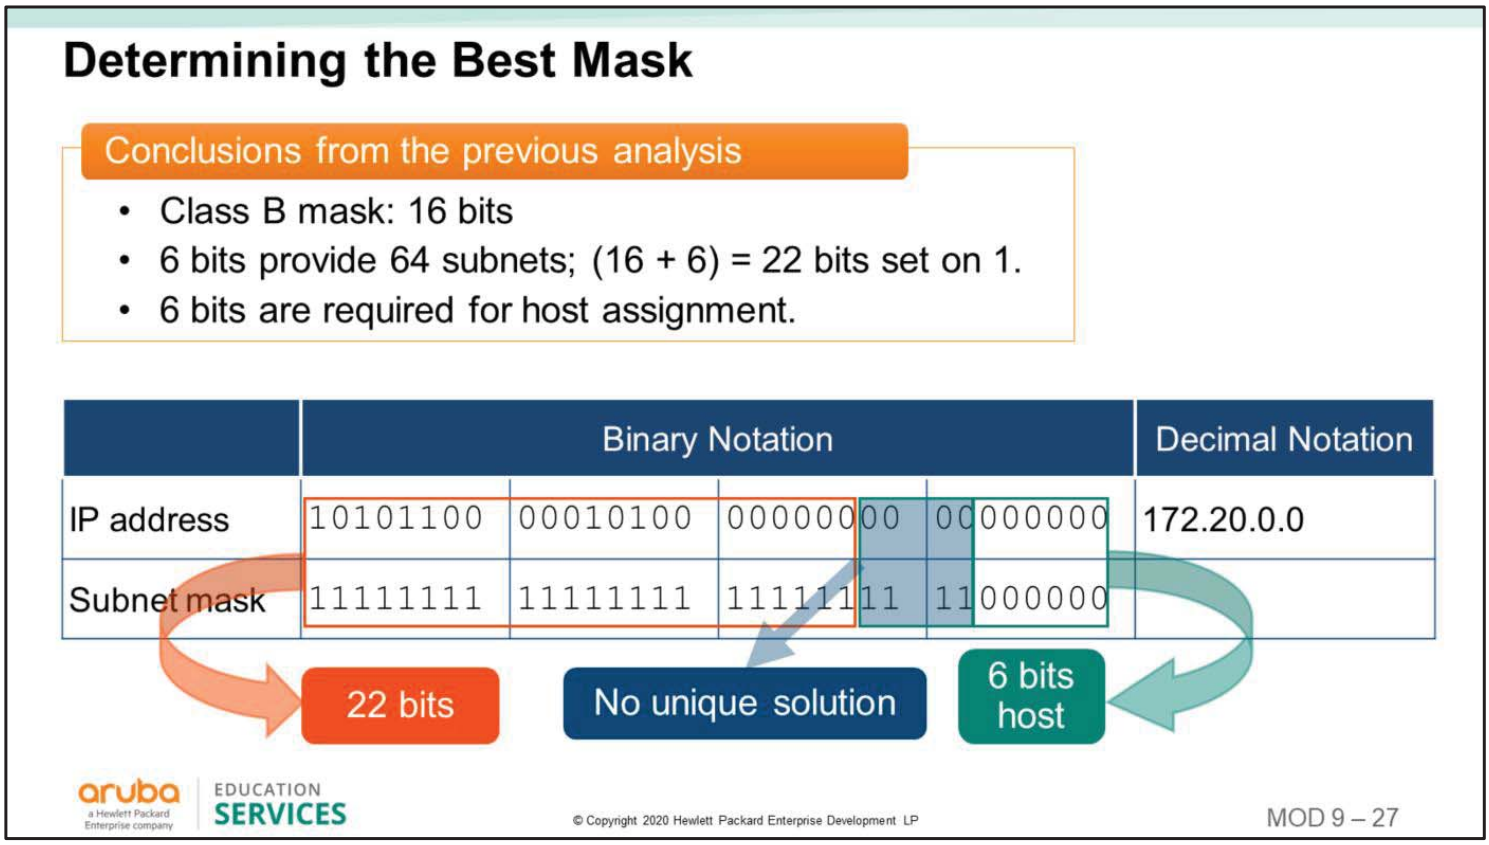

From the previous results you can conclude the following:

- Since is a Class B address the network mask is 16 bits set to 1.
- You will need 6 borrow bits.
- You will need 6 host bits.

Now, you may wonder, what about the rest of the bits? How can you use them? This example simply does not have a unique answer that meet the original requirement. Continue on to see how this plays out.

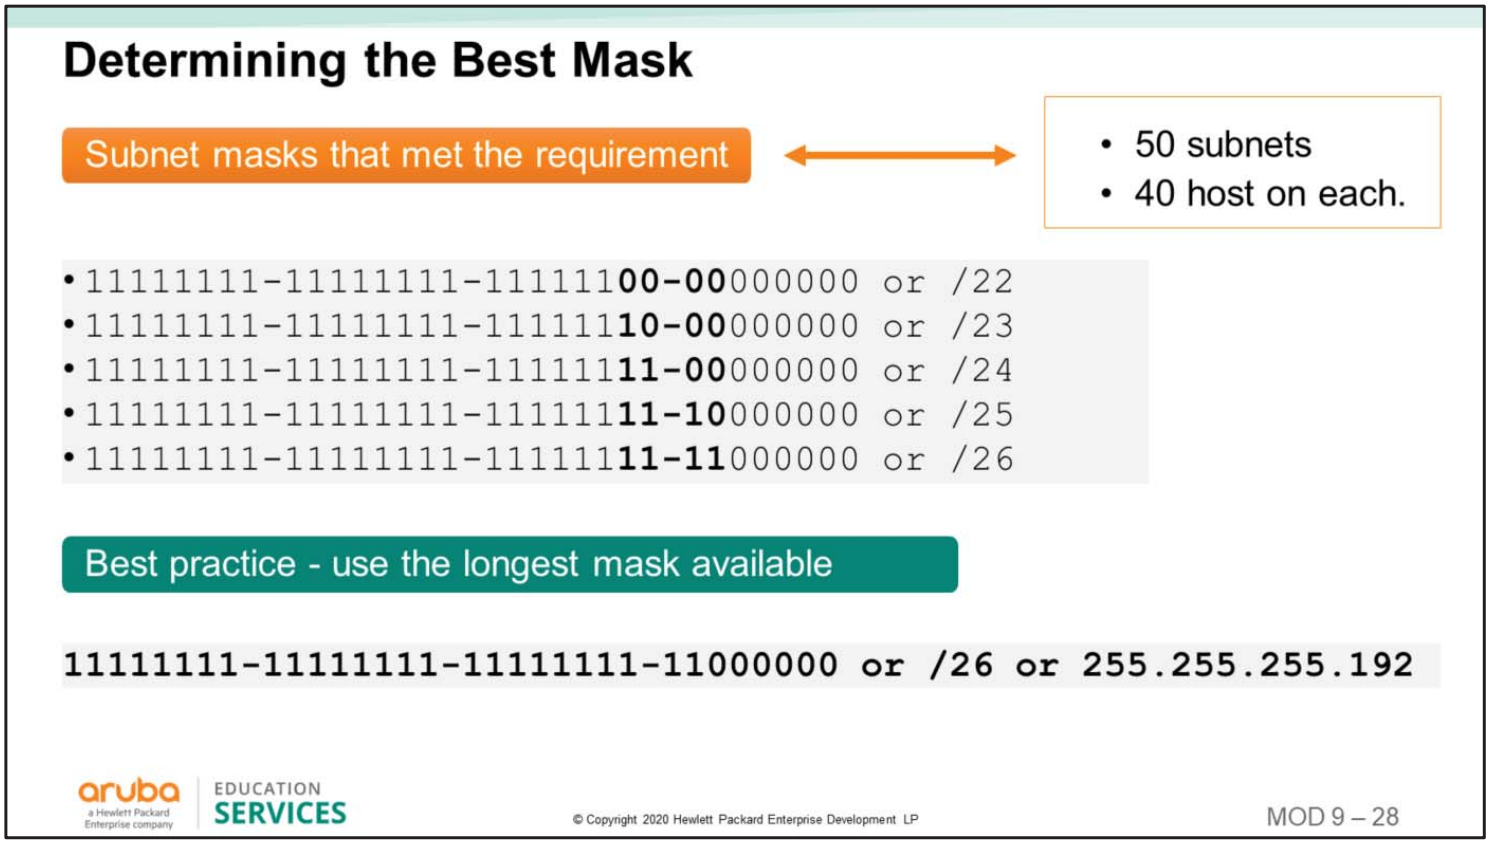

- 11111111-11111111-11111100-00000000 or /22
- 11111111-11111111-11111110-00000000 or /23
- 11111111-11111111-11111111-00000000 or /24
- 11111111-11111111-11111111-10000000 or /25
- 11111111-11111111-11111111-11000000 or /26

In case you run into this situation the best practice is to use the subnet mask that uses the longest number of bits set to one. In this case the subnet mask /26 would be the recommended answer.

Answer: **11111111-11111111-11111111-1100000**0 or **/26** or **255.255.255.192**

Question #1

IPv4 Address Classes, Reserved Addresses, Private and Public IPv4 Addressing

Which statements are true about classful IP addressing?

a. Class A addresses support more networks and more hosts than Class C addresses.

b. If you must support 255 hosts on a single network, you can use a Class C address.

c. Your ISP's routers will analyze the network portion of destination address 172.16.3.7, and route that packet toward its destination.

d. Public address ranges are globally unique.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 29

Correct answer:
d.

Question #2

Subnet Mask – Class B example

You are using a /24 subnet mask with the Class B address 172.20.0.0. Which options describe a valid result of this scenario?

a. You cannot do this. Class B addresses use a /16 mask.

b. You can assign addresses to over 200 networks.

c. Each subnet can support 65,534 hosts.

d. If you must support more than 300 subnets, you must choose a different mask

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 30

Correct answers:
b.
d.

VLSM and CIDR

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

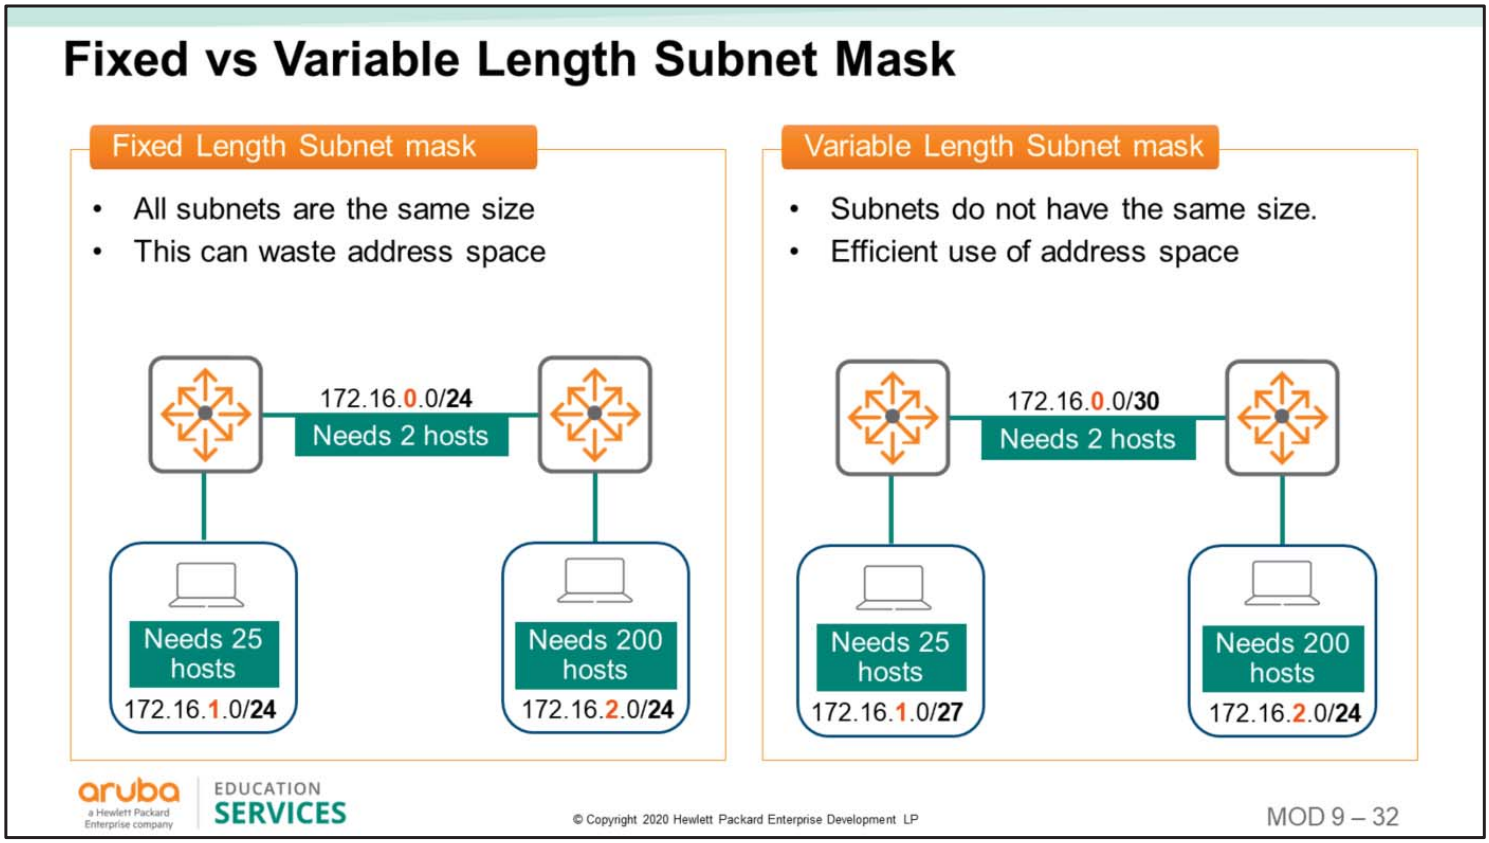

You have learned how to use subnetting to divide one network into smaller subnetworks. Some routing protocols, like RIPv1, only support fixed length subnet masks. This means if you must use a /24 mask to accommodate a subnet with 200 hosts, then you must use a /24 mask on all other subnets of the 172.16.0.0 network. This Fixed Length Subnet Masking (FLSM) is shown in the figure.

What a waste of address space. Many of the other networks only need 25 hosts, but the /24 mask means that those subnets can accommodate 254 hosts – over 200 IP addresses are wasted – never to be used. This is even worse on the point-to-point links between switches. These subnets will never need more than 2 IP addresses. One router has IP address 172.16.0.1, and the other has 172.16.0.2. The other 252 addresses are wasted. Because of RIPv1's FLSM requirement, you have used up subnets 172.16.1.0/24 – 172.16.5.0/24.

More sophisticated protocols like OSPF support Variable Length Subnet Masking (VLSM). Although one or more subnets may require a /24 mask to accommodate 200 hosts, you remain free to use different masks in the 172.16.0.0 address space.

In the figure, the bottom, right subnet must support 200 hosts, so you assign 172.16.2.0/24. Now, to accommodate the subnet with 25 hosts, you create subnets 172.16.1.0/27. Each can accommodate up to 30 hosts, and you still have subnets 172.16.1.32/27, 172.16.1.64, 172.16.1.96/27, 172.16.1.128/27, 172.16.1.160/27, 172.16.1.192/27 and 172.16.1.224/27 available for future growth.

To accommodate the point-to-point link, you create 172.16.0.0/30 again with plenty of subnets remaining for future growth. As you can see, this is far more efficient than VLSM. The only disadvantage of VLSM is that it requires you to think in binary, but you will soon

master that anyway. Let us explore this concept further and show how you can derive network designs like this on your own.

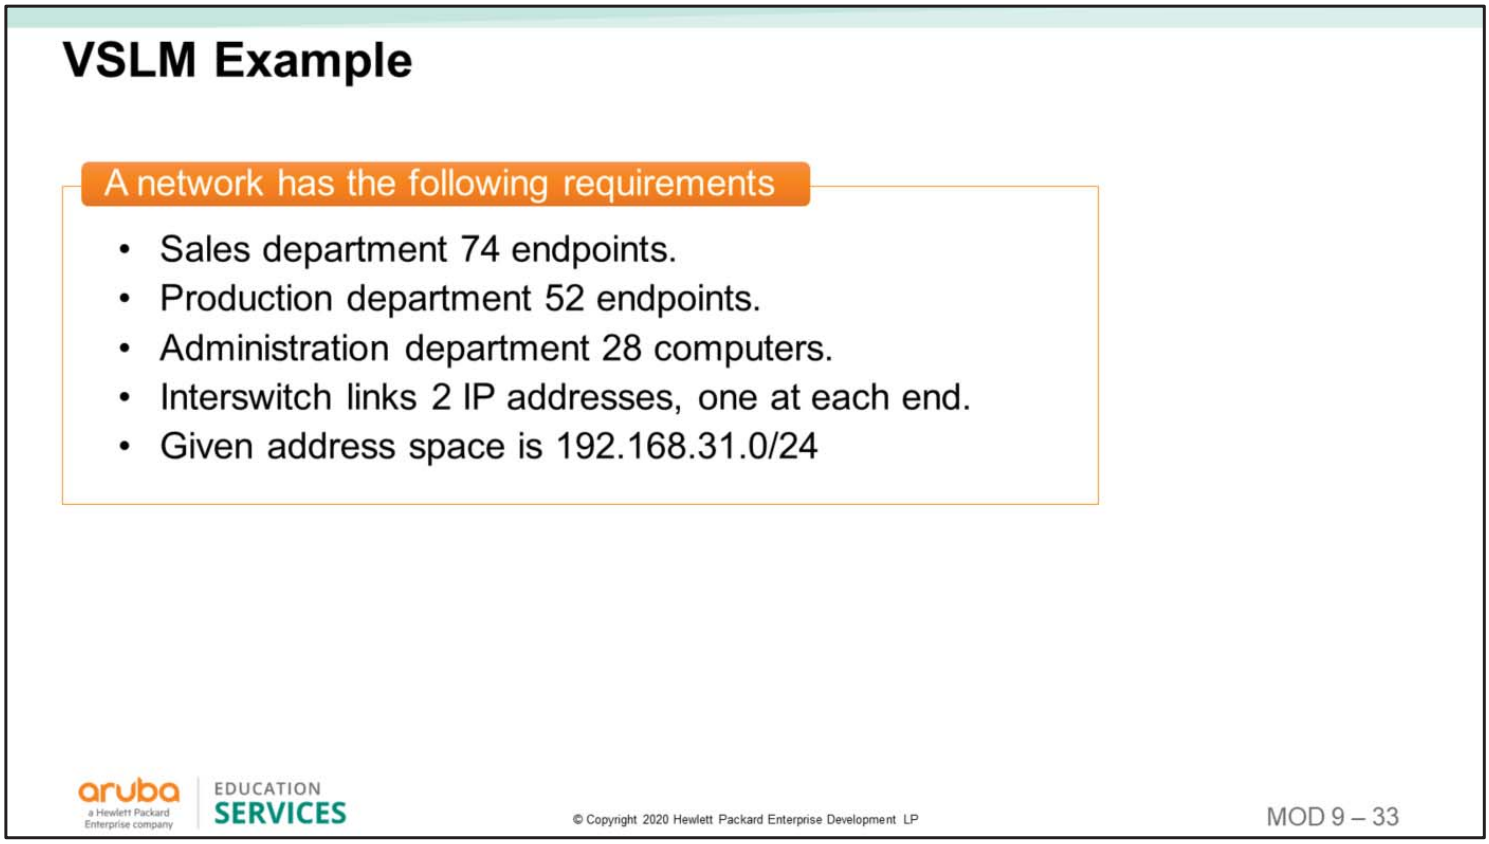

An example can help you to fully understand how you can use VLSM. In this case you have a network with the following requirements:

- Sales department has 74 endpoints.
- Production department has 52 endpoints.
- Administration department has 28 computers.
- Interswitch links require 2 IP addresses, one for each link.
- Given Address space is 192.168.31.0/24

VSLM Example

1. Real requirement, add subnet and broadcast addresses.
2. Select the block size using the following formula.

Real requirement <= 2^h (block size)

3. Arrange segments in descender order.

Segment	Host requirement	Real requirement (Subnet and Broadcast addresses	Block size
Sales	74	76	$2^7 = 128$
Production	52	54	$2^6 = 64$
Administration	28	30	$2^5 = 32$
Link 1	2	4	$2^2 = 4$
Link 2	2	4	$2^2 = 4$
Link 3	2	4	$2^2 = 4$

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 34

Step by step VLSM Subnetting

1. Consider the real requirement, including the reserved subnet and broadcast addresses. This number is obtained adding 2 to the Host requirement (Subnet address and the Local Broadcast address).
2. Select the appropriate block size. This number can be obtained from the formula.

Real requirement <= 2h, where **h** is the number of hosts.

Note. You have to select the highest value of h where the formula is true.

In this case, for Production subnet you have 54 <= 2h in this case the answer is h = **6**, and 26 = 64. Therefore the block size need it is 64.

3. The next step is arranging segments in descender order, based on the block size.

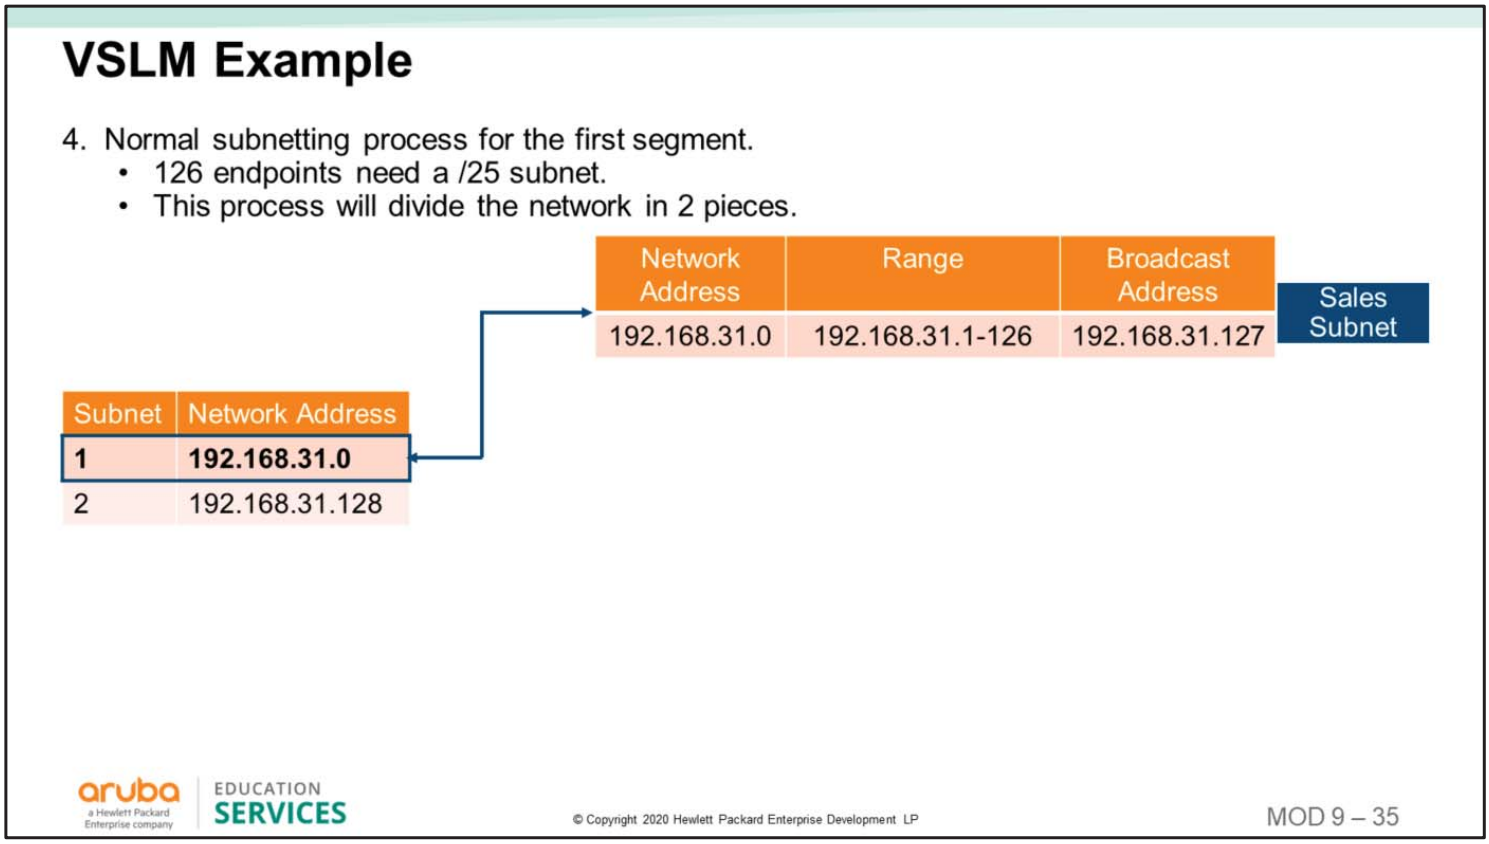

4. Do the normal subnetting process for the first segment, the one that has the larger block size.

By analyzing the host needs, you can conclude that a /25 subnet mask is needed to allocate 126 endpoints. Notice that this subnet mask will divide the original network into 2 subnets. The first one can serve the Sales department needs and the second one becomes the subnet that will be divided to serve the next block size.

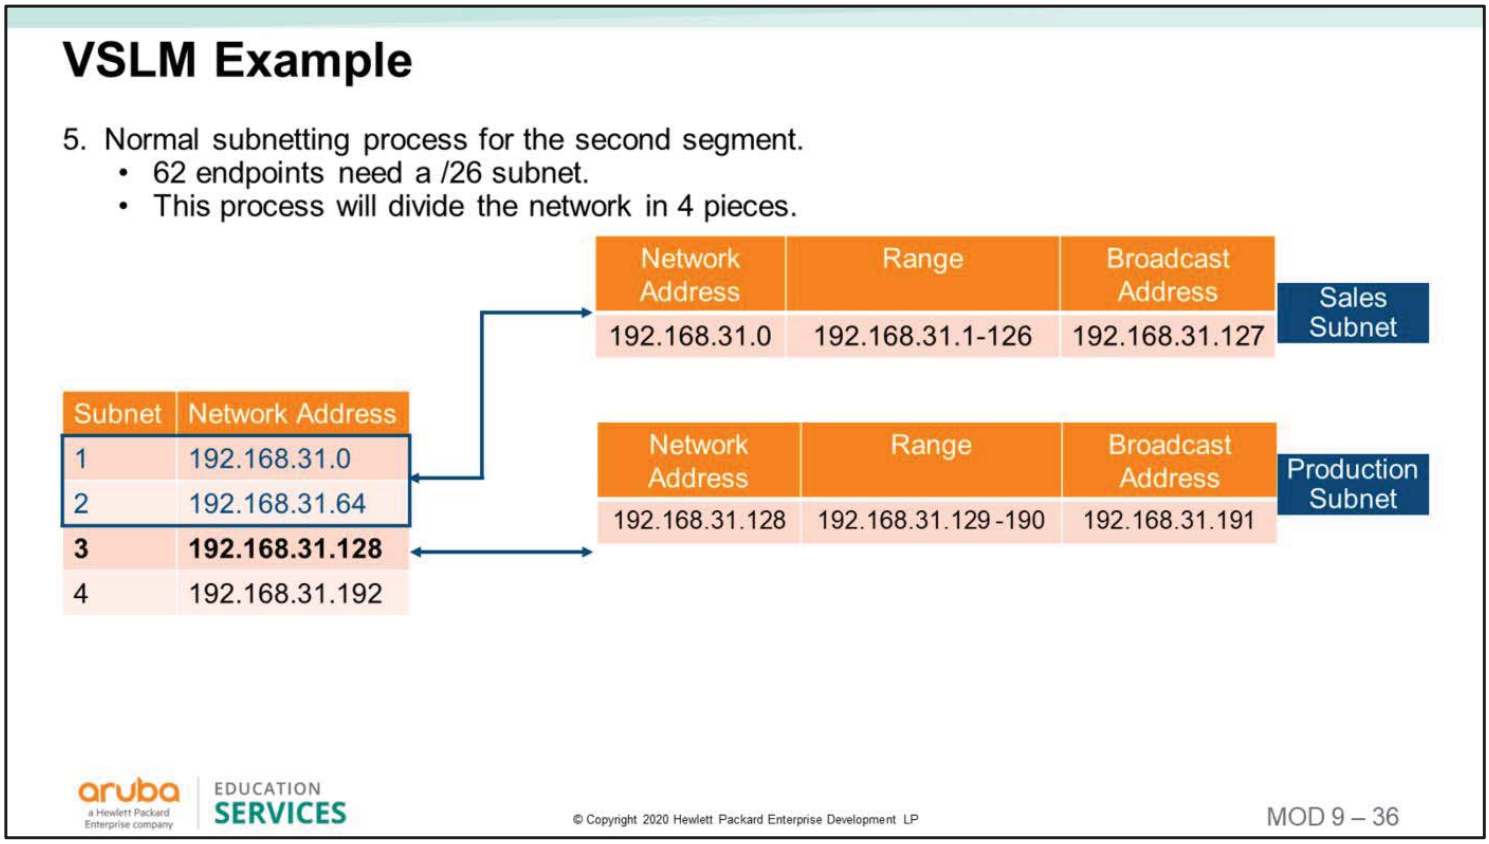

5. Apply the subnetting process for the next block size. The production department in this case.

By analyzing the needs you conclude that a /26 network mask is needed to accommodate 62 hosts. Using this subnet mask it divides the original Network into four different subnets. However the first two subnets have been used by the previous department. So you have to start using the third one.

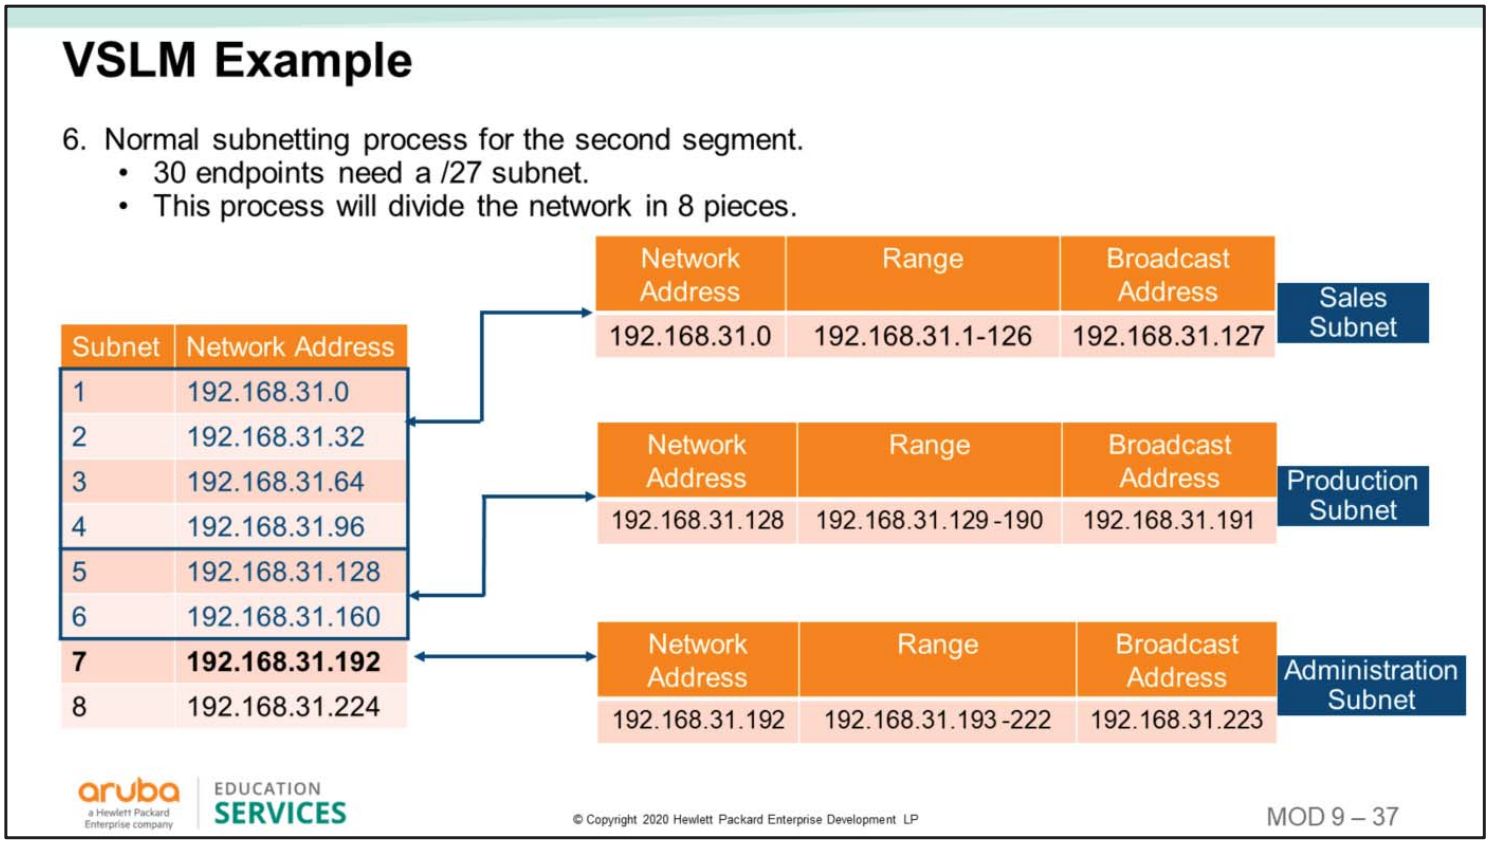

6. Apply the subnetting process for the next block size. The administrative department in this case.

By analyzing the needs you can conclude that a /27 network mask is needed to accommodate 30 hosts. Using this subnet mask it divides the original network into eight subnets. However the first two subnets have been used by the previous department. So you must start at the seventh one for this use case.

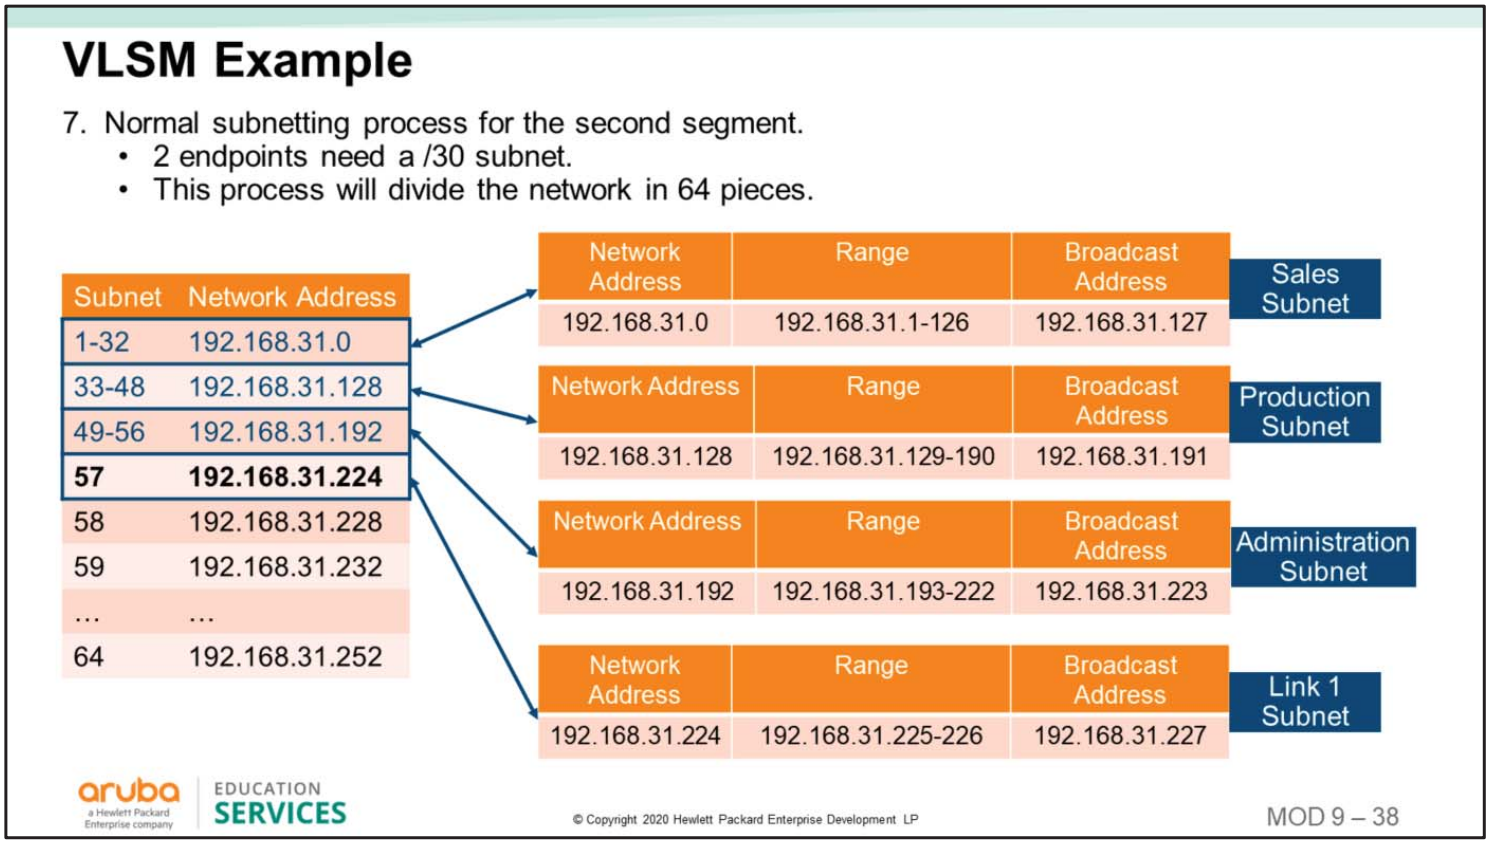

7. Apply the subnetting process for the next block size – Link 1 in this case.

By analyzing the needs you can conclude that a /30 network mask is needed to accommodate 2 hosts. Using this subnet mask it divides the original Network into 64 different subnets. However the first 56 subnets have been used by the previous department. So you must start at the 57th one.

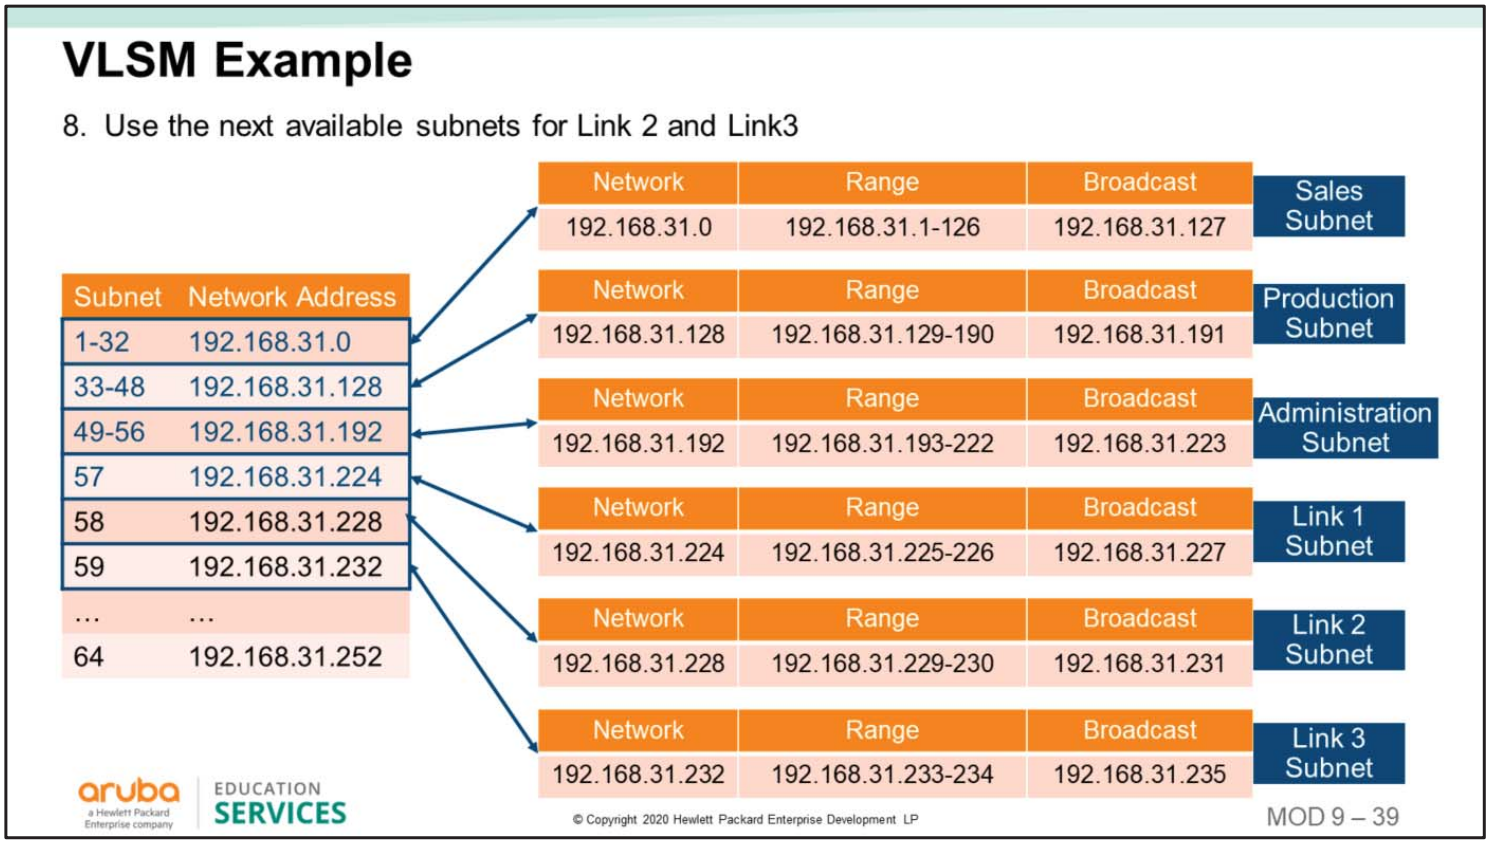

8. Since the next block size is the same as the previous one, simply use the next available subnet for the Link 2 and Link 3.

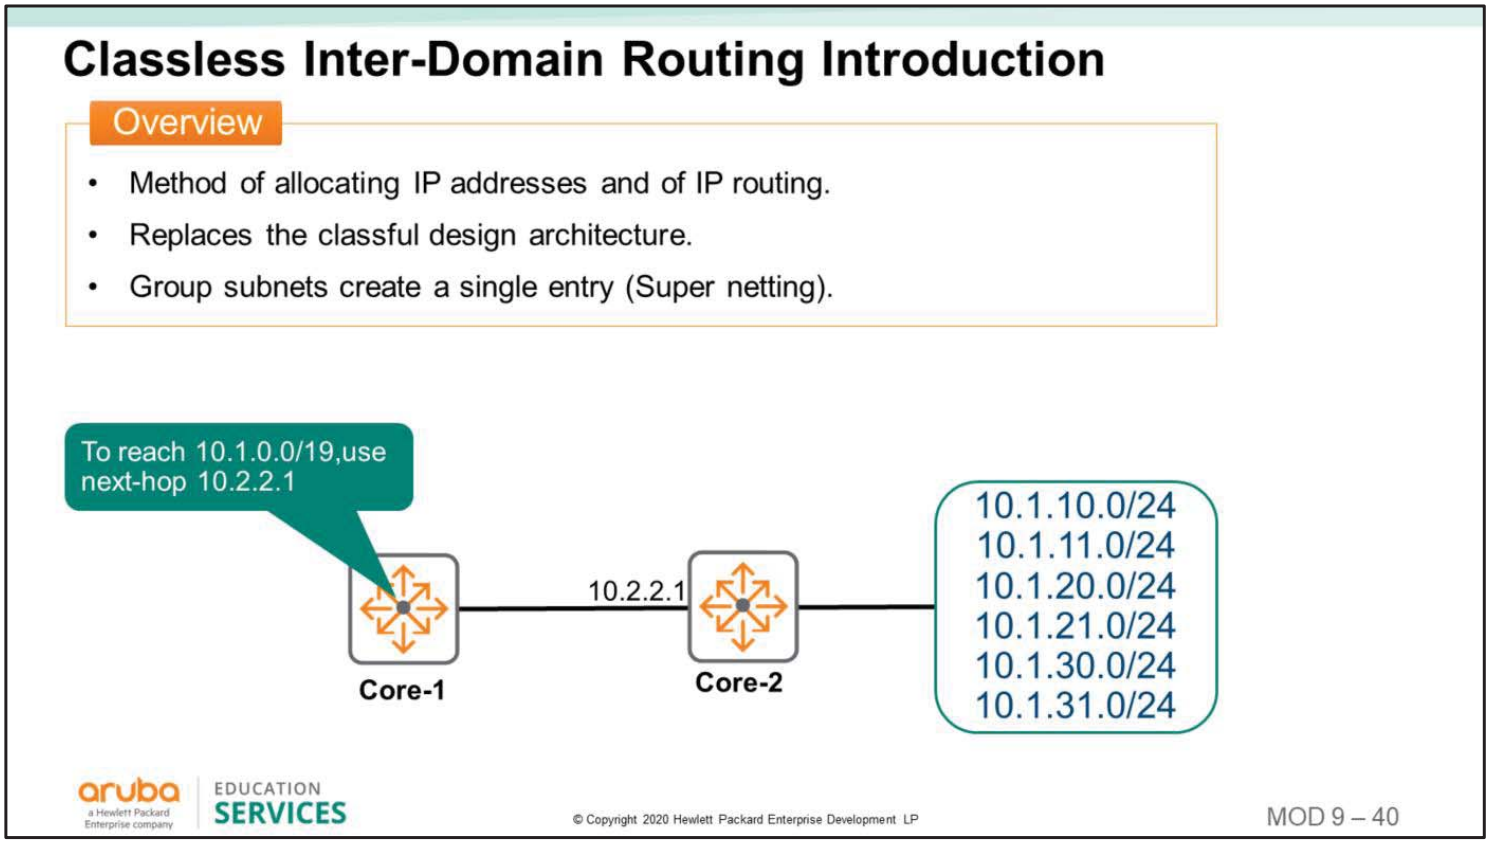

CIDR is a method for allocating IP addresses and IP routing that replaces the previous architecture of Classful network design. CIDR standardizes the use of an **IP address/subnet mask** notation. CIDR also helps to aggregate routes to make a more efficient network.

You have learned how to divide a network address space into subnets to allocate organizations needs. The ability to break up one network into multiple subnets has great advantages for you as a design or administrative expert, and for the network. However, from the router's perspective, some new inefficiencies arise. Now they must store each individual subnet is as a separate entity in their routing tables, and in their topology databases. This can use memory, processing, and bandwidth. It is not very efficient, especially when the next hop for all those subnets is the same device. The individual subnets are maintained, but several entries can be grouped into a single, unique entry that summarizes all individual entries.

The figure shows a scenario where Core-2 has six subnets, which you have recently created. Without CIDR implemented in Core-1, each individual network will be listed as an entry in the Routing table. So, six entries in Core-2 means 6 entries in Core-1. Six hundred entries in Core-2 means 600 entries in Core-1.

Instead, your routers can evaluate the common bits among these addresses and perform a route summarization.

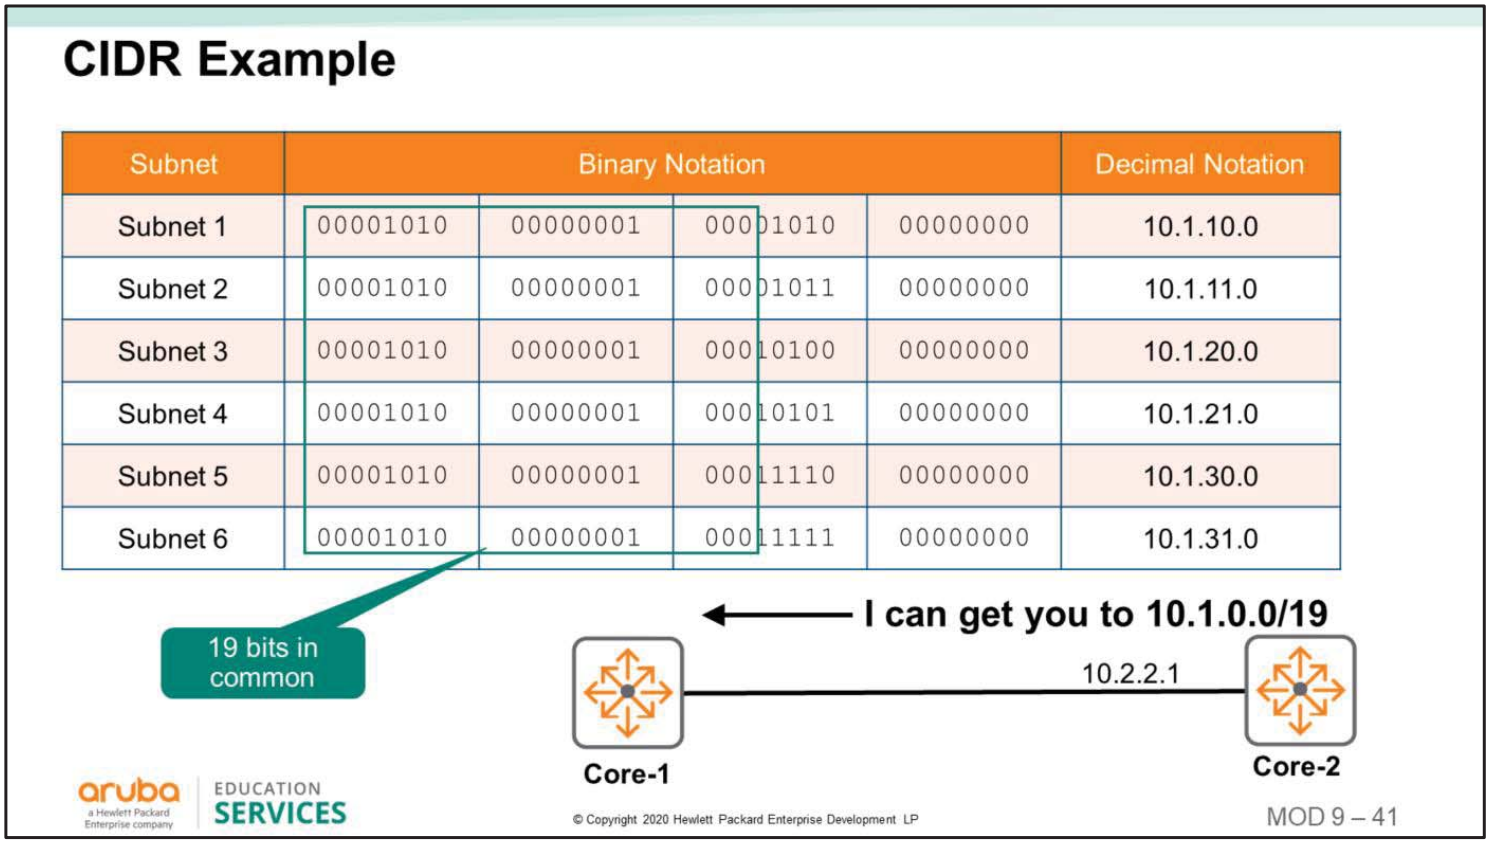

The figure shows the six addresses, converted into binary. This makes it clear – for the range of addresses between 10.1.10.0 and 10.1.31.0, the first 19 bits are identical. Stated another way, if you ignore the last 13 bits of the address, this range of addresses is identical.

This means that Core-2 can perform this calculation, and instead of advertising six addresses with a /24 mask, it advertises one address with a /19 mask – 10.1.0.0/19. Understand that reality has not changed. The individual subnets still exists - unchanged. Core-2 is simply summarizing what it says to Core-1. Core-2 is saying, “if you need to route a packet to any destination where the first 19 bits matches the pattern 00001010.00000001.000, just send those packets to me.”

Note. CIDR implicitly uses VLSM when summarization takes place in a router or multilayer switch. Not all the entries in routing tables on these devices will have the same subnet mask.

Question #3

VLSM Example, CIDR Example

Which statements below are true?

a. Given address 172.20.66.5/19, the subnet is 172.20.64.0

b. Subnetting a Class A address will increase number of possible host addresses.

c. Given address 10.1.187.5/16, this scheme allows for a total of 65,534 hosts on the subnet.

d. When you use CIDR to aggregate a group of routes, assigned IP addresses are modified.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 9 – 42

Correct answers:
a.
c.

Lab Activity

9.1 Subnetting and VLSM

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Congratulations! This module should have significantly elevated your ability to design and deploy more complex, scalable IP networks. You leveraged your knowledge of binary and decimal number systems to explore subnetting – how to more efficiently use available IP addresses. Related concepts included the various address classes and reserved address space, and why classful routing can seem quite rigid when compared to classless routing. Then you applied your subnetting knowledge to network design and analysis scenarios.

Then you explored Variable Length Subnet Masking (VLSM) and Classless Interdomain Routing (CIDR). You learned that VLSM help you to assign address space more efficiently, while CIDR helps the routers to advertise and work with that address space more efficiently.

Finally, a lab activity gave you more experience with these concepts.

aruba
a Hewlett Packard Enterprise company
IP Routing Part 3
ArubaOS-CX Switching Fundamentals
Rev #20.21
aruba EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 10 – 1

Objectives

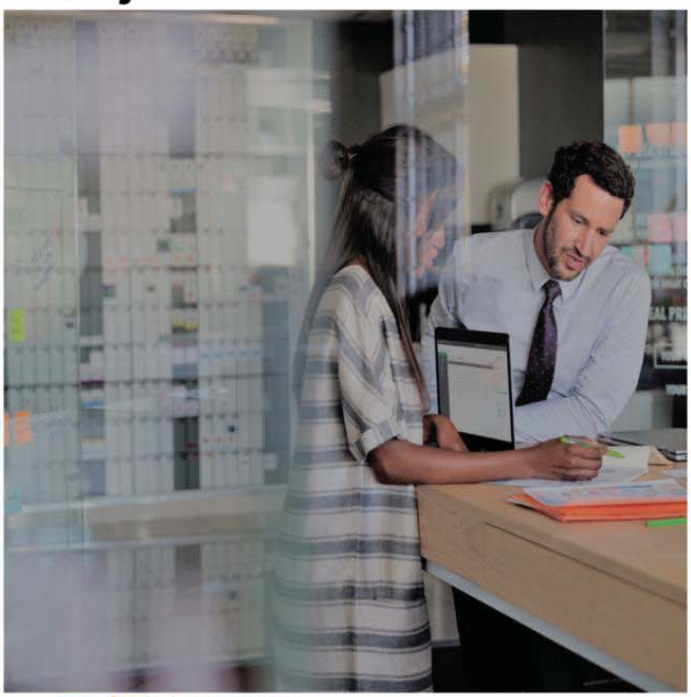

- Learning Objectives
  - Describe the operation and use of route types and Administrative Distance
  - Explain Layer 3 scalability issues
  - Routing Protocols
  - Distance vector routing protocols
  - Link-state routing protocols

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 10 – 2

After completing this module, you will be able to describe the operation and use of route types and how routing decisions are sometimes made based on an Administrative Distance value.
You will be able to explain Layer 3 scalability issues.
This module also focus on learning about routing protocols in general, and the differences between interior and exterior routing protocols, as well as the difference between Distance Vector and Links State routing protocol.
You will learn that Link State-based routing protocols are more efficient, scalable, and secure, and that the most common implementation is OSPFv2.

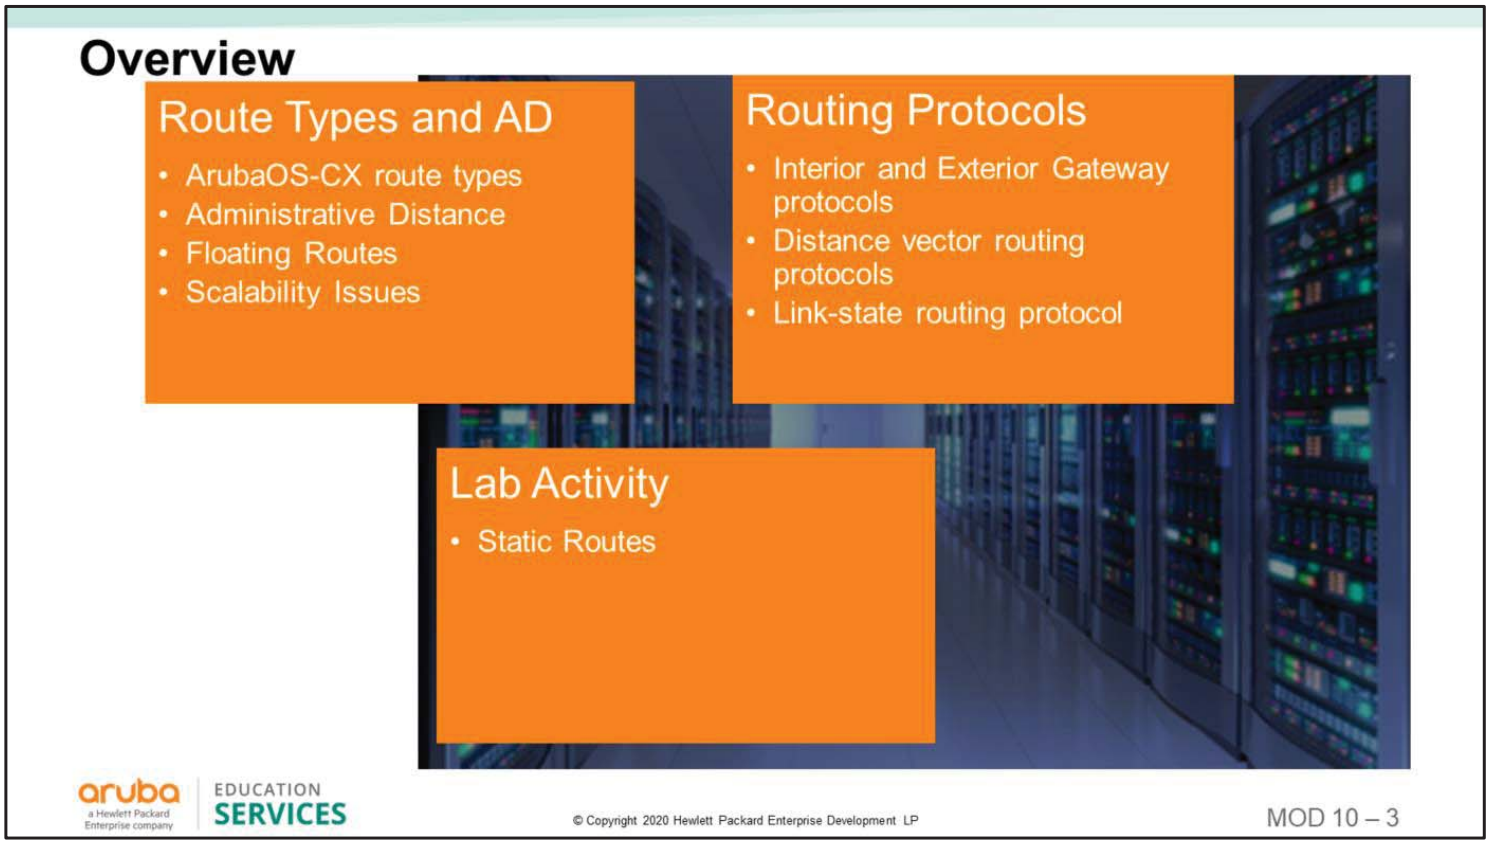

In this module, you will learn about Route types – how a router can learn about the same path from different sources, and then use Administrative Distance to choose the most trustworthy source. Next you will learn to apply this knowledge with a technique known as floating static routes, before exploring some Layer 3 scalability issues.

You will learn about the routing protocols that help to scale the routing solution. You will also know the difference between distance vector and link-state routing protocols.

Finally, a lab activity will give you more experience with these concepts.

Route Types and Administrative Distance

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

Route Types in ArubaOS-CX

ArubaOS-CX supported route types

- **Connected**: subnet is physically connected to the device
- **Local**: subnet is configured inside the device (loopback or SVI)
- **Static**: route added manually as a static route
- **RIPv2**: route learned via RIPv2
- **OSPF**: route learned via Open Shortest Path First
- **BGP**: route learned via BGP

SVI Configuration:

```
Switch(config)# interface vlan 1
Switch(config-if)# ip address 10.1.1.1/24
Switch(config)# interface 1/1/1
Switch(config-if)# vlan access 1
```

Routed port Configuration:

```
Switch(config)# interface 1/1/1
Switch(config-if)# routing
Switch(config-if)# ip address 10.1.1.1/24
```

```
Switch# show ip route
Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.1.1.0/24, vrf default
        via  vlan1,  [0/0],  connected
10.1.1.1/32, vrf default
        via  vlan1,  [0/0],  local
```

aruba a Hewlett Packard Enterprise company EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 10 – 5

You learned about the IP routing table in a previous module, and that a routing device makes Layer 3 forwarding decisions based on the best-path entries in this table. The entries in the routing table are populated manually or through a routing protocol. Route types that can exist in an ArubaOS-CX switch include the following:

- Connected:
- Local
- Static
- RIPv2
- OSPF
- BGP

Connected and Local entries

Consider the following configuration, where the Switching Virtual Interface (SVI) for VLAN1 receives an IP address and port 1/1/1 is assigned to VLAN 1.

Switch(config)**# interface vlan 1**
Switch(config-if)**# ip address 10.1.1.1/24**
Switch(config)**# interface 1/1/1**
Switch(config-if)**# no routing**
Switch(config-if)**# vlan access 1**

This configuration generates some entries in the routing table, you can verify it using the **show ip route** command. In this case you can see that 2 entries are created in the routing table.

The first entry is listed as **10.1.1.0/24** and is a **connected** type route. This entry means that the Subnet is available physically connected to the switch and there is no need for a next hop device.

The second entry is listed as **10.1.1.1/32** and is a **local** type route. This entry means that the IP address that was previously configured is available for the routing process inside the switch. This could be a loopback interface, or as shown here, a Switch Virtual Interface (SVI) for VLAN 1.

Note. You may notice that the local entry uses a /32 subnet mask. This is a special mask that indicates that subnetting cannot happen (no divisions are permitted) in the address space provided. This mask is useful to refer to individual hosts. In this case only one IP address is permitted to be configured in the SVI.

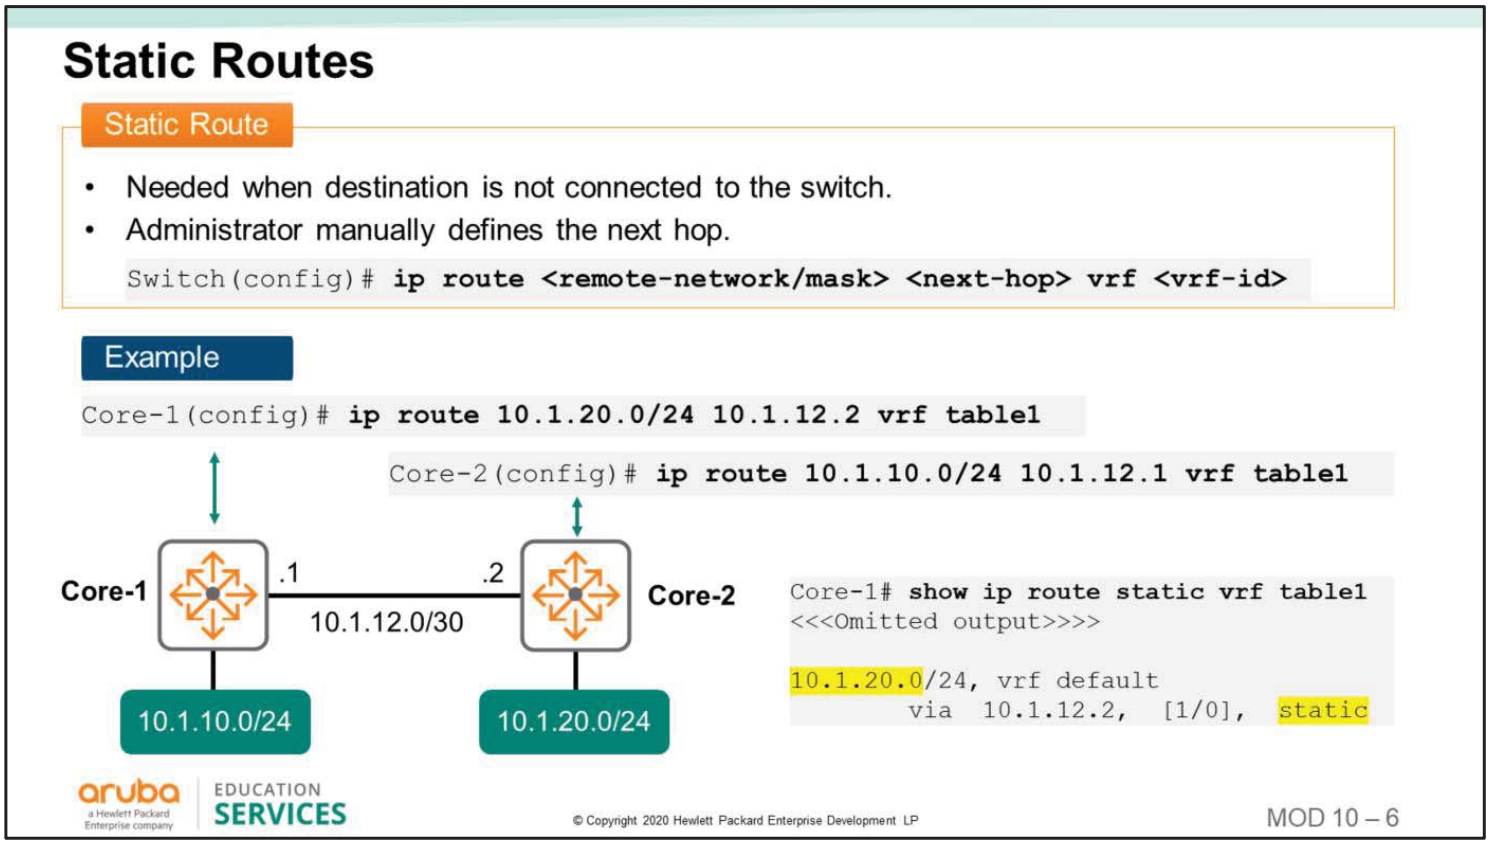

Static routes

You need a static route when the destination network is not directly or physically connected to the switch or router, and no routing protocol like OSPF is advertising that route. In this case, you manually define the destination network and the next hop in the path. The next hop path must be directly connected to the device where you apply the static route.

In ArubaOS-CX, use the **ip route** command to build a static route.

Switch(config)**# ip route <remote-network/mask> <next-hop> vrf <vrf-id>**

Suppose that you need endpoints in subnet 10.1.10.0 to communicate with others in 10.1.20.0. Core-1 switch has two physically connected networks: 10.1.10.0/24 and 10.1.12.0/30. However, the switch is not physically connected to 10.1.20.0/24 subnet. So you add the route below:

Core-1(config)**# ip route 10.1.20.0/24 10.1.12.2 vrf table1**

Core-2 switch also needs a static route to complete the path, from its perspective it only has entries in its routing table for 10.1.20.0/24 and 10.1.12.0/30 networks. You must add a new route:

Core-2(config)**# ip route 10.1.10.0/24 10.1.12.1 vrf table1**

To verify this entry exists on the routing table issue the **show ip route vrf table1** command.

Administrative Distance

Overview

- A trust rating for route entries
- Used if one route is learned by two or more methods
- Lower distance is more preferred

```
Switch# show ip route

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.1.20.0/24, vrf default
        via  10.1.12.2,  [1/0],  static
10.1.20.0/24, vrf default
        via  10.1.12.2,  [110/200],  ospf
```

Populating method	Administrative Distance
Directly connected	0
Static Route	1
External BGP	20
OSPF	110
Internal BGP	200
Unknown or unused	255

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 10 – 7

Sometimes a router will learn a route from two different sources. Maybe a router is running BGP, and OSPF, and you have entered some static routes. Perhaps all three methods have taught the router about network 10.1.20.0/24. Which source for routing information should the router trust? The source with the lowest Administrative Distance.

This is very much like humans learning information. You are hopefully very close to your mother...you trust her. You might say she has a very low administrative distance. If some stranger gives you conflicting advice, you may choose to listen to your mother. The stranger has a higher administrative distance.

The table shows the Administrative Distances for each routing protocol.

For example, suppose that Core-1 learns about network 10.1.20.0 from OSPF and from a static route. Core-1 will use the path specified by the static route, due to the lower, more trusted administrative distance.

Floating Static Routes

Overview

- Manipulate admin distance to create primary/backup routes
- Improves network resiliency

Example

- Prefer ISP1 Internet connection over ISP2
- In case ISP1 router fails, ISP2 should be used

```
Core-1(config)# ip route 0.0.0.0/0 10.1.251.2
Core-1(config)# ip route 0.0.0.0/0 10.1.252.2 distance 10
```

```
Core-1# show ip route

0.0.0.0/0, vrf default via  10.1.251.2,[1/0],static
0.0.0.0/0, vrf default via  10.1.252.2,[10/0],static
```

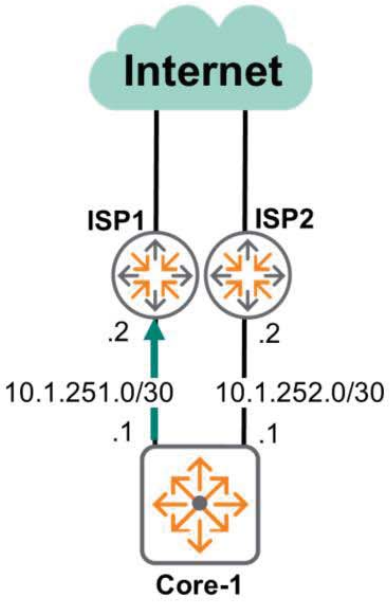

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 10 – 8

Manipulating the administrative distance can be used to create primary and backup routes, to improve network resiliency.

Imagine that you have two Internet connections that use different ISPs. The first connection has a higher bandwidth than the second one and therefore the administrator wants to use it as the Primary connection. The secondary ISP connection will be a backup, waiting to be used in case the primary link fails.

To do this, configure two static routes. One to the “all routes” destination of 0.0.0.0 with ISP-1 as the next hop. You let its default administrative distance remain at 0.
Configure a second static route to the “all routes” destination of 0.0.0.0 with ISP-2 as the next hop, changing the administrative distance to10.

Core-1(config)**# ip route 0.0.0.0/0 10.1.251.2**
Core-1(config)**# ip route 0.0.0.0/0 10.1.252.2 distance 10**

Since the ISP-1 route has a lower administrative distance it will be used exclusively. The ISP-2 path shall only be used if the primary link fails. You can verify this behavior using the **show ip route** command.

Note. The Network 0.0.0.0 with a subnet mask of 0.0.0.0 is a super network that summarizes all possible addresses in the IPv4 scope. It means any address.

Note. Is important to differentiate between default gateway and default route. The first concept refers to a device that does not have the routing feature enabled, such as a host or a Layer 2 switch. In that situation we just provide the information about where to send

packets that are not in the same subnet. Default route defined as 0.0.0.0/0 is used in devices where the routing process is enabled, such as in routers and MultiLayer switches.

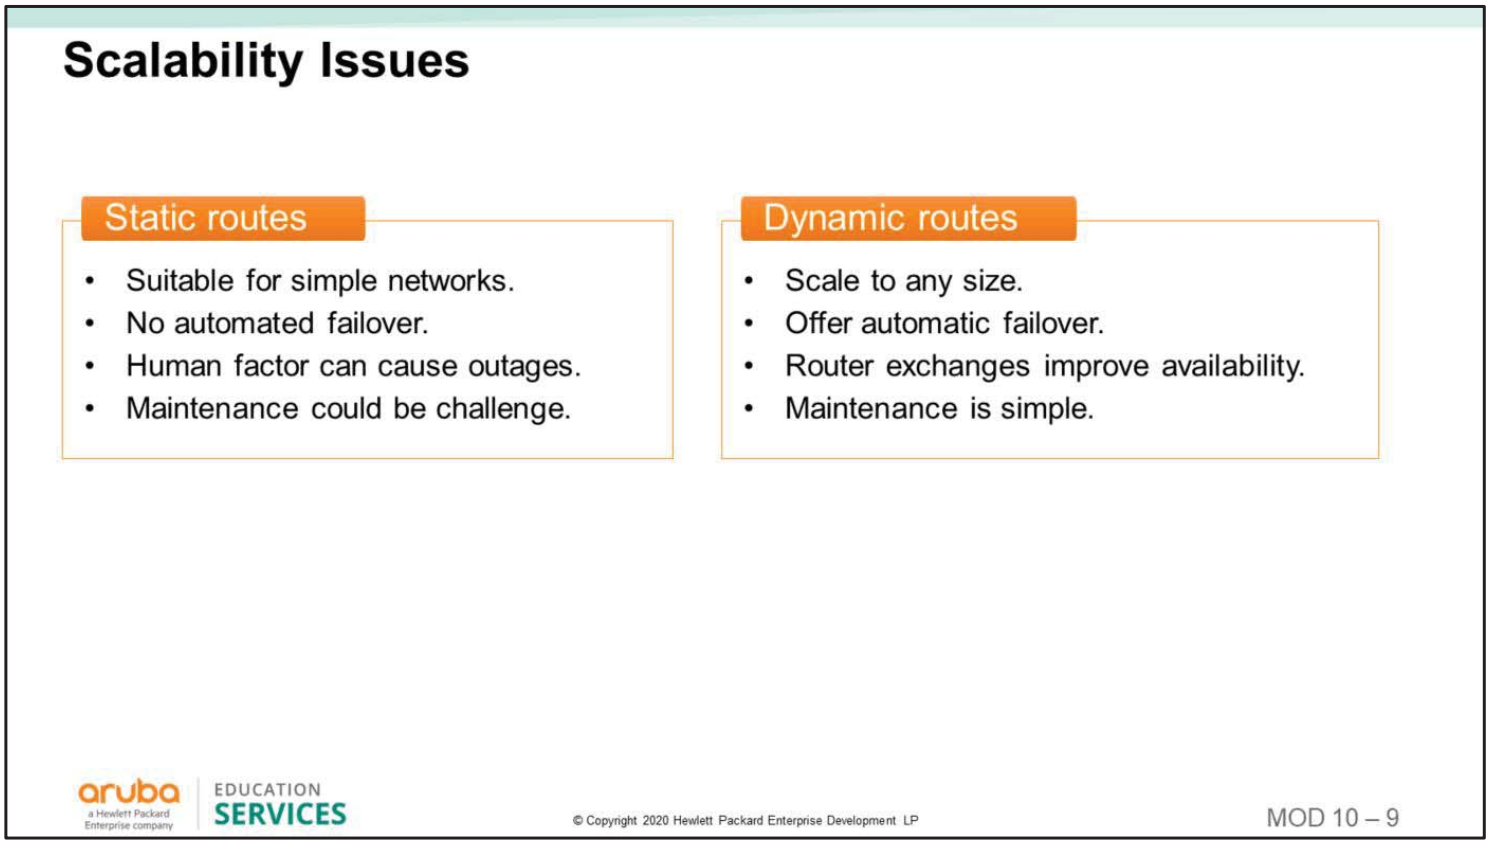

Working with static routes is simple and easy for small networks that have a few subnets. However, when an organization has hundreds or even thousands of subnets, static routing is not an efficient method to administer subnets. There is no automatic route advertisement, you must manually configure everything. Other than the simple Floating Static routes you just learned, there's no dynamic failover mechanism.

The human factor is also a big consideration. Administrators can cause serious problems if a route is placed in the incorrect device or if the route is not properly configured. This can create Layer 3 route loops, traffic blackholing, and lost connectivity.

Recovering from a routing device failure could be slow if Static routing is used. You must manually configure new alternative paths. This is inconvenient at best, especially when dynamic routing would automatically do this for you.

Dynamic routing protocols are far more scalable, handling thousands or even millions of routes across multiple routing devices. They automatically failover to alternate paths with little to no downtime. This is because routers constantly exchange messages that keep the network available and minimize your managerial workload.

Question #1

Administrative Distance

Suppose that a router has learned about network 172.18.37.0/24 from three sources – OSPF, Internal BGP, and a static route. Which statements are true about this router's path selection?

a. The router chooses the OSPF route, because it is trusted more than BGP.

b. The router chooses the best path based on the lowest cost.

c. The router chooses the option with the highest Administrative Distance.

d. The router chooses the static route.

Correct answer:
d.

Routing Protocols

Interior and Exterior Gateway Protocols

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 10 – 11

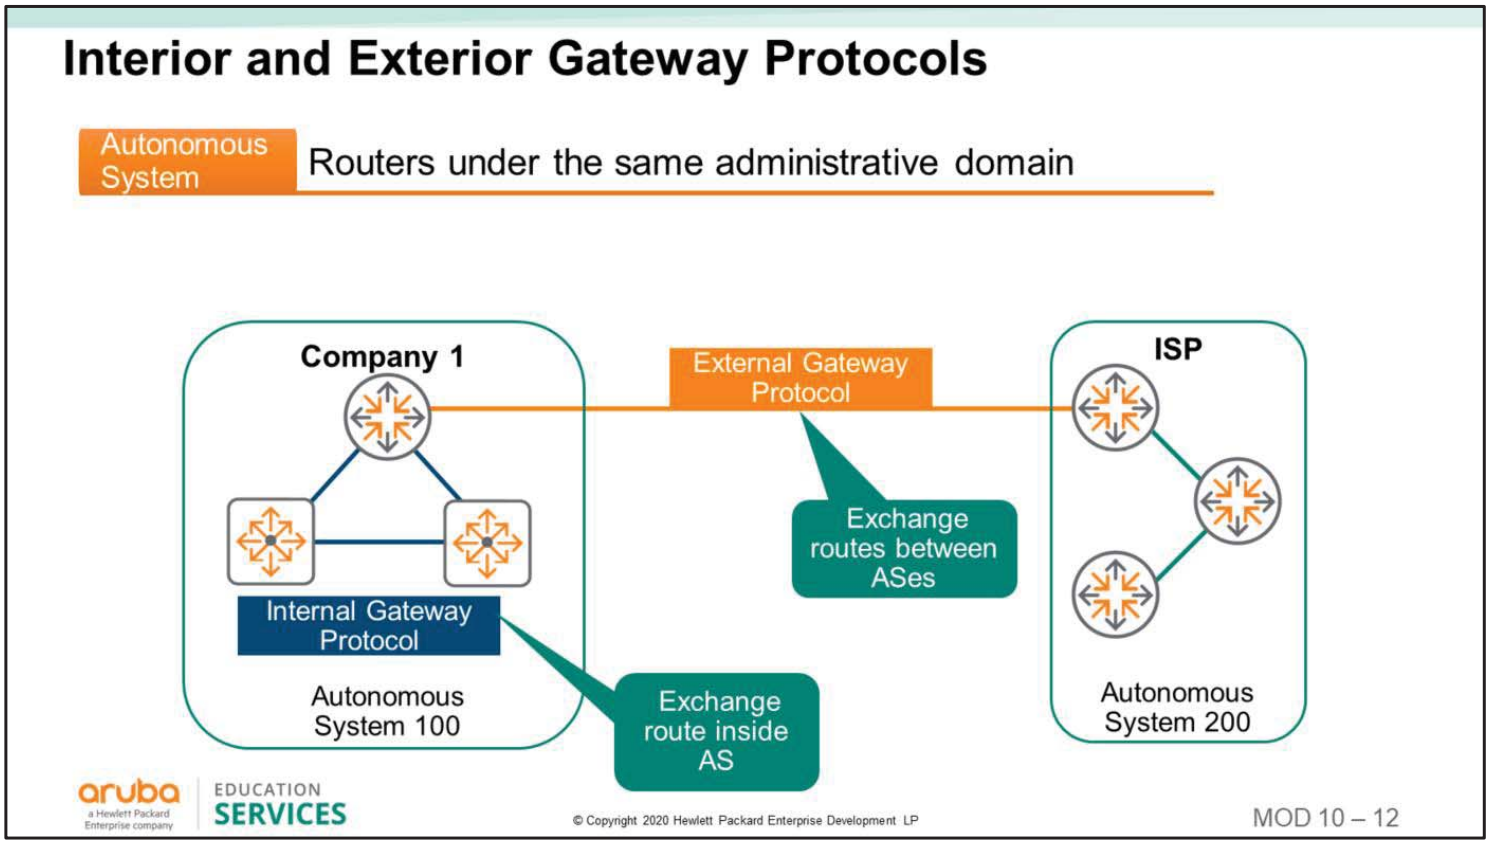

An Autonomous System (AS) is a collection of routers under a common administrative domain. Your Internet Service Provider (ISP) owns their own internal network – they have autonomy over this system – it is their AS. Each corporation has autonomy over their network – it is their AS.

To route packets inside an AS, each company chooses an Interior Gateway Protocol (IGP). An IGP is simply a routing protocol that runs inside an AS. Examples of IGPs include the Routing Information Protocol (RIP), Intermediate System – Intermediate System (IS-IS), and Open Shortest Path First (OSPF). Each company can use the IGP that works best for them.

In the figure AS 1 has three routing devices. They use RIP to exchange network information to discover all paths inside the AS, and then find the best path for each network. AS 2 has a similar arrangement, but they have chosen to use OSPF.

These companies might be partners and want to share information. They connect to the internet, via an ISP that owns AS 3. They wish to route between AS 1, AS 2, and AS 3. Robust, enterprise-class routing between these Autonomous Systems requires an Exterior Gateway Protocol (EGP). The only EGP currently in use today is the Border Gateway Protocol (BGP).

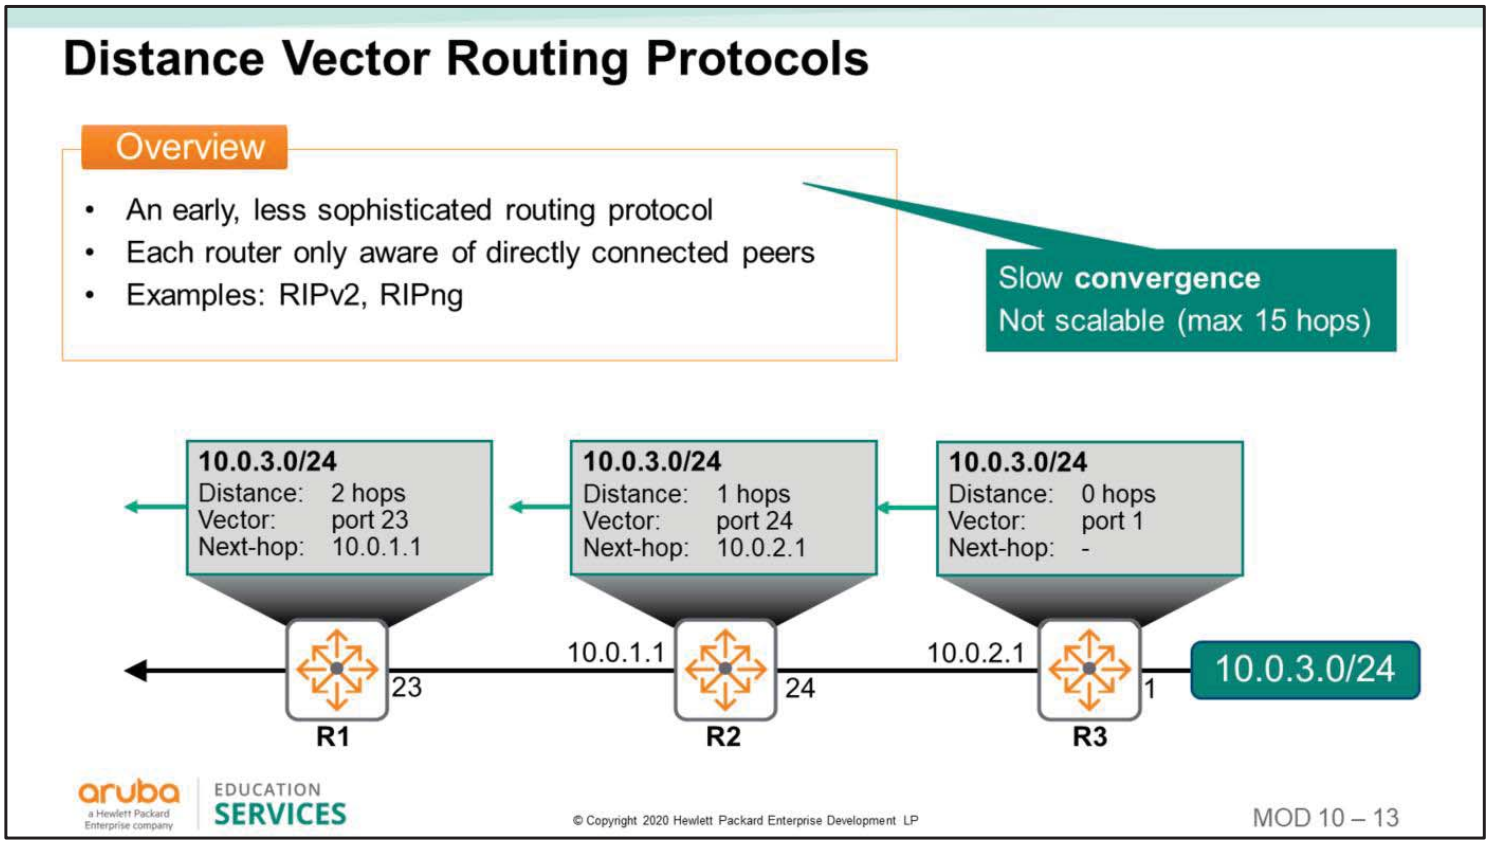

The earliest methods devised to exchange IP information use Distance Vector algorithms. Each router advertises its distance from each network, and the vector or direction packets must travel to get to that network.

The distance is measured in “hop count”. In the figure, how many routers must R3 “hop over” to get to network 10.0.3.0/24? The answer is 0 routers – it is directly connected to that network.
What vector or direction must R3 forward packets to reach 10.0.3.0/24? The answer is “out port 1”. It is as if the router says, “to get to any host on destination network 10.0.3.0/24, which is 0 hops away, I must forward packets out my local Port 1. Using something like RIP or RIPv2, R3 advertises its distance and vector to R2.

R2 receives this and knows, “If R3 is 0 hops away, and I must hop over R3 to get to that network, then I am 1 hop away. To get to destination 10.0.3.0/24, which is 1 hop away, R2 must forward packets out local port 24, to next-hop router 10.0.2.1. R2 advertises its distance and vector to R1.

R1 receives this and knows, “If R2 is 1 hop away, and I must hop over R2 to get to the destination, then I am 2 hops away”. To get to 10.0.3.0/24, which is 2 hops away, R3 forwards packets out local port 23, to next-hop router 10.0.1.1”.

Distance Vector routers are not aware of the entire network topology. They only perceive their directly connected routing peers. This simplifies operation, but also limits operation:
- Slow convergence: When a failure occurs, Distance Vector protocols can take minutes to converge, depending on network size, complexity, and architecture.
- Limited scalability: Each router can be no more than 15 hops away from any other router.

Understand that you could still have dozens, or even hundreds of routers, as long as it is a relatively flat architecture.

Distance Vector protocols include the Routing Information Protocol (RIP), RIP version 2 (RIPv2), and RIP Next Generation (RIPng). RIP and RIPv2 do not support security feature to protect the advertisement messages that are exchange between routing devices.

Due to their limited scalability, performance, and security, the ArubaOS-CX does not support any Distance Vector protocols. Instead, Aruba supports the more robust Link State routing protocols.

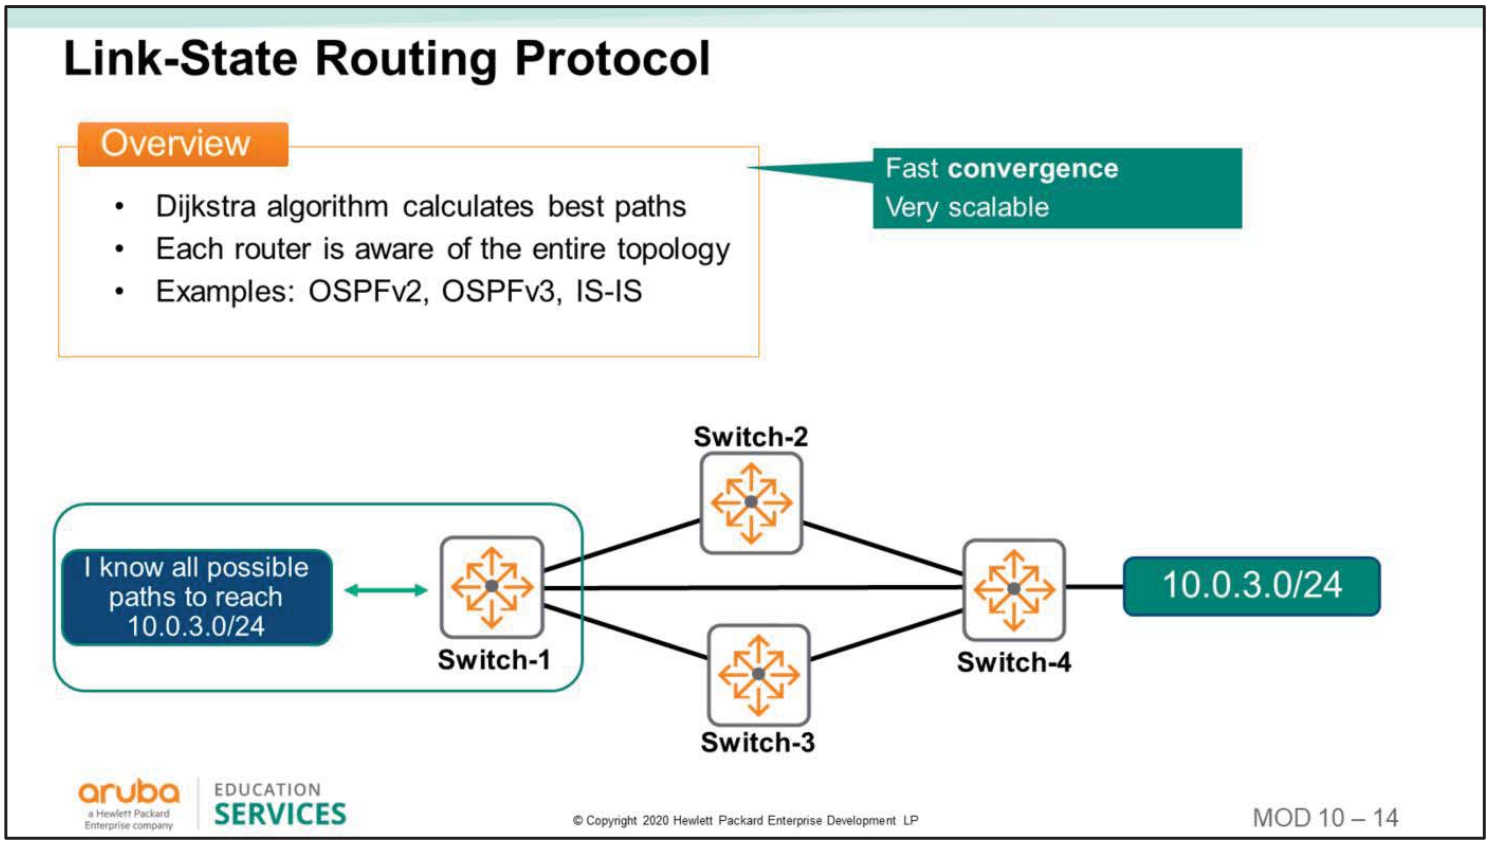

Link-State protocols resolve limitations and disadvantages of the Distance Vector protocols. Routers that run Link-State protocols possess information about the complete network topology. This approach enables them to independently calculate their routing table using the more sophisticated Dijkstra algorithm to select the best path for every network destination.

Actual implementations of Link-State protocols include Open Shortest Path First version 2 (OSPFv2), which is used for IPv4; OSPFv3, which is used for IPv6, and Intermediate System to Intermediate System (IS-IS). IS-IS is far less common than OSPFv2, and so this module is focused on OSPF.

These protocols provide faster convergence times compared with Distance Vector protocols. Another advantage is scalability.

Many years ago, one disadvantage of protocols like OSPF is that they created more CPU, memory, and bandwidth utilization than protocols like RIP. However, routers have become more capable over the years, and our ability to design more efficient network architectures has greatly reduced this concern. Furthermore, the ArubaOS-CX multilayer switches are designed to meet today's network requirements. You can verify this information by running the **show capacities ospfv2** command.

```
Switch-6300# show capacities ospfv2

System Capacities: Filter OSPFv2
Capacities Name
Value
```

```
-------------------------------------------------------------
-----------------
Maximum number of OSPFv2 areas configurable in the system
128
Maximum number of OSPFv2 interfaces configurable in the
system         128
Maximum number of OSPFv2 interfaces per area in the
system             128
Maximum number of OSPFv2 neighbors allowed across all
VRFs             128
Maximum number of OSPFv2 processes supported across each
VRF              1
Maximum number of routes in OSPFv2 supported across all
VRFs         64000
```

Note. This training will focus the attention on OSPFv2

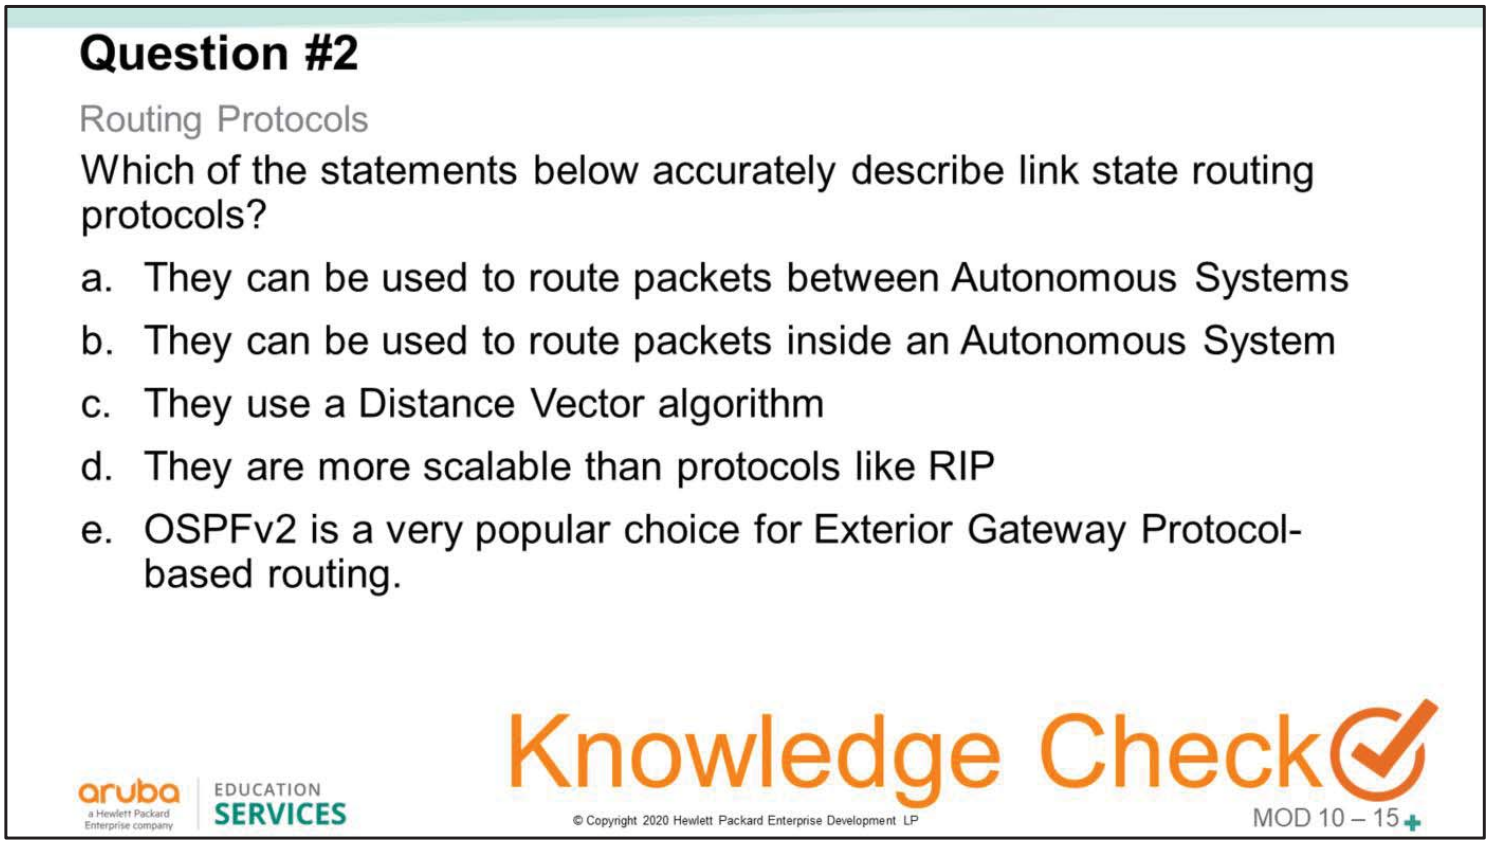

Correct answers:
b.
d.

Lab Activity

10.1 Static Routes

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Summary

Route Types and AD

- ArubaOS-CX route types
- Administrative Distance
- Floating Routes
- Scalability Issues

Routing Protocols

- Interior and Exterior Gateway protocols
- Distance vector routing protocols
- Link-state routing protocol

Lab Activity

- Static Routes

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

aruba
a Hewlett Packard Enterprise company
OSPFv2 Single Area
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 11

Objectives

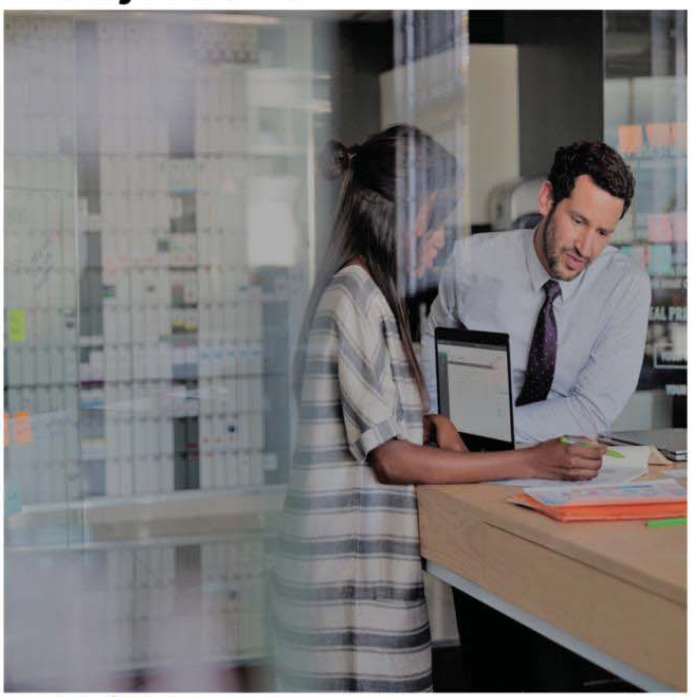

–Learning Objectives

- –Describe OSPF general operation
- –Explain areas and Router IDs
- –Describe hello messages
- –Describe OSPF network types
- –Describe DR election and neighbor states
- –Explain OSPF LSA types, path selection, and convergence
- –Configure OSPF
- –Use Path costs to manipulate routes
- –Describe and configure passive interfaces

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 11 – 2

After completing this module, you will be familiar with routing protocols, with a focus on Open Shortest Path First (OSPF).
You will be able to describe OSPF general operation and explain areas and Router IDs.
You be able to describe the purpose of hello messages, network types, and Designated Router Election.
You will learn to explain OSPF LSA types, path selection, and convergence, before learning how to configure OSPF, including path costs and passive interfaces.

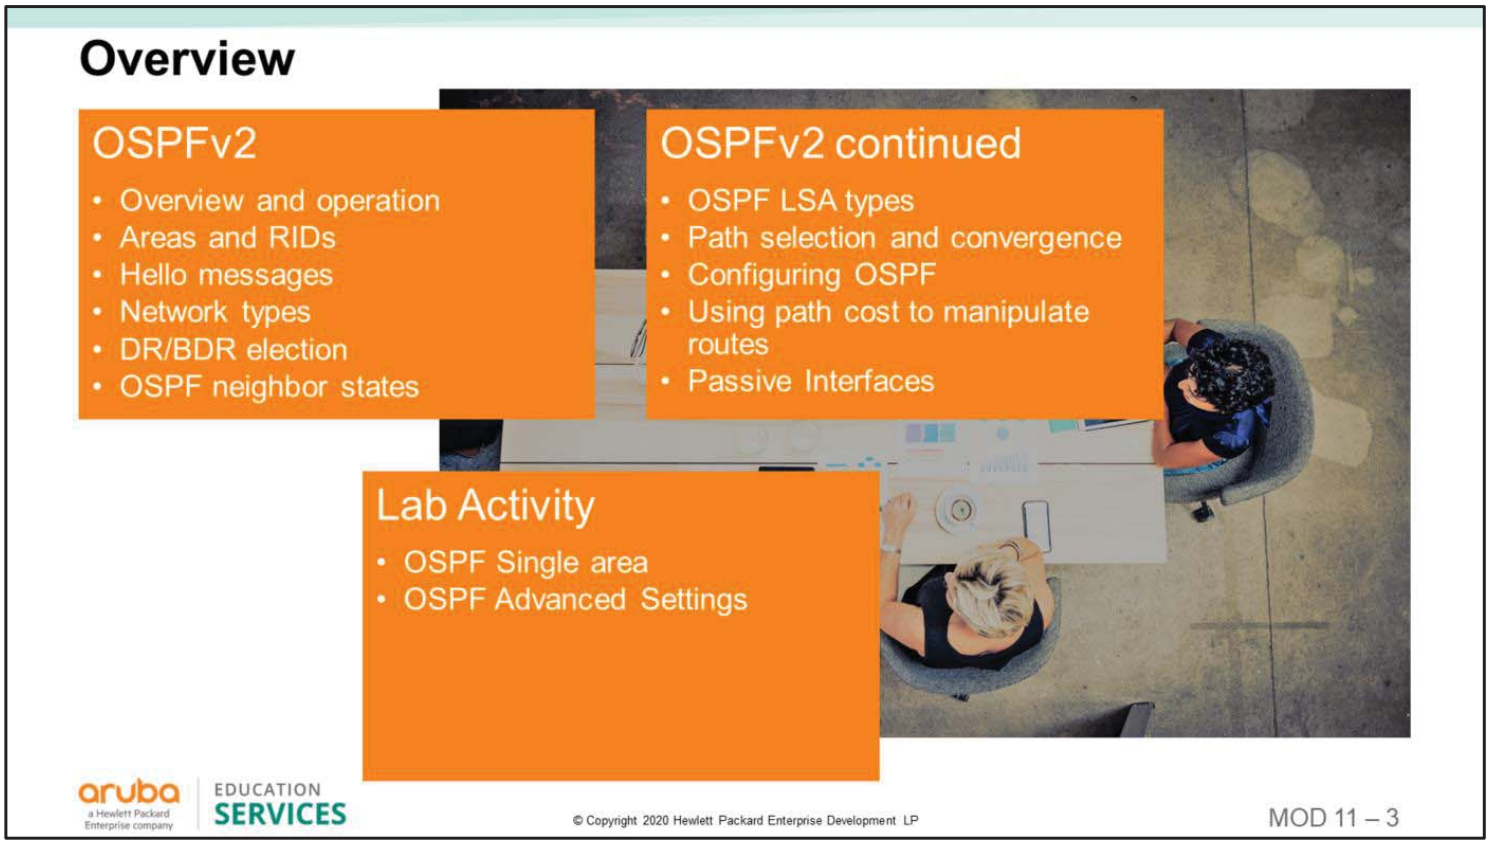

OSPF may be the most popular option for corporations to route traffic within their organization.
OSPFv2 is used to route IPv4 packets within a corporate internetwork. You will explore OSPF operation, including OSPF areas, Router IDs, and various message types and neighbor states.
You will learn how subnet information is advertised by various types of Link State Advertisements (LSA). Then you will learn how these LSAs build a database of all paths, and how that database is then use to build a routing table – a list of best paths over which to forward end user data traffic.

Finally, you will learn how to configure OSPF, and apply that knowledge with a Hands-on lab activity.

OSPFv2

Router ID, Hello messages and Passive Interfaces

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

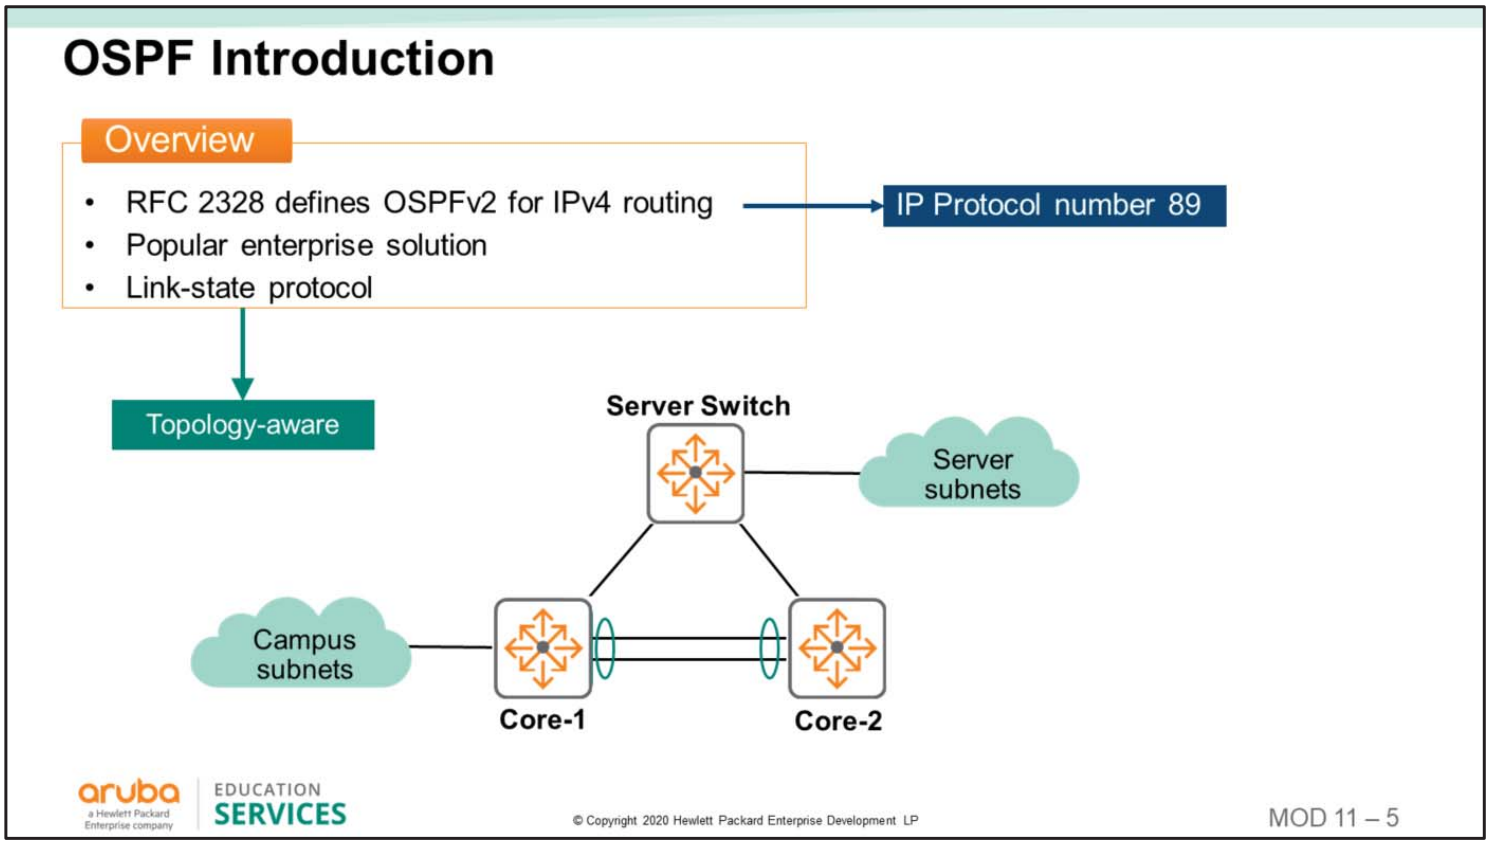

RFC 2328 defines OSPFv2 to route IPv4 packets. OSPF does not use TCP or UDP – OSPF advertisements are placed directly inside an IP packet. Therefore, it does not have a TCP or UDP port number, it has IP protocol number 89.

Note. The **IP protocol number** is a number associated to the protocol that is in use in Layer 4. The IP protocol number is included as a field in the Layer 3 header. This announcement in Layer 3 helps network devices to be aware of the Layer 4 protocol that is in used without decapsulating the packet. TCP uses the protocol number 6 and UDP uses the protocol number 17.

This is a very popular enterprise IGP routing solution, due to its hierarchical scalability and security mechanisms. For example, OSPF peers can authenticate packets exchanges.

As a Link-State protocol, OSPF enabled routers advertise information about their connected Layer 3 interfaces and networks (links) and the cost associated with that interface. Let’s see how it works.

Router ID Overview and Selection Criteria

Overview

- Unique, 32-bit address for each router
- Humans see “Core-1, Core-2, or Server Switch”
- Routers see 10.1.100.1, 10.1.100.2, 10.0.100.0

Determine RID

1. Manually configured.
2. Highest loopback IPv4 address
3. Highest non-loopback IPv4 address

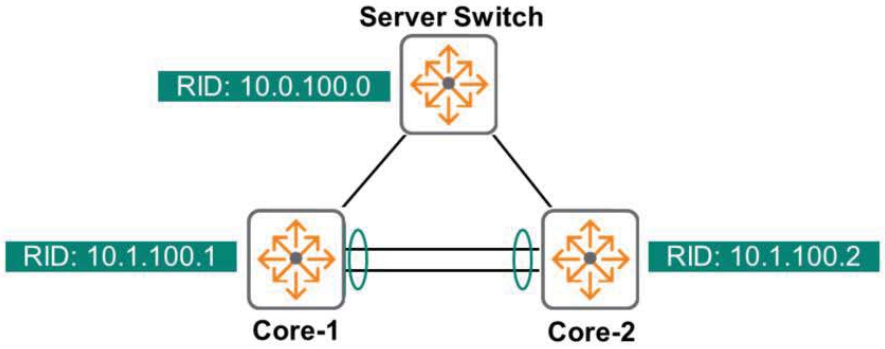

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 11 – 6

All OSPF routers require a unique Router Identifier (RID) - a 32-bit IP address in dotted-decimal notation. Routers include their ID in all OSPF packets that they send. Humans might look at the figure and see ”Core-1, Core-2 and Sever Switch”. The routers might identify themselves as 10.1.100.1, 10.1.100.2 and 10.0.100.0

You can depend on each router to automatically identify itself, or you can take control of the situation and manually assign RIDs. Automatic assignment is tempting, since it requires so little effort. However, many experienced engineers prefer to manually assign RIDs, due to certain documentation, troubleshooting, and management advantages.

AOS-CX uses the following sequence to determine the Router ID:

1. If you manually specify the RID, then that is what the router uses.
2. If you do not specify the RID, the loopback interface with the highest IP address becomes the RID.
3. If no loopback interfaces exist, the regular non-loopback interface with the highest IP address becomes the RID. Non-functional Interfaces that are in a down state are not considered.

A Loopback interface is defined as a logical interface that is always is in an up state (unless you manually disable it). This interface is useful for processes and protocols that depend on the interface status to work. The IP address associated to a loopback interface is routable, which means that external devices can initiate communication to it. As you gain more education and experience, you will continue to learn about the advantages of loopback interfaces – for scalability, troubleshooting, and network management.

OSPF General Operation Overview

Phase	Objective	Result
1. Neighboring	Build a neighbor relationship	Neighbor table
2. Topology database	Exchange of routing information	Link State Data Base
3. OSPF routing table	Run SPF to get the best path	OSPF routing table
4. Building routing table	OSPF process publish its results into the Routing Table (FIB)	Routing table entries

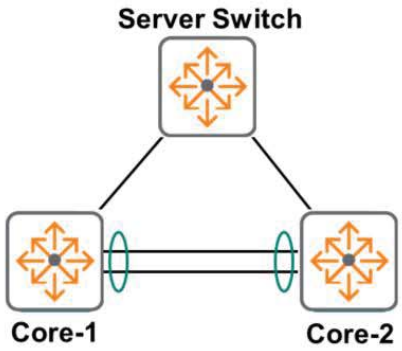

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 11 – 7

Remember that a router's job is to learn about every path, choose the best paths, and then route packets along those paths. To do this, OSPF uses a four-phased approach. Each phase results in a kind of table - built, maintained, and stored in local memory on each router. The more deeply you understand each phase and its associated table, the better you will be at designing, deploying, and troubleshooting OSPF networks. The four phases are shown in the figure, and listed below:

- **Phase 1: Build a Neighbor table** – Each OSPF router establishes a relationship with an OSPF neighbor. This information is available in the Neighbor table.
- **Phase 2: Build a topology database** - also called a Link State Database (LSDB), OSPF enabled routers exchange network prefixes and parameters. This information is stored in the LSDB table.
- **Phase 3: Build OSPF route table** – OSPF enabled routers run the SPF algorithm to get the best path for a particular destination. The result populates the OSPF routing table
- **Phase 4: Build the route table** – also known as a forwarding table or a Forwarding Information Base (FIB). In this phase the OSPF internal process publishes its results into the routing table.

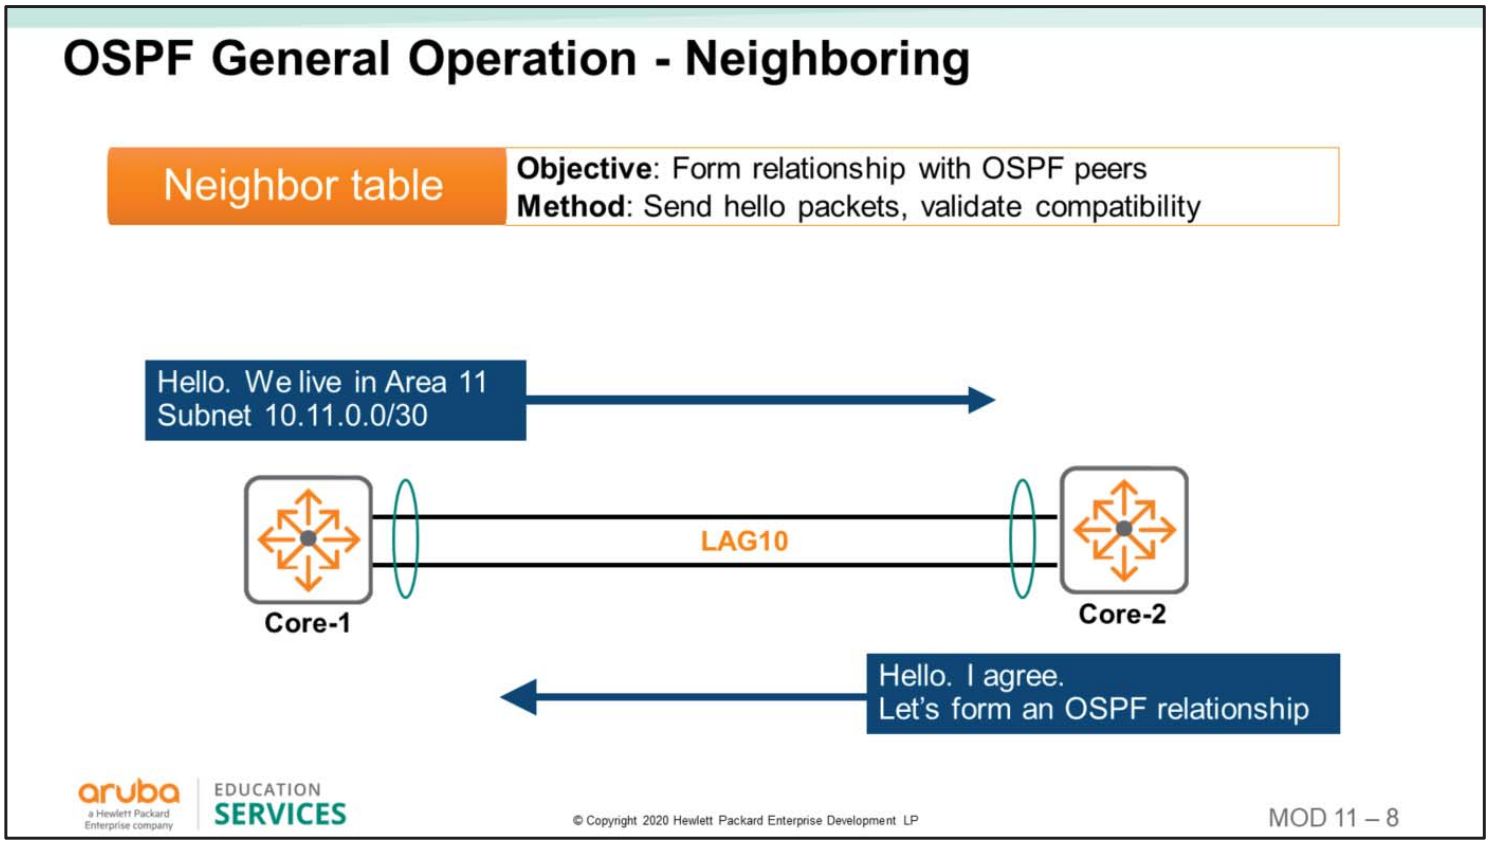

Phase 1: Build a Neighbor table

First, directly-connected OSPF neighbors introduce themselves to their directly connected neighbors. The objective is to validate that they are compatible. The method is to send OSPF "Hello" packets out each OSPF interface.

For example, Core-1 sends hellos out all of its interfaces where OSPF has been enabled shown in the figure, including out LAG10, on subnet 10.11.0.0/30, "Hello. I am Core-1 and I think that we live in Area 11, a normal area, and that we do not need to use secure authentication. We are connected on subnet 10.110.0/30". Core-2 sends similar information in its hello packet out LAG10. As long as the criteria match, the OSPF routers agree to be neighbors – they form an OSPF neighbor relationship. If any of these parameters do not match (typically due to misconfiguration), the routers refuse to form a neighbor relationship, and your network is broken. Let us assume that all routers in the figure successfully formed an adjacency with each directly-connected peer. This will be reflected in each router's neighbor table.

Note: Only a small part of Hello packet contents are shown and discussed here, to convey the main idea. There is much more information exchanged in hello packets, about which you will soon learn.

Note. The term adjacency refers to a couple of routers that are physically connected to each other. This term is different from OSPF neighboring since an OSPF neighbor relationship can be established between a couple of routers that are not physically connected.

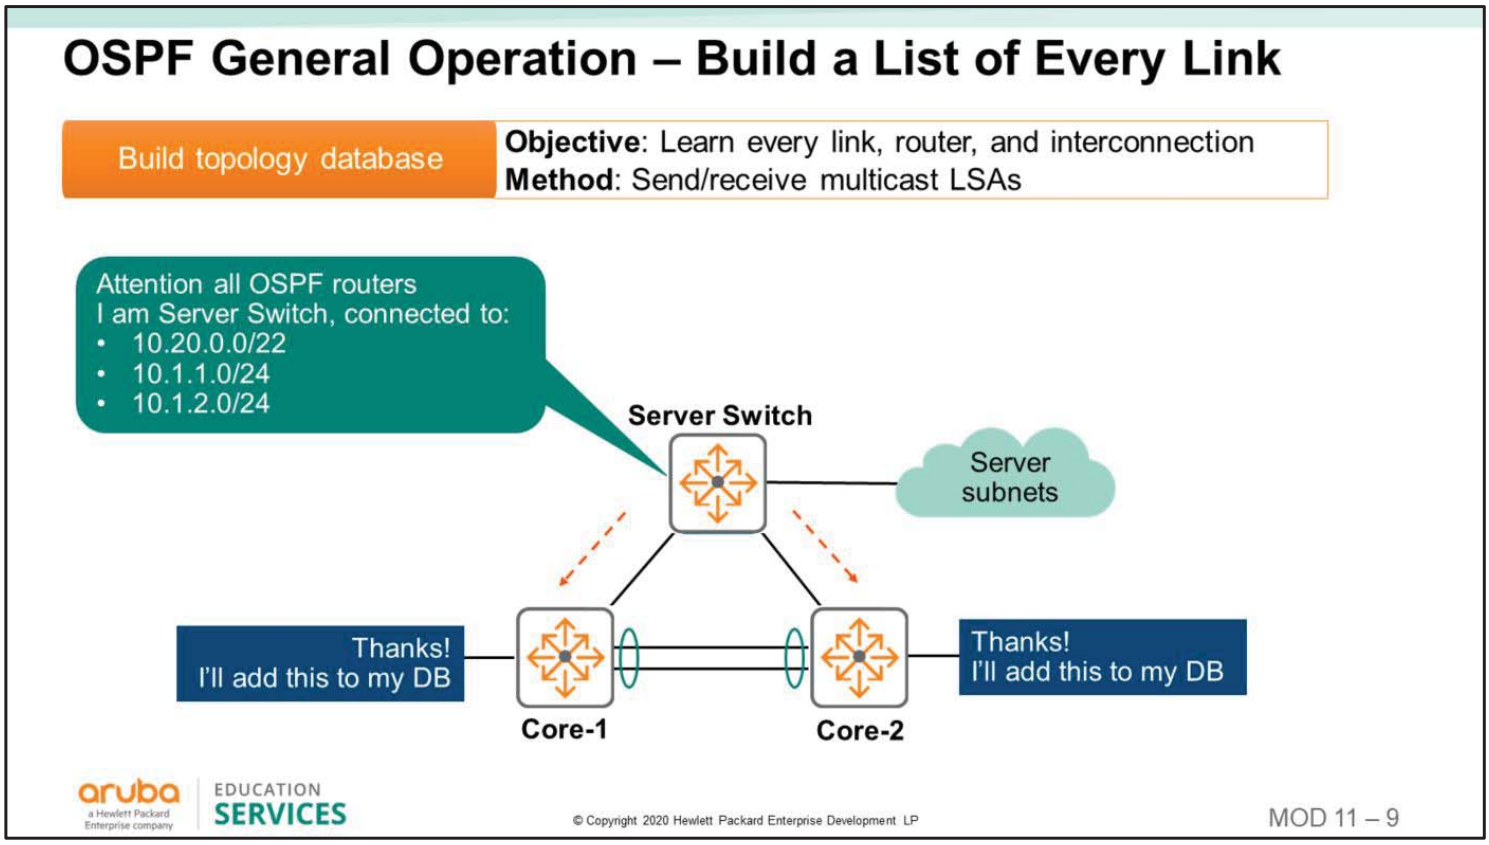

Phase 2: Build a topology database – also called a Link State Database (LSDB)
A router shares known subnets with other routers If and only if an adjacency is fully formed. The objective is to build a topology database – a list of every single link (subnet), every single router, and exactly how those routers and subnets are interconnected. The method is to send multicast Link State Advertisements (LSAs). Server Switch sends LSAs out all of its interfaces, "Attention all OSPF routers. I am Server Switch, and I am directly connected to 10.1.1.0/24, 10.1.2.0/24, 10.20.0.0/22. Every other OSPF router receives these LSAs and adds the information to its Topological database. Soon, every router has received LSAs from every other router. Thus, each router has a full topology database. Again, this database is essentially the diagram you see in the figure, in a numerical format. There is only one reality – the figure shown here in this example. Therefore, every OSPF router has an identical topological database.

Note: The figure only shows Server Switch advertising its directly-connected subnets, because that is all that it currently knows. However, once Server Switch receives LSAs from the other routers (as they have just received from Server Switch) it will advertise those routes as well. Routers advertise the entire contents of their topology database to all other routers. The topology database is also known as the link-state database.

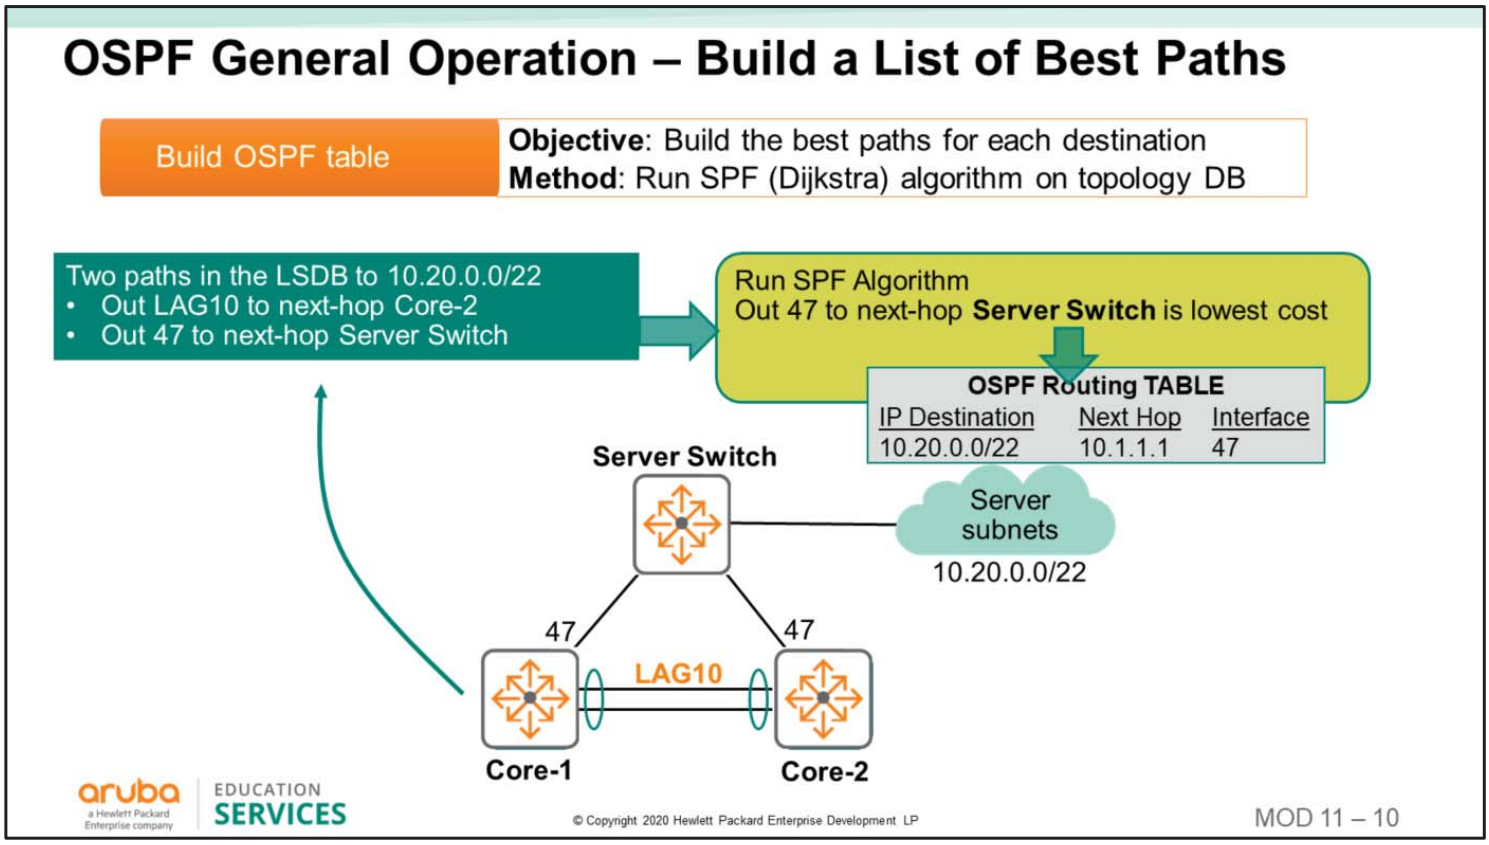

Phase 3: Build a route table

Although there is only one topology, each router's position and connectivity within that topology is unique. Thus, each router's objective is to determine the best path to each link from where they sit within the topology. The method is to run the SPF (Dijkstra) algorithm on the topology Database or LSDB.

For example, Core-1's SPF algorithm analyzes the topology database and sees that Core-2 could reach destination 10.20.0.0/22 by routing packets out port LAG10. And Server Switch can reach the same destination out of port 47. The algorithm considers the cost associated with each path (based on bandwidth) and chooses the lowest cost (fastest) path. This best path is added to the list of best paths – the route table.

In the example the best path is to send the traffic directly to Server Switch.

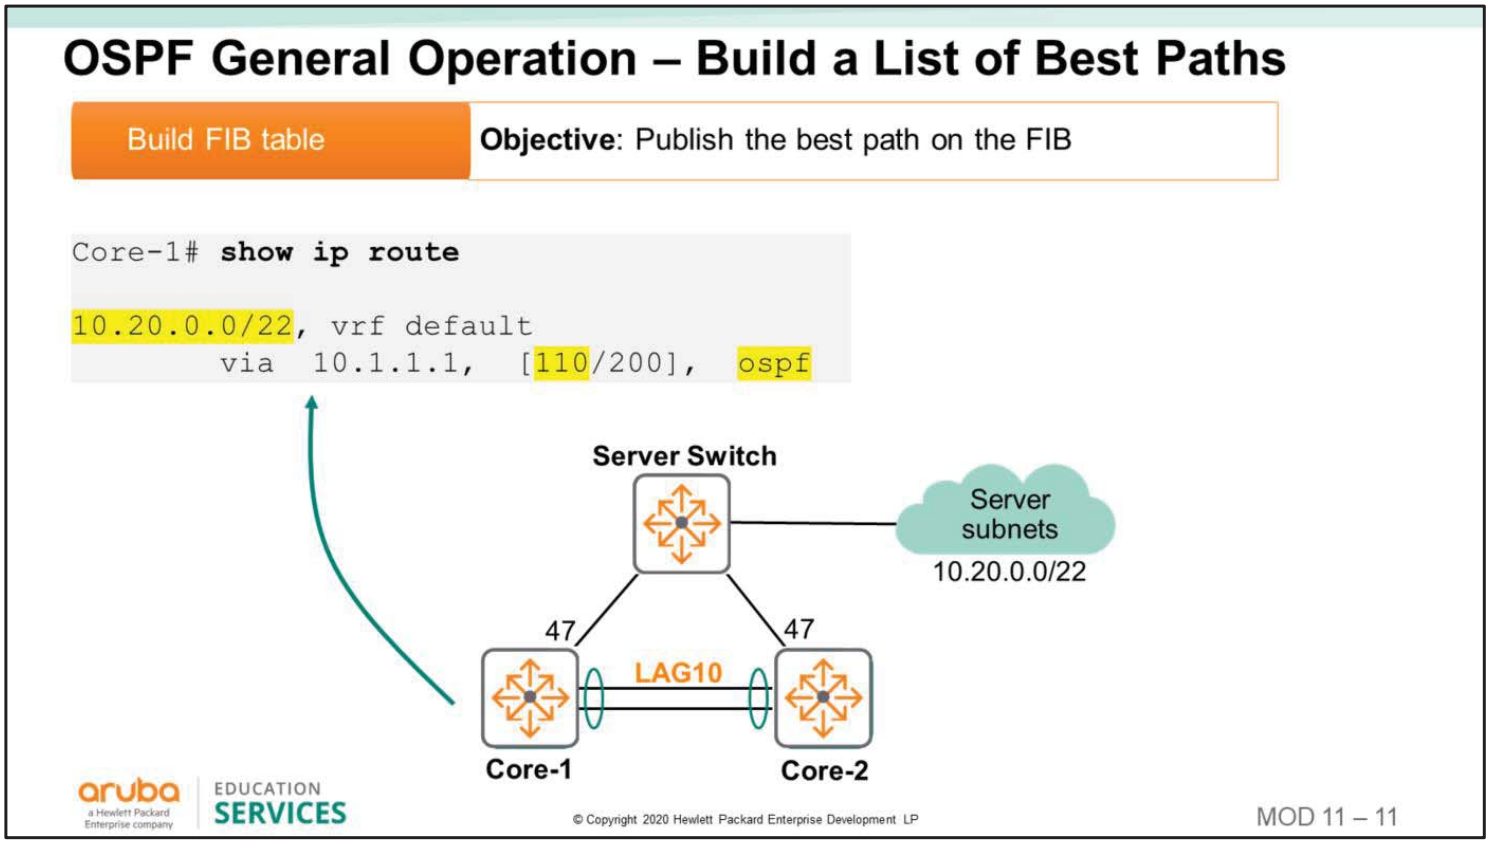

Phase 4: Build a route table

The best path previously calculated by OSPF is published on the Routing table (FIB). In this case the best path from Core-1 is to Server Subnets uses the path through Server Switch

OSPFv2

Neighboring

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

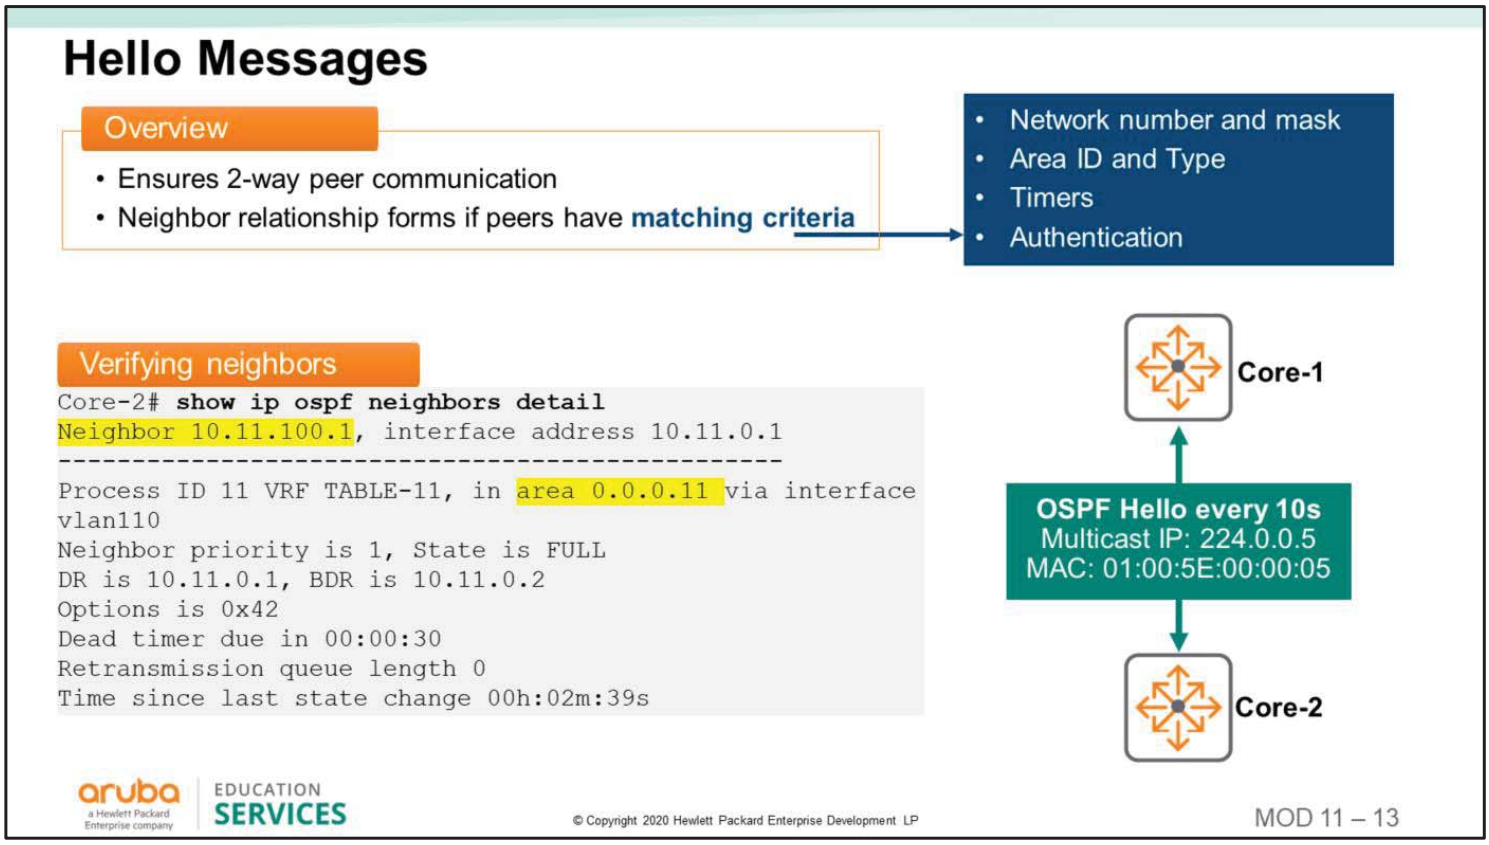

Directly-connected OSPF routers send hello packets to ensure 2-way communication, and to act as a failure detection mechanism. By default these packets are sent every 10 seconds to Multicast IP address 224.0.0.5. This is the reserved “attention all OSPF routers” multicast address. Its associated MAC address is 01:00:5E:00:00:05. Recall that all Layer 2 switches flood broadcast and multicast frames out all ports in the broadcast domain, so all directly connected routers exchange Hellos – the first packet to be sent in the OSPF exchange process.

Remember that another main purpose of Hello packets is to build and maintain a neighbor table. Peers only form a neighbor relationship if they are compatibly configured. They must agree that they are on the same subnet and have the same subnet mask. They must be in the same area and agree on the area type. Their timers must match (10-second OSPF hello timer, etc.), and they must be configured to use the same authentication type.

To verify the hello interval you can use the **show ip ospf neighbor detail** command.

In AOS-CX the hello interval is 10 seconds however you can customize the interval from 1 to 65553 seconds.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf hello-interval <seconds>**

Dead interval is the interval of time after which a neighbor is declared dead. Typically this value is 4 times the Hello interval. This value also must match to create a neighbor relationship.

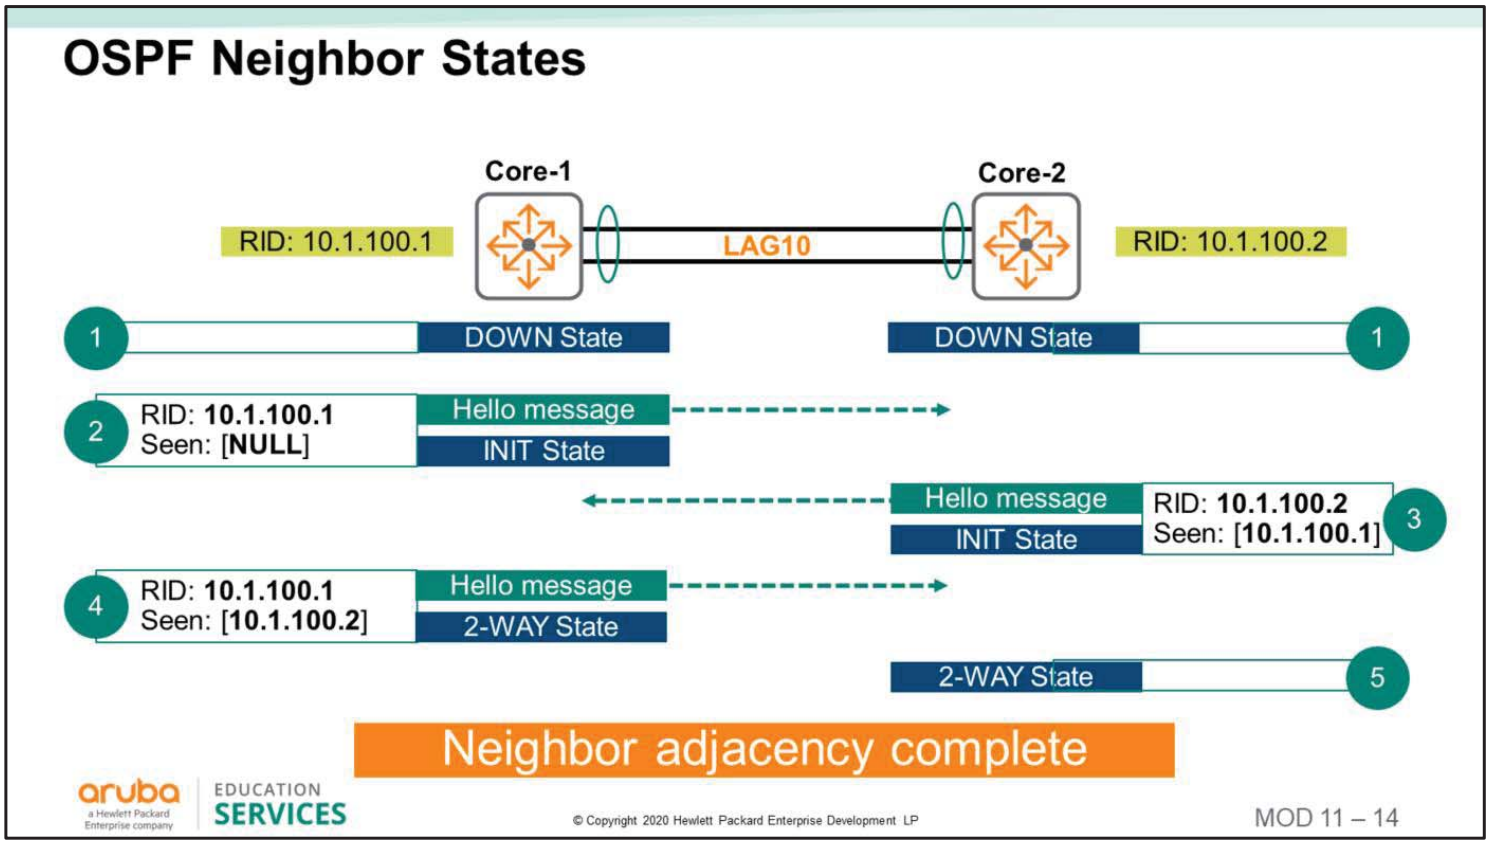

OSPF uses a Finite State Machine (FSM) to process the neighbor state transitions between routers, when certain conditions are satisfied. This process can be divided into two main phases:

- Establish Neighbor Adjacencies
- Synchronize OSPF databases

Is important to clarify that in a Broadcast Network, database synchronization only happens between the Designated Router (DR) and DROTHER routers and between the Backup Designated Router (BDR) and DROTHER routers. In other words, DROTHER routers only establish neighbor adjacencies and stay in the 2-WAY state (which is consider a stable state) but never synchronize databases.

As an example, consider that two core switches attempt to become OSPF neighbors over a Broadcast Network.

Establish Neighbor adjacencies:

1. Core-1 and Core-2 start with the DOWN neighbor state.
2. Core-1 transits to the INIT state when it sends the first Hello message. This message includes Router ID = 10.1.100.1; the Hello message includes a seen field – From which routers has Core-1 received hellos. Since this is the first Hello message the field has a NULL value.
3. Core-2 receives Core-1's Hello message and responds, indicating that the message has been seen and values are compatible. Core-2 transits to INIT state.
4. Core-1 receives the Hello message and transit to the 2-WAY state. This device sends

again the Hello message to Core-2 but this time it includes both Router IDs in the seen field.

5. Core-2 receives this message and moves to the 2-WAY state.

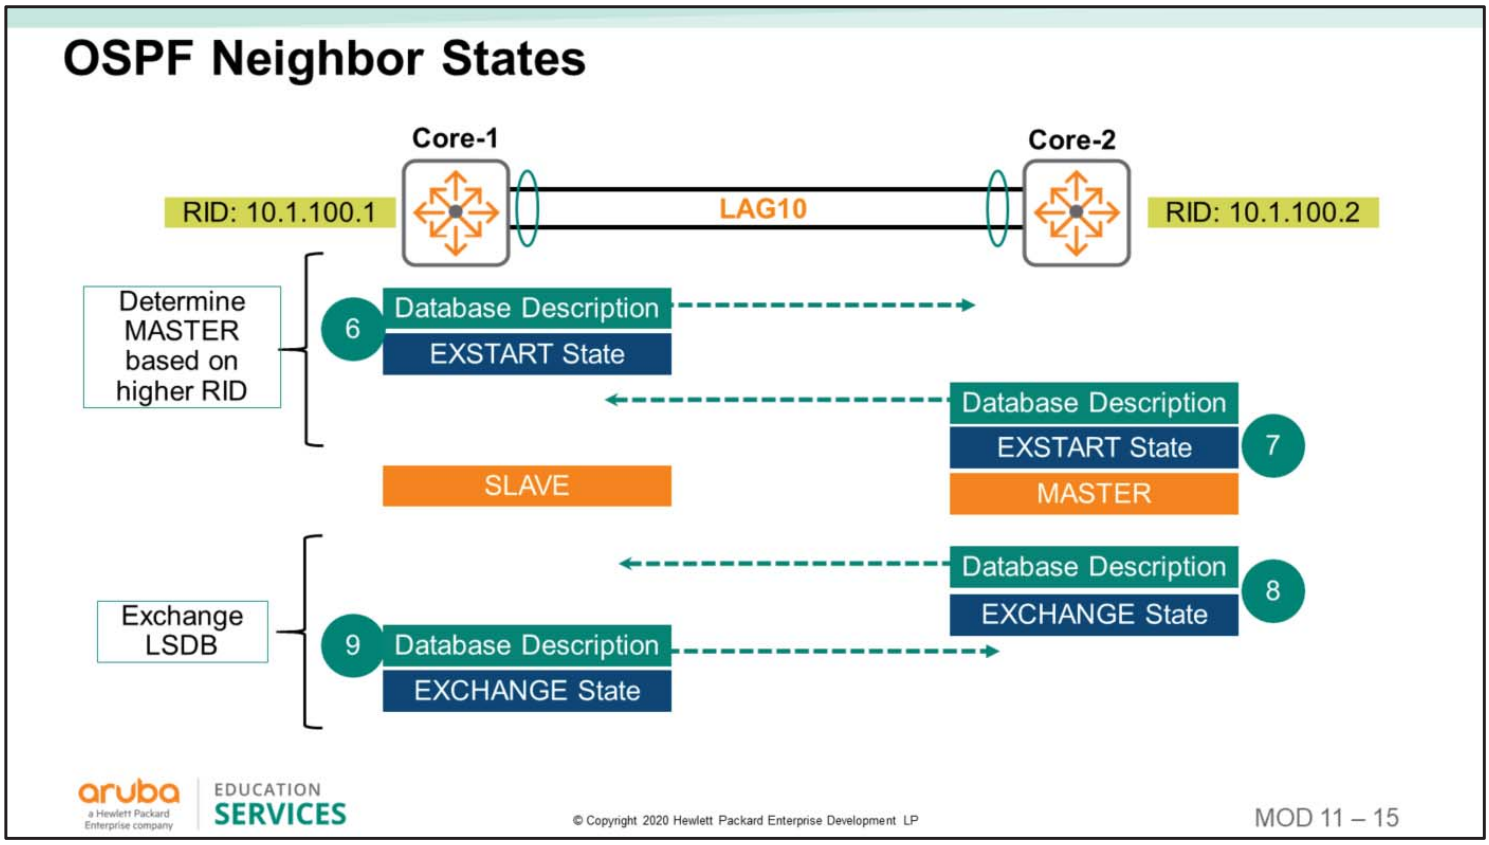

In the example, Core-1 and Core-2 are the only two devices in the network, and so they continue to the database synchronization process.

Synchronize OSPF Databases.

6. Core-1 initiates the synchronization process by sending a Database Description packet. The switch transits to the EXSTART neighbor state.
7. Core-2 does a similar process - it moves to the EXSTART state by sending a Database Description packet.

The goal of the EXSTART state is to determine which switch will become the MASTER on the link, based on the highest Router ID. Understand that the MASTER role defines only which switch will initiate the Database exchange process. This this role has nothing to do with DR/BDR – it is only for Link State Database (LSDB) exchange. In this example, Core-2 has a higher Router-ID and becomes the MASTER.

8. Core-2 sends another Database Description packet as it transitions to the EXCHANGE state. Core-2 is now sharing a summary of the contents of its LSDB with Core-1
9. Core-1 also sends a Database Description packet and moves to the EXCHANGE state, sharing its LSDB with Core-2.

After several packets sent and received in the EXCHANGE state, each router will have a summary of the others entire LSDB. They compare the information received and, in the next step, request the missing information.

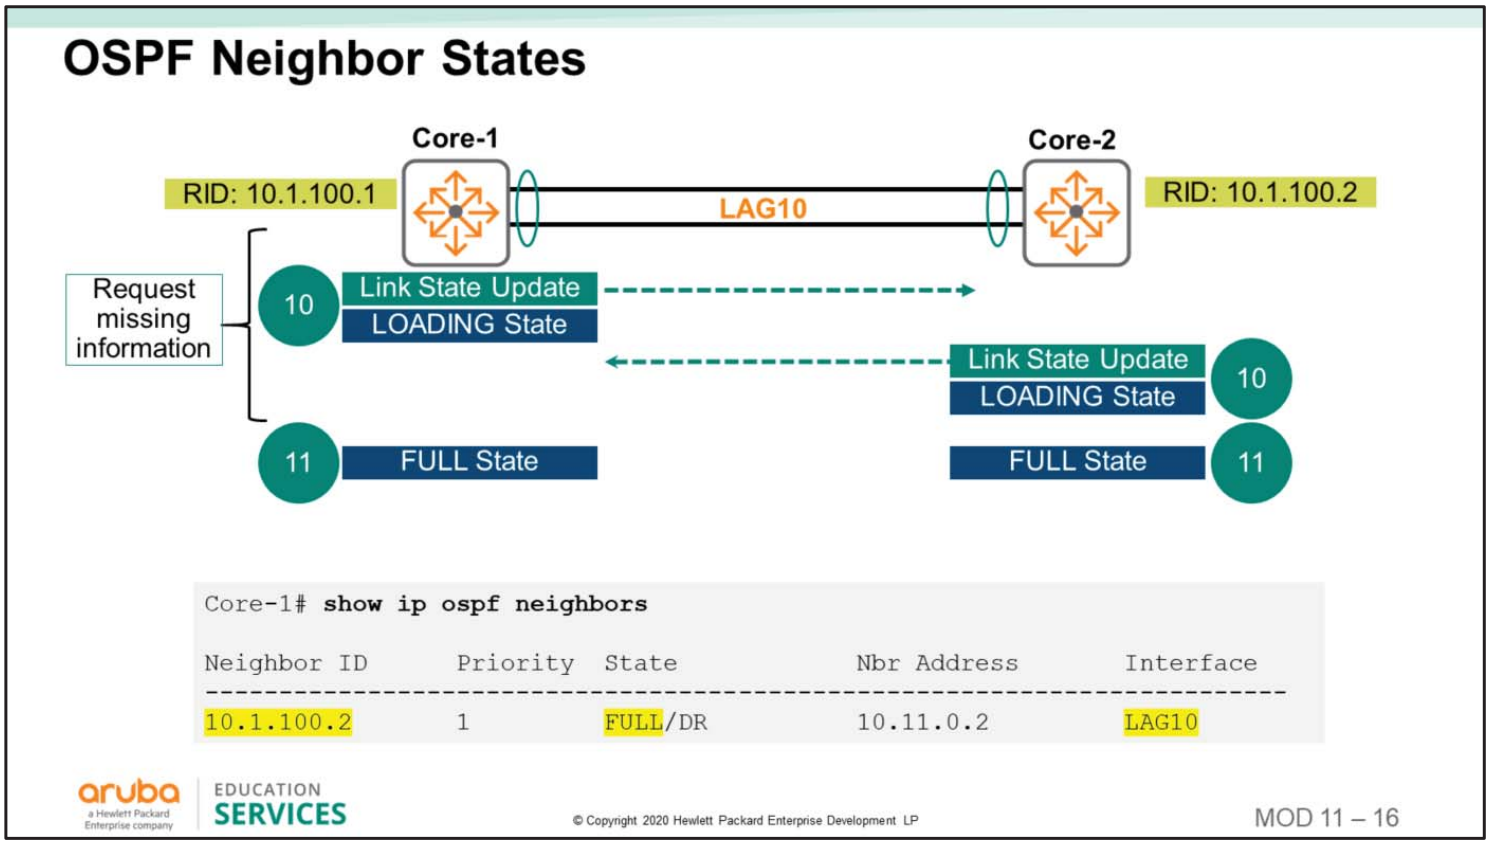

10. Core-1 and Core-2 request the missing information using Link State Update (LSU) packets. The other peer answers by sending the requested database information. Both switches transition to the LOADING neighbor state.

This state is the actual routing information exchange.

11. Core-1 and Core-2 move to the FULL state when there is no more information to be exchanged, and both devices have the same Link-State Database (LSDB).

You can verify the OSPF neighbor state using the show ip ospf neighbors command.

OSPFv2

Operation

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

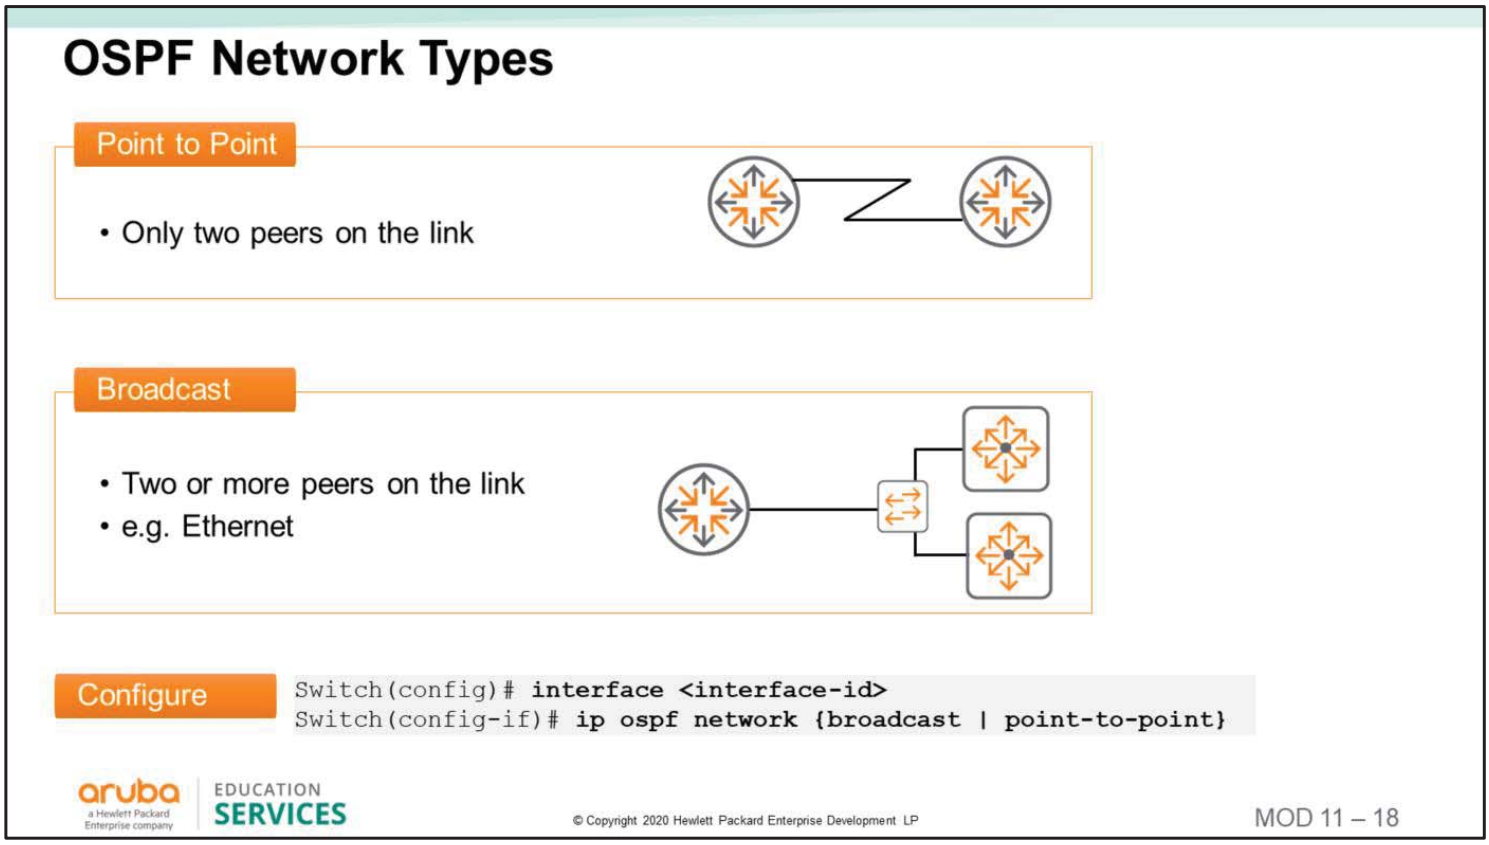

From the protocol perspective two options exists for network types:

- **Point-to-Point Network:** Only two peers are on the link. When an interface is configured to be part of a Point-to-point link, OSPF knows that a single neighbor device is expected on the interface. PPP Serial links (not used too much in networks now) is an example of this network type.
- **Broadcast Networks:** Two or more peers might be on the link. When an interface is configured to be part of a broadcast network, OSPF knows that more than one neighbor might be discovered on the interface.

In ArubaOS-CX, interfaces default to the broadcast network type, but you can modify this with the **ip ospf network** command.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf network {broadcast | point-to-point}**

To verify the type of network that is in use you can use the **show ip ospf interface** command.

Unless you truly have multiple routers in the same broadcast domain, it is a recommended to configure switch-to-switch links with the point-to-point network type.

Note. The network type is a concept that was used many years ago when Ethernet was not fully accepted as the unique Layer 2 protocol. Back in those days, serial communication was used more often, where it is only possible to have one device at each end of the link. This led to the concept of point-to-point network type. Meanwhile, you know that Ethernet is

designed to have multiple devices connected to the same broadcast domain. This led to the concept of a Broadcast Network type.

Note. OSPF can also use the Non-Broadcast Multiaccess (NBMA) Network type. This network can support multiple devices (multi-access) but does not support the broadcast capability. Frame Relay is an example of this network type. The use of this type of network is deprecated for modern networks, and so will not be covered in this course.

Broadcast Network scalable problem

Scalable problem

- Each peer exchange information
- Link State database can easily increase
- Decrease the performance of the network

Solution

- Exchange LSDB information with a unique Designated Router

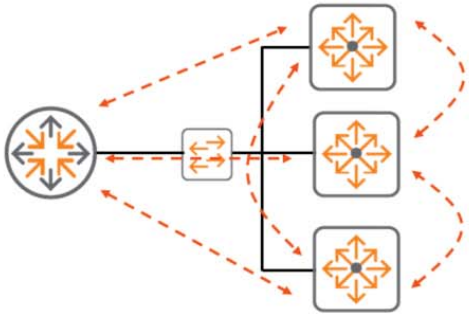

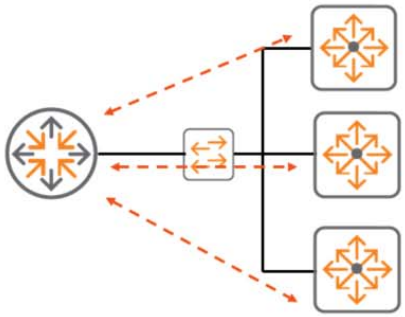

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 11 – 19

The amount of information that can be exchanged between OSPF peers could be high in a medium-sized enterprise network. In a Broadcast network type this could easily increase, since each router can have multiple peers. This could impact router performance when hundreds or thousands of routes must be computed for each OSPF peer. We need a solution.

OSPF solves this scalability challenge by electing a Designated Router (DR) in the broadcast domain. This device maintains a complete full neighbor state with the rest of the devices (which implies that databases are exchanged with the peers). However, the non-designated routers do not exchange database information with each other. This helps to reduce the amount of information that each router in the domain must process.

To maintain high availability, you can elect a Backup Designated Router (BDR) to avoid a single point of failure. This device also maintains a full state with all devices in the Broadcast network. However, it only advertises information when the primary DR is no longer available.

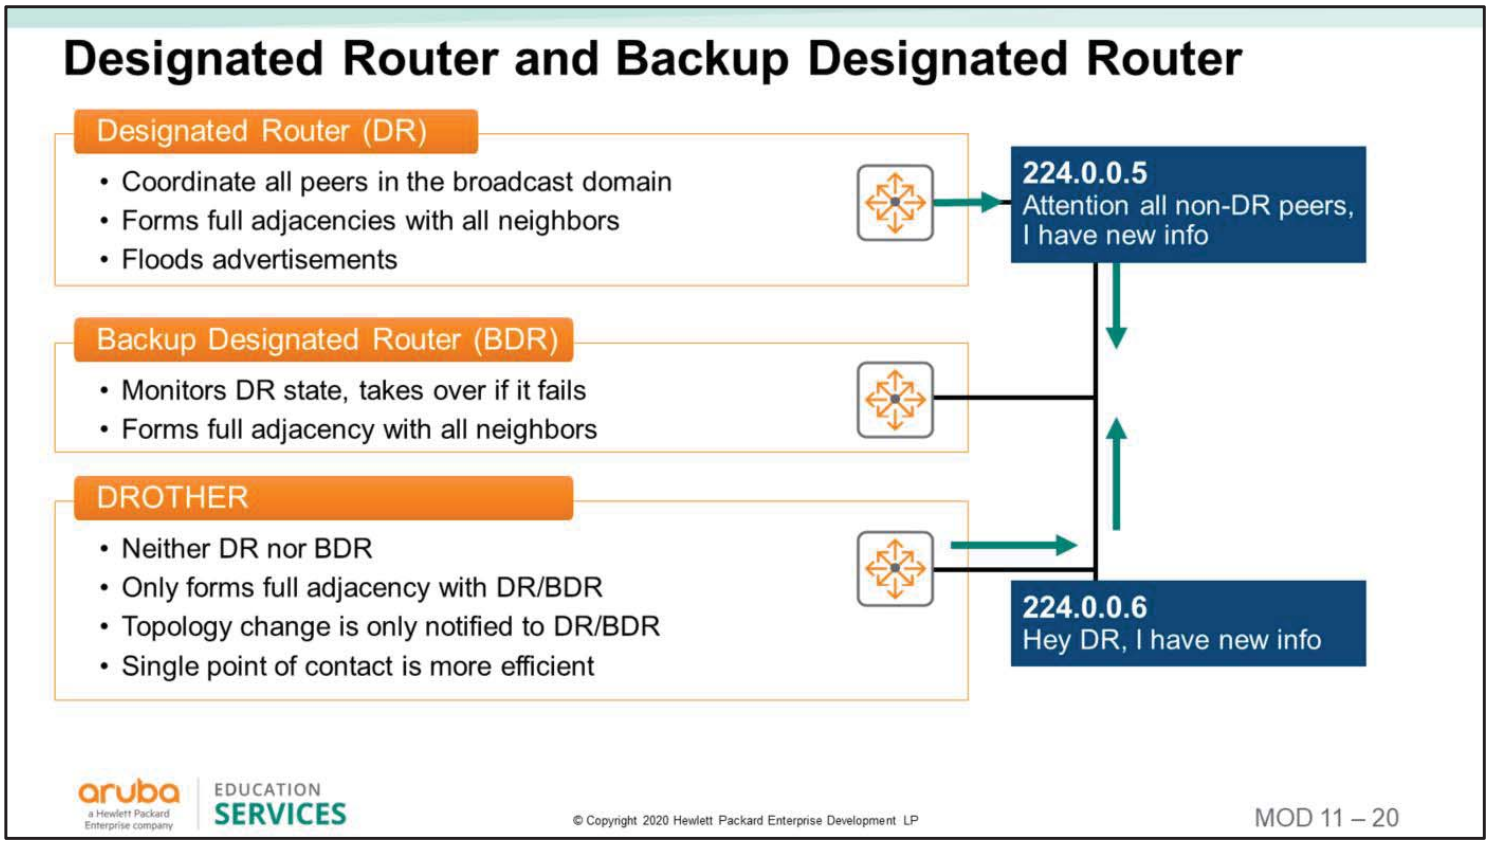

In a Broadcast network OSPF routers expect to have form neighbor relationships with more than one peer. This situation could potentially be inefficient – when there is a topology change, LSAs can be flooded and re-flooded in the broadcast domain. The solution is to elect a single point of contact – a Designated Router (DR), which helps to coordinate peer interactions in the broadcast domain. It forms full adjacencies with all other OSPF peers in the broadcast domain, so that it can receive and flood LSAs.

A Backup Designated Router (BDR) is elected to take over if the DR fails. The BDR also forms a full adjacency with the DR, and all other peers in the broadcast domain.

Any router not elected to be the DR or BDR is labeled as DROTHER, “Some router other than me is the DR”. These routers only form a full adjacency with the DR and the BDR. This is where the efficiency is realized. Whether there are 3 routers or 30 routers in the broadcast domain, DROTHERs need only form adjacencies with 2 peers – the DR and BDR.

In the figure, one of the DROTHERs detects a topology change on one of its other interfaces. It need not communicate to every other router in the broadcast domain, only the DR and BDR. So it sends a multicast LSA using destination IP address 224.0.0.6 – the reserved address that says, “attention DR and BDR”.

The DR is functional, so it receives this packet. It updates its topology database and informs all other DRs in the link with a multicast LSA to 224.0.0.5, “Attention all DROTHERs, I have new information”.

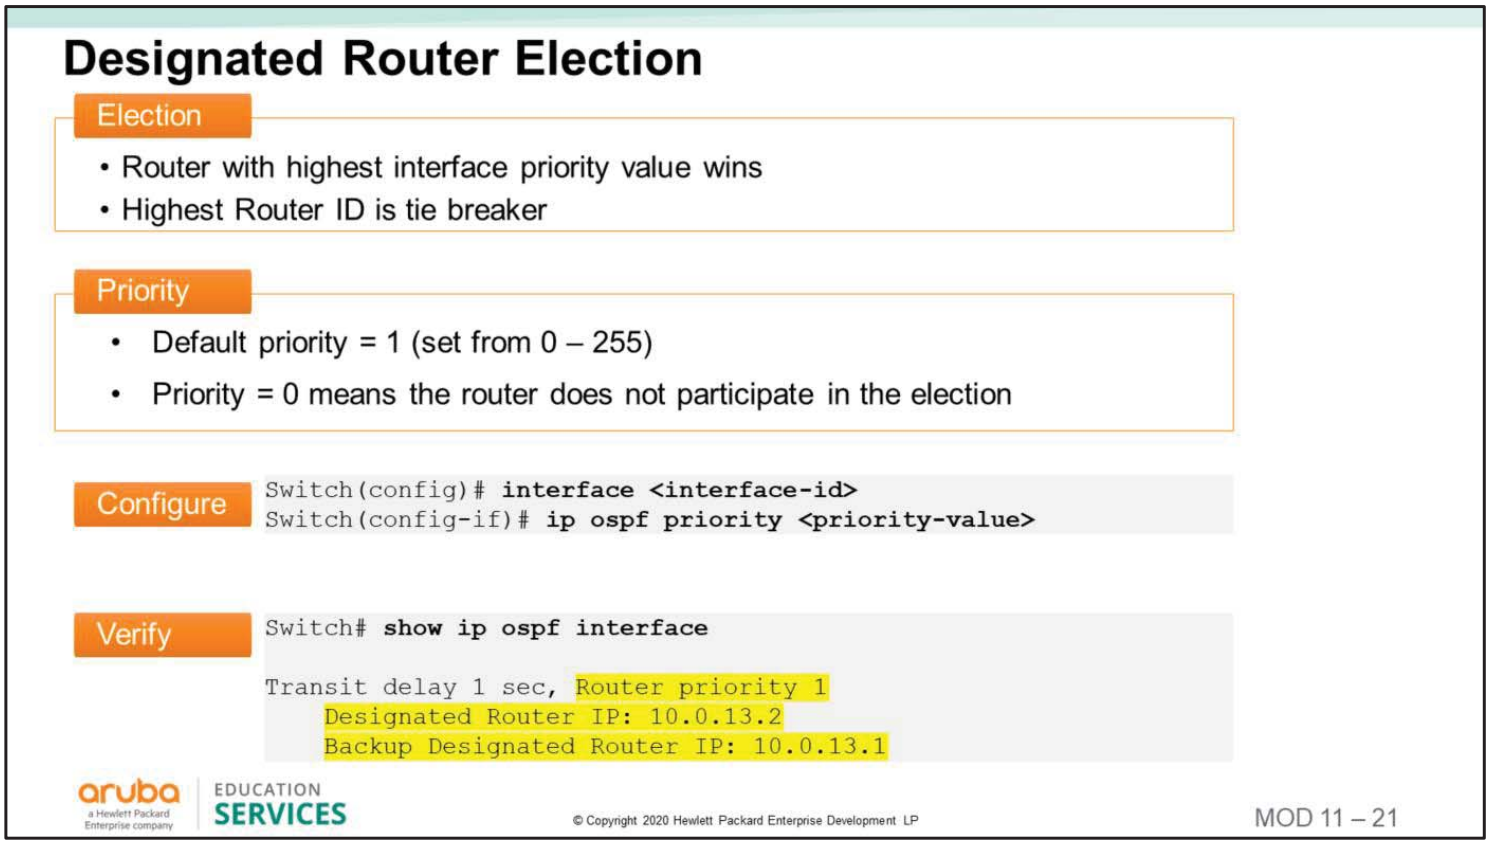

DR and BDR election is based on a priority value assigned to an interface - the highest priority value wins the election. In case of a tie, the router with the higher Router-ID becomes the DR.

AOS-CX follows certain rules related to the priority value:
- The default priority value is 1.
- Valid priority values are 0-255.
- A priority value of 0 means that the router does not participate in DR election.

As shown in the figure, you perform this configuration at the interface level.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf priority <priority-value>**

To verify the priority value and the elected DR, use the command **show ip ospf interface,** or **show ip ospf neighbors**.

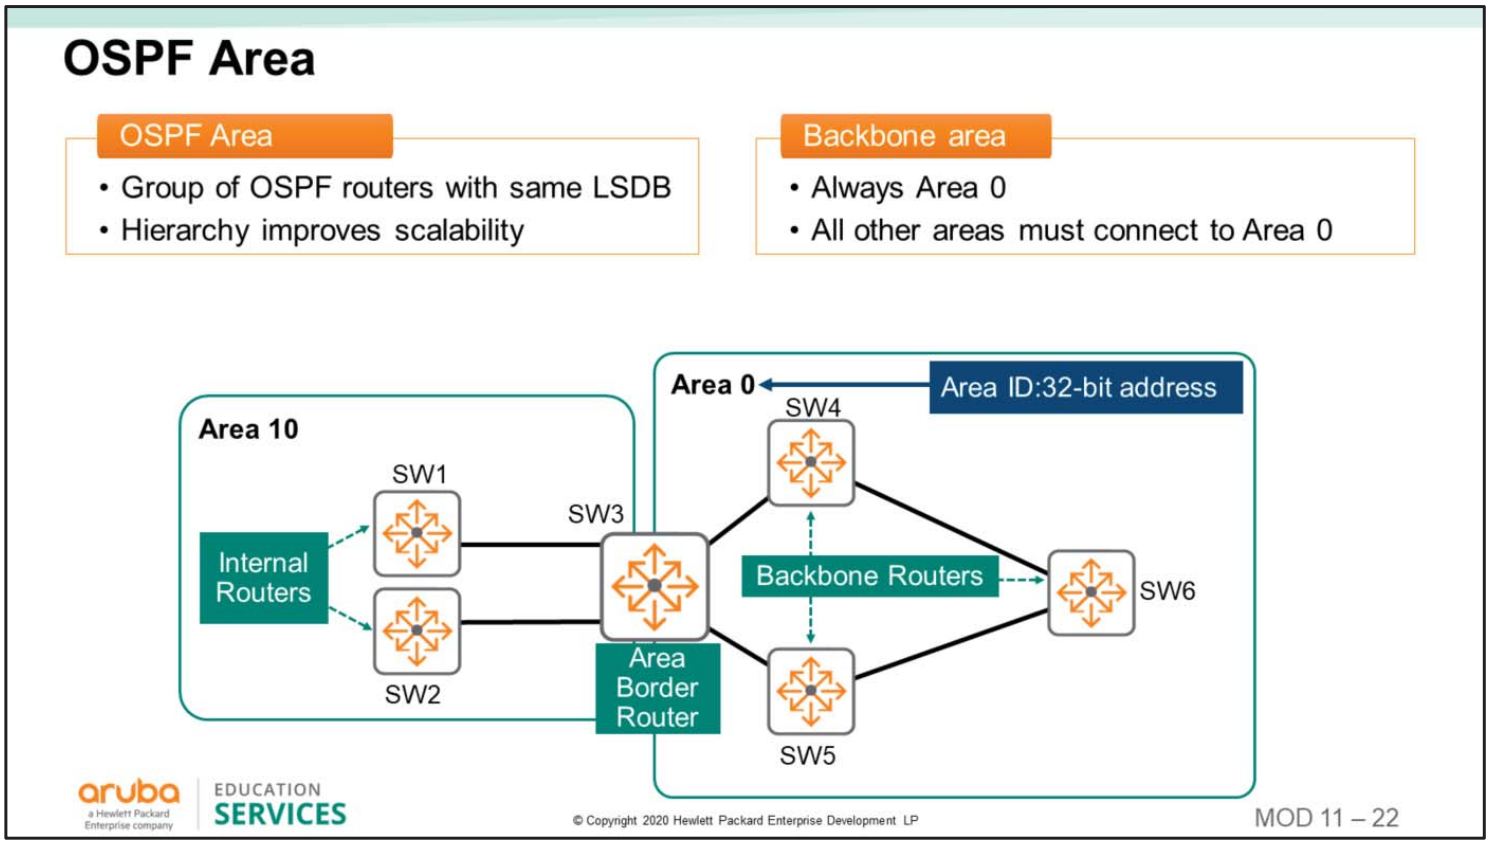

An Area is a group of OSPF routers that share the same Link State database. All router must be part of an area.
When you split a large network into separate areas, you reduce the size of the LSDB in each router, lower CPU utilization, and increase overall network stability. This is because each router in an area must only maintain the topology for that area.

For example, SW1, SW2, and SW3 need not learn about the entire topology – they only needs to know about Area 10 topology.

Routers SW1 – 2 are called “Internal routers” – all their interfaces are in a single area. If an internal router must route outside of its area, it simply forwards packets to its Area Border Router (ABR). An ABR is a router connected to two or more areas. SW3 is the ABR for Area 10, with two interfaces in Area 10 and two interfaces in Area 0.

By definition, ALL areas must connect directly to the special backbone area – Area 0. This is the most important area in the hierarch, and so should have a redundant design. So, Area 10 cannot connect directly to any other area that is not 0. Communication between 2 non-backbone areas MUST communicate via the backbone area 0. Routers with interfaces in Area 0 are called backbone routers. Routes SW4, 5, and 6 have all interfaces in a single area, so they are also Internal routers. You can call them Internal Backbone routers.

You assign interfaces to an area by assigning them to an area ID - a 32-bit long address that can be written in dot-decimal notation (Area 0.0.0.0), or in decimal notation – Area 0, as shown in the figure. AOS-CX supports both notations.

This course is only focused on single-area designs, in which all router interfaces are in Area 0. You can learn more about hierarchical OSPF solutions in the advanced courses.

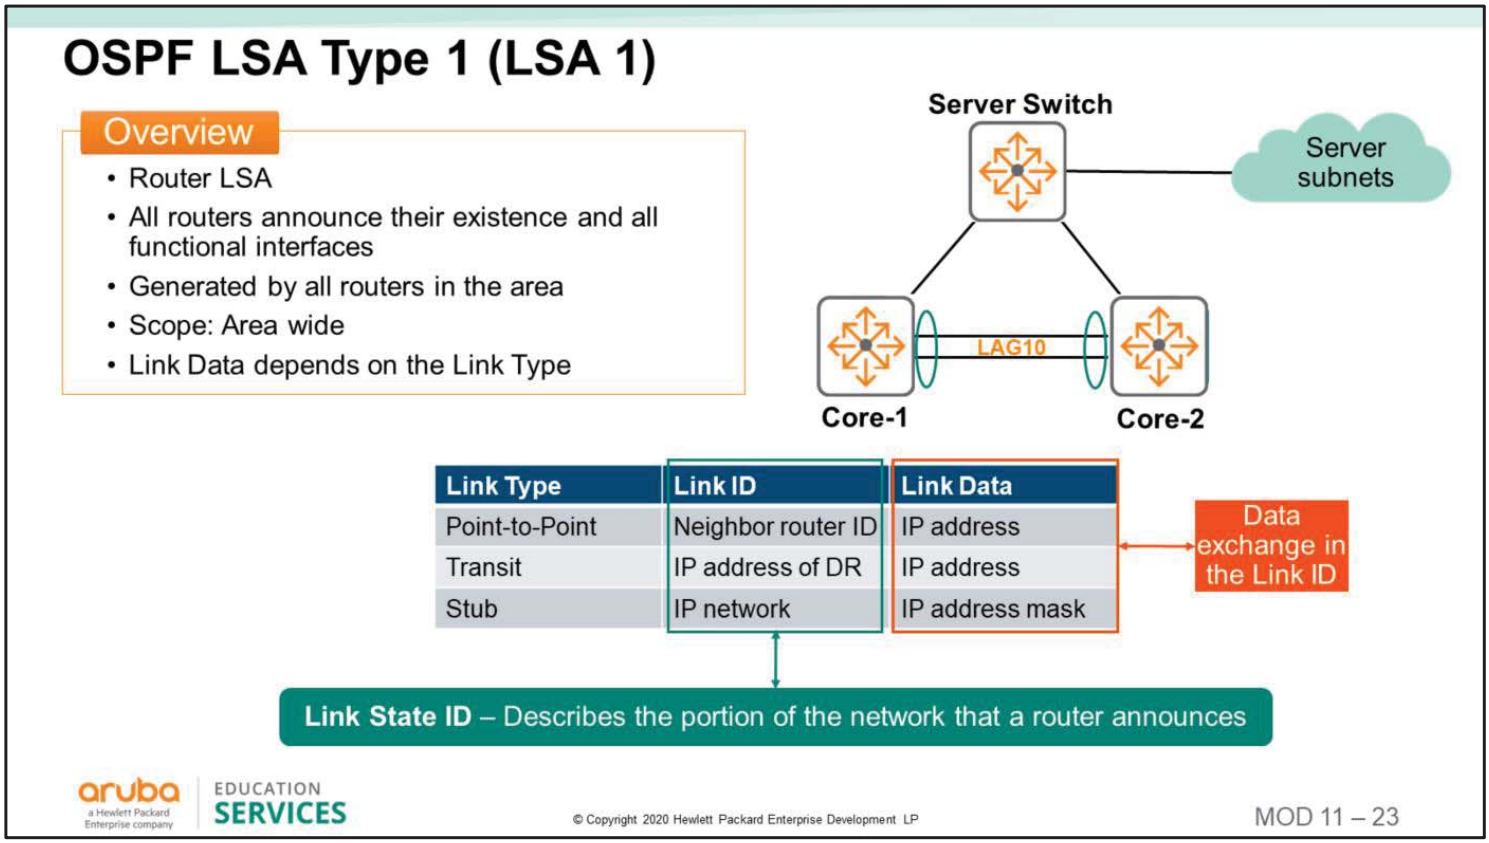

OSPF routers generate different types of Link State Advertisements (LSAs), each with a different purpose and scope in the OSPF routing domain. Each LSA includes the Link-State identifier which describes the portion of the network that a router announces. You will learn about these LSA types, starting with LSA Type 1.

The purpose of LSA Type 1 is for a router to announce itself. That is why they are called “Router LSAs”. You might walk into a room and say, “Hello everyone, my name is James Bond”. Similarly, routers uses a Type 1 LSA to say, “Hello everyone, I am RID 10.1.100.1, and I have 3 interfaces that are functional, and are participating in this OSPF area”.

Is important to remember that the information shared on LSA Type1 depends on the Link Type. Consider what information will be shared by the Server Switch when different link types are configured.

In addition to defining the network type, OSPF also defines the link type, which are two different concepts. Link types are primarily used to describe the interfaces or neighbors of an OSPF router.

- **Stub Link**. Used when OSPF is enabled on an interface and no OSPF neighbor exists on the interface. For example, a Loopback interface is considered stub Network.
- **Transit Link**. Used in a Broadcast network with 2 or more OSPF neighbors.
- **Point-to-point Link**. Used in Point-to-point networks, only one OSPF neighbor is expected in the link.

Note. The Link Type is a consequence of set up the network-id and the number of neighbors that are in the link. When the network Type is set to point-to-point then the router expect a

single device on the link. When the network type is set to Broadcast then the stub or transit link type is used. Stub for no neighbor and Transit for 1 or more neighbors.

In AOS-CX you can verify this information by using the **show ip ospf lsdb** command. In the example that router Core-1 has learned that there are three routers in the area. This can be a powerful troubleshooting command. If your network diagram shows that there are five routers in area 0, but you only see 3 listed, you know that there is a problem.

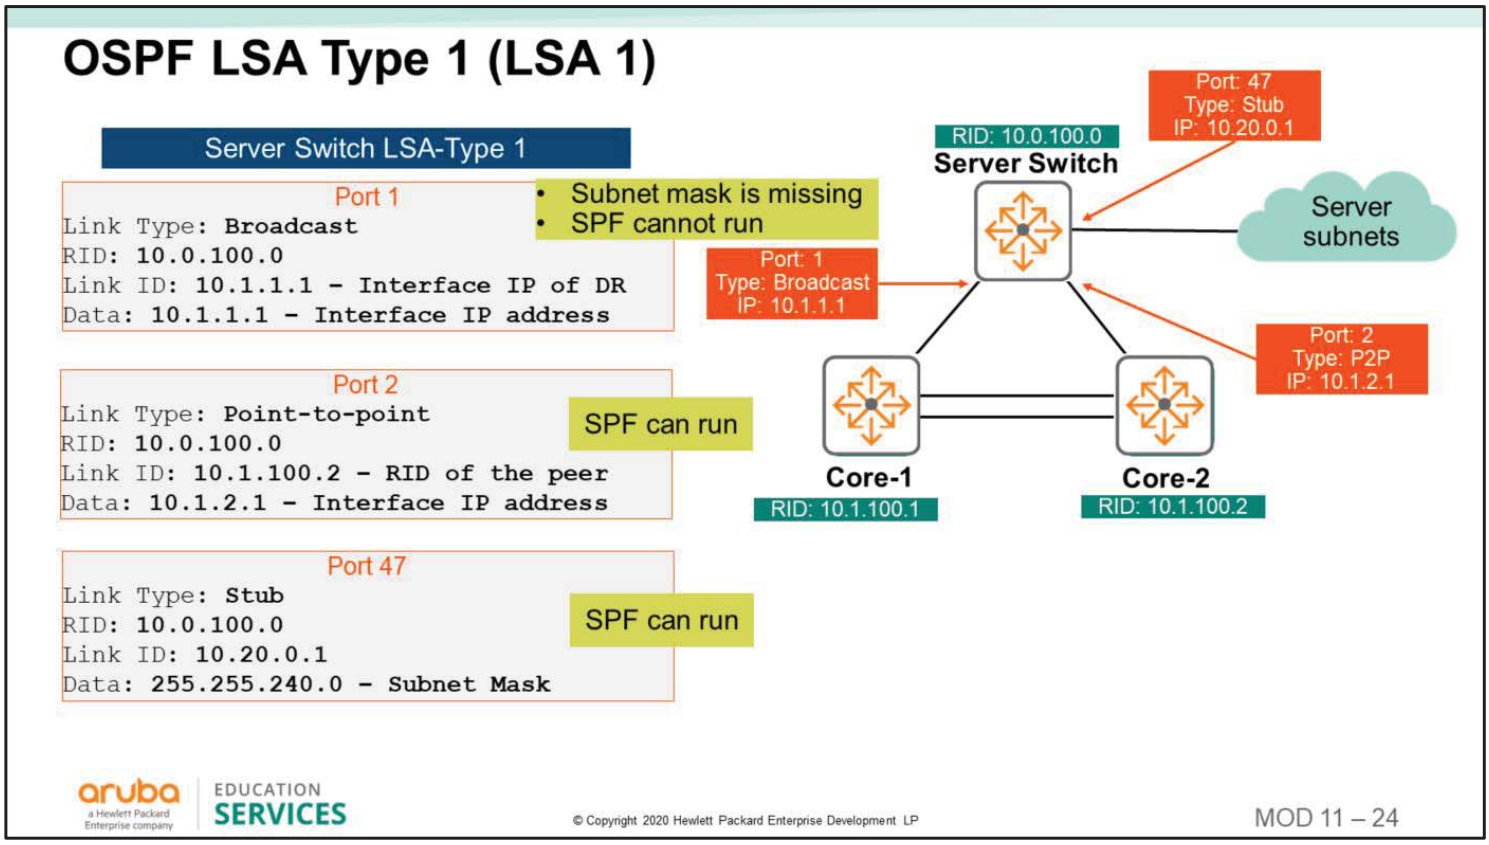

Let's analyze the information shared by the Server Switch's LSA-Type1 advertisements when different link types are configured.

Port 1:
- Link Type: Broadcast
- RID: 10.0.100.0
- Link ID: 10.1.1.1 – Interface of DR
- Data: 10.1.1.1 – Interface IP address

Port 2:

- Link Type: Point-to-point
- RID: 10.0.100.0
- Link ID: 10.1.100.2 – RID of the peer
- Data: 10.1.2.1 – Interface IP address

Port 47:
- Link Type: Stub
- RID: 10.0.100.0
- Link ID: 10.20.0.1
- Data: 255.255.240.0 – Subnet mask

Now analyze the information received from Core-1 and Core-2 perspective.

First, the Stub Link type includes the Subnet and the Mask, this information is enough to run

the SPF algorithm locally to reach the destination.

Second, the Point-to-point interface includes the Router ID of the peer and the local interface used. From Core-1 perspective the link details is not complete, Core-1 needs the data of the other peer on the link. This information will be known when Core-1 receives the LSA Type 1 from Core-2. When we receive the information from both parties then Core-1 can run the SPF algorithm.

Finally the broadcast interface will include only the IP address of the Router ID, but in this case no subnet information will be included using LSA Type1. How can we solve this problem?

The answer is to use the LSA Type 2

In AOS-CX you can verify this information by using the **show ip ospf lsdb** command.

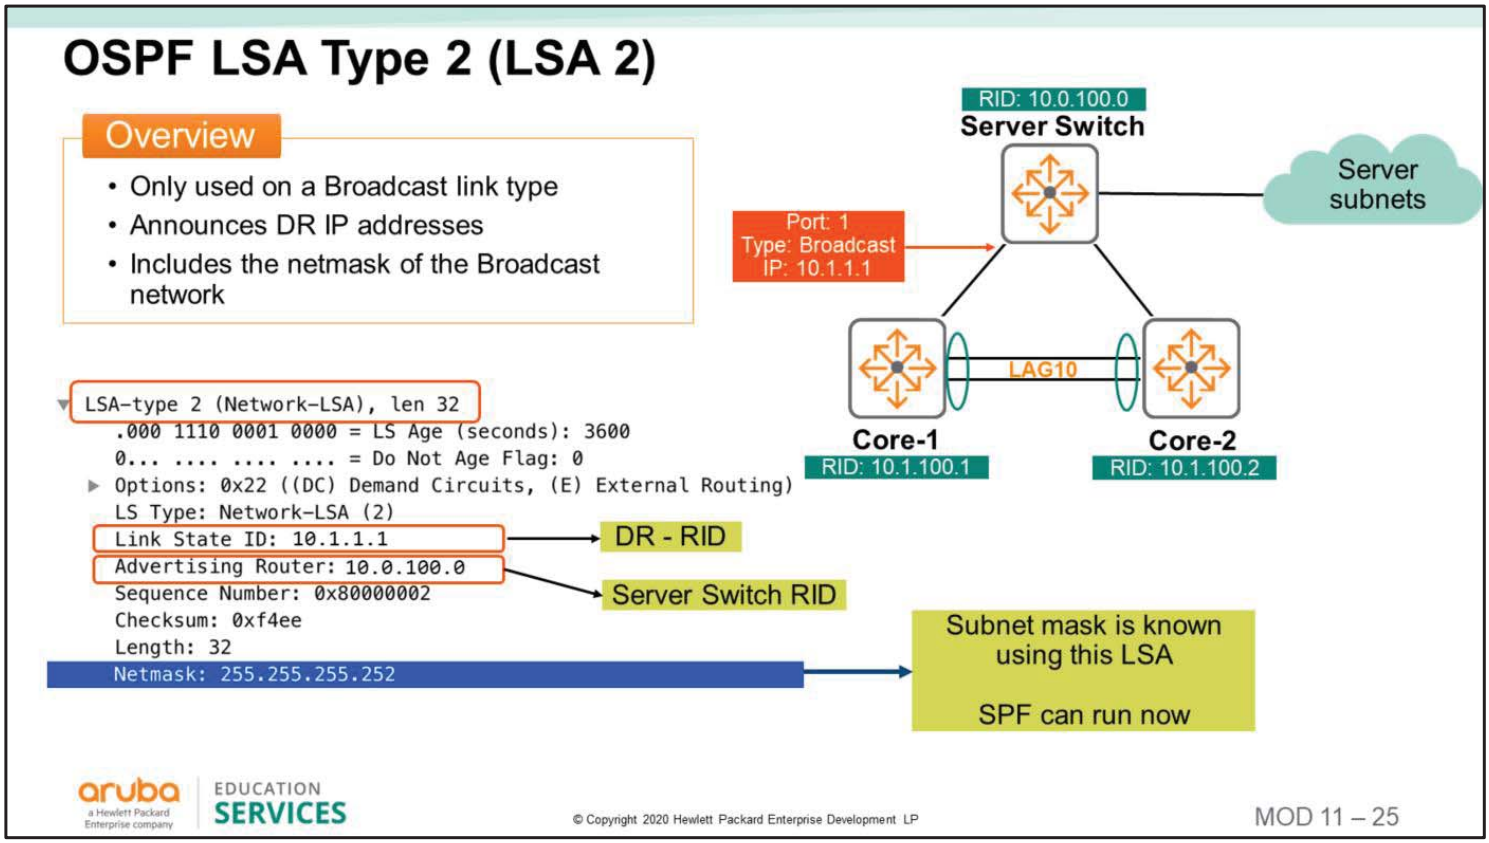

The LSA Type 2 is used when there is a Broadcast network type link, which have an elected DR and BDR. Routers on point-to-point network types do not generate a Type 2 LSA. In this case the Link-state ID is the IP address of the Designated Router. These LSAs are known as Network LSAs, because routers advertise the networks for which it is the DR.

The image shows a packet capture using Wireshark where the netmask is included in the LSA. This information helps to run the SPF algorithm. The missing information that was missing in Type 1 is now known.

In the **show ip ospf lsdb** output, the Network Link State Advertisement section shows you a list of DRs.

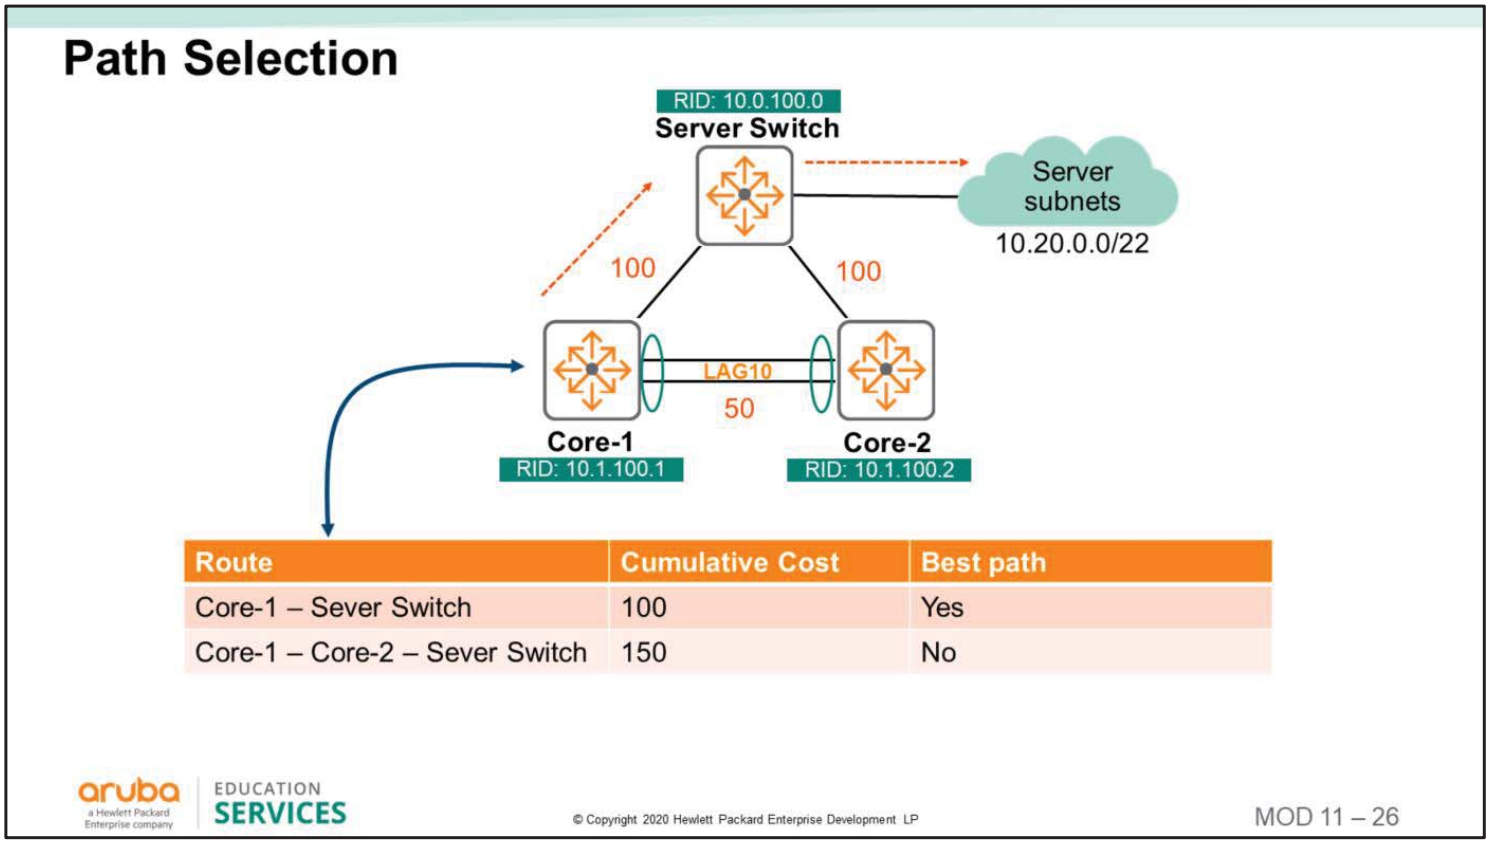

After all routers have successfully exchange LSAs and LSUs, they all have an identical LSDB – a topology database. Remember, the LSDB is a list of very link, every router, and how those routers and links are interconnected. It is a list of **every path**.

Now routers must run the Shortest Path First (SPF) algorithm – Dijkstra's algorithm – to find the **best paths** to each destination subnet. The best path is the one with the lowest cost, and cost is based on bandwidth. Therefore, if there are multiple paths to a single subnet, OSPF chooses the past with the lowest cumulative cost – the fastest path.

Consider the topology shown in the figure. Core-1 has two paths to reach destination subnet 10.20.0.0/22. To determine each paths's cost, simply add the values indicated for each link. Here's what you get:

- Using Core-2 as the next hop, the cost is 100 = 100.
- Using Server-Switch as the next hop, the cost is 100+50 = 150

Obviously, using Server-Switch results in the lowest cost, and so that path is added to the route table.

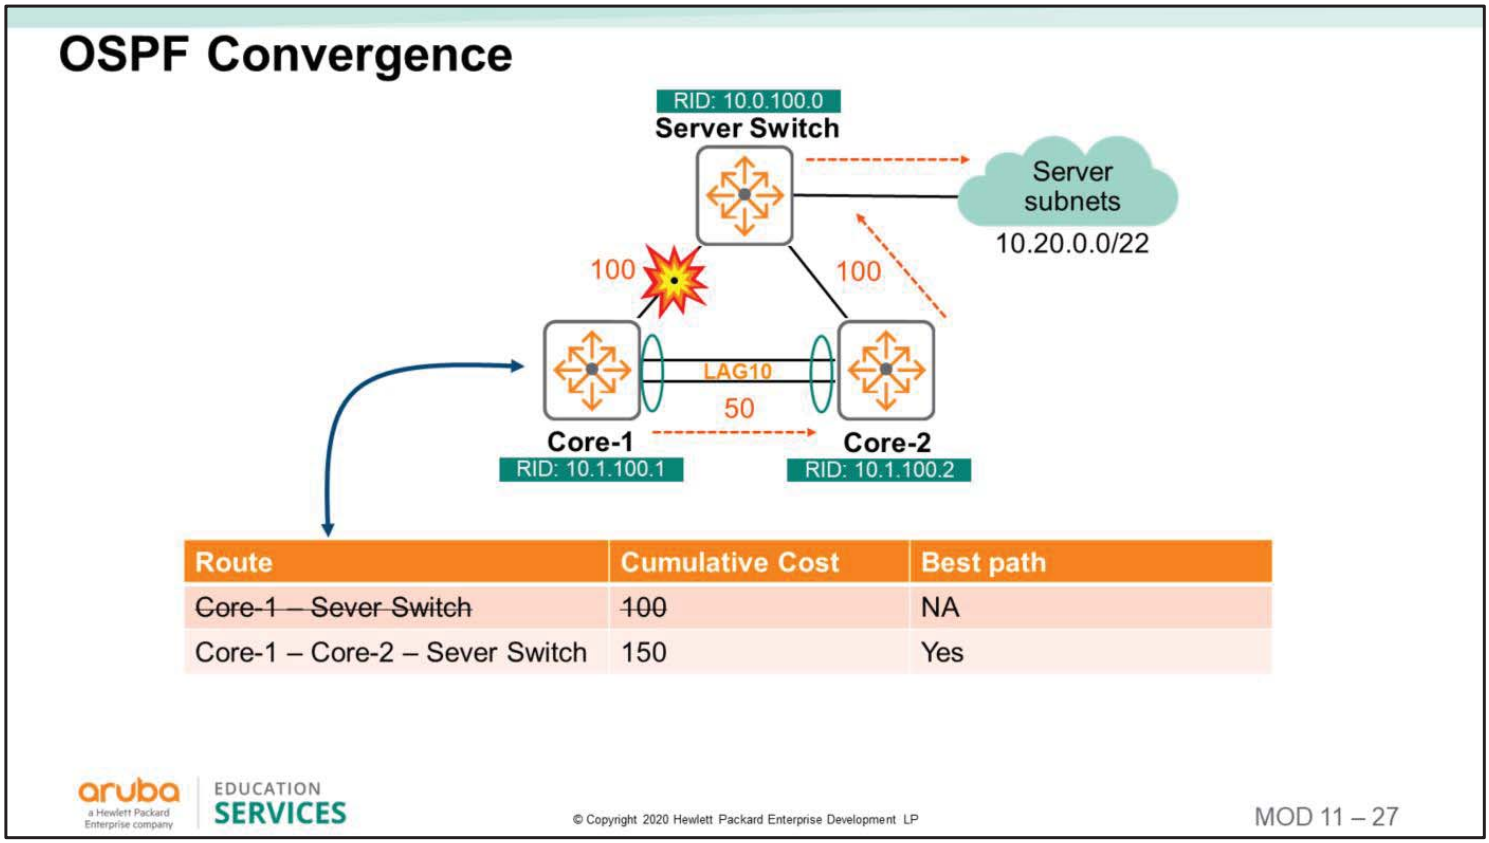

There are two components to OSPF routing convergence:

- **Detect** topology changes
- **Recalculate** routes

Topology change detection is supported in two ways by OSPF. The first, and quickest, is a failure or change of status on the physical interface. The second is a timeout of the OSPF hello timer. An OSPF neighbor is deemed to have failed if the time to wait for a hello packet exceeds the dead timer, which defaults to four times the value of the hello timer. The default hello timer is 10 seconds, so the default dead timer is 40 seconds.

When a change is detected, an LSA is sent to all routers in the OSPF area to signal the topology change. In the example shown, Link between Core-1 and Sever Switch has failed. Server Switch and Core-1 detect this outage – their links go from an UP state to a DOWN state. Thus, Server Switch and Core-1 originate the topology change LSA. Then they run the SPF algorithm to calculate their new best paths to any affected network, such as 10.20.0.0/22. Server switch will no change anything the best path is directly connect it to it. Core-1 on the other hand will recalculate its best path. In this case it will use the path using Core-2.

Ultimately, Core-2 receives LSA's about the change, but no recalculation is needed from its point of view.

Each router performs route recalculation after a failure has been detected. This causes all routers to recalculate all their routes using the Dijkstra (SPF) algorithm.

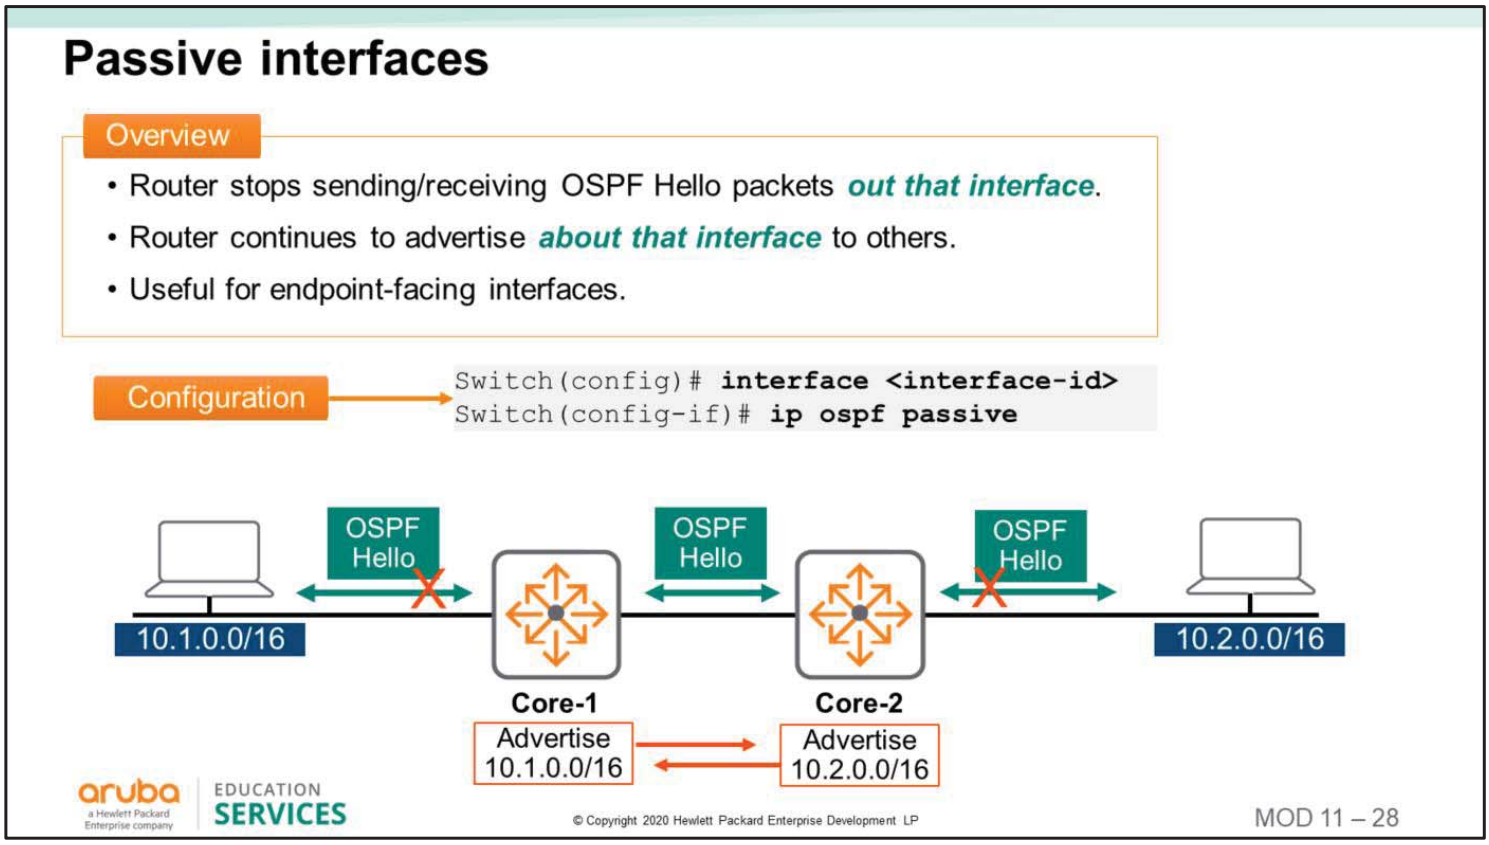

OSPF configuration involves enabling the protocol on logical and physical interfaces. Thus, the router can generate LSAs to advertise subnets to other routers. This implies that the router sends periodic Hello messages on all OSPF-enabled interfaces.

In some cases, this is not desired. The most common example is when only hosts exist on that subnet. Router constantly send OSPF packets on a network where nobody cares. Hosts do not respond to OSPF multicast 224.0.0.5 and 224.0.0.6. It is simply wasted processing cycles and bandwidth on the link. If bad actors are on the link, they could learn information about your network, and potentially mount an attack.

The solution is to configure these host-facing router interfaces as passive. When an interface is passive, it stops sending and receiving OSPF packets out that interface. However, it continues to advertise about that interface to other routes. In the example Switch-1 does not send hello packets out its 10.1.0.0/16 network, but it will tell Switch-2 about that network. Switch-2 has a similar arrangement, as relates to the 10.2.0.0/16 link.

In AOS-CX you can use the **ip ospf passive** command in the interface context to enable a passive interface.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf passive**

Lab Activity

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Additional Content

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

OSPF Scalability

Challenge

- Every router process information about every other router
- LSAs use bandwidth, LSDP uses memory, SPF algorithm uses CPU

Solution

- Split the network up into hierarchical OSPF **areas**.

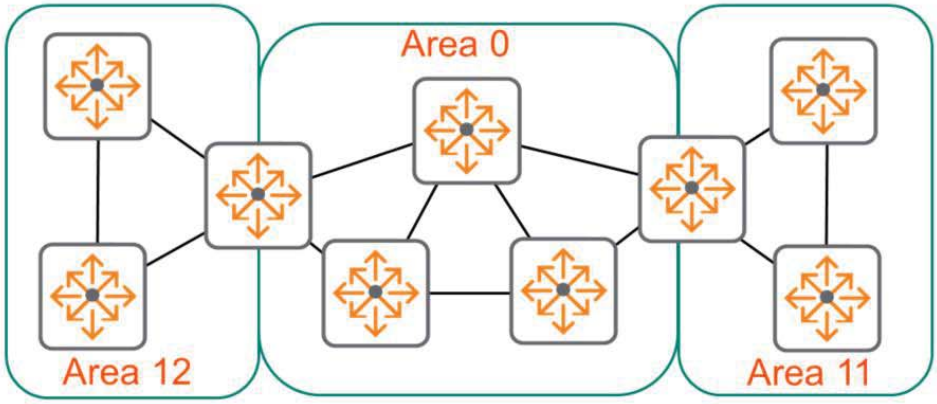

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 11 – 31

During the previous discussion, did you notice a potential scalability challenge with OSPF? You learned that every single router sends LSAs that advertise nearly the entire contents of their LSDB to every other router. With hundreds of thousand of routers, these LSA packets can begin to consume significant bandwidth. The LSDB (topology database) can grow quite large, consuming memory resources on each router. The SPF algorithm must then run on this very large database, consuming CPU cycles. This is not scalable.

The solution is OSPF hierarchy – implemented by splitting the network up into Areas.

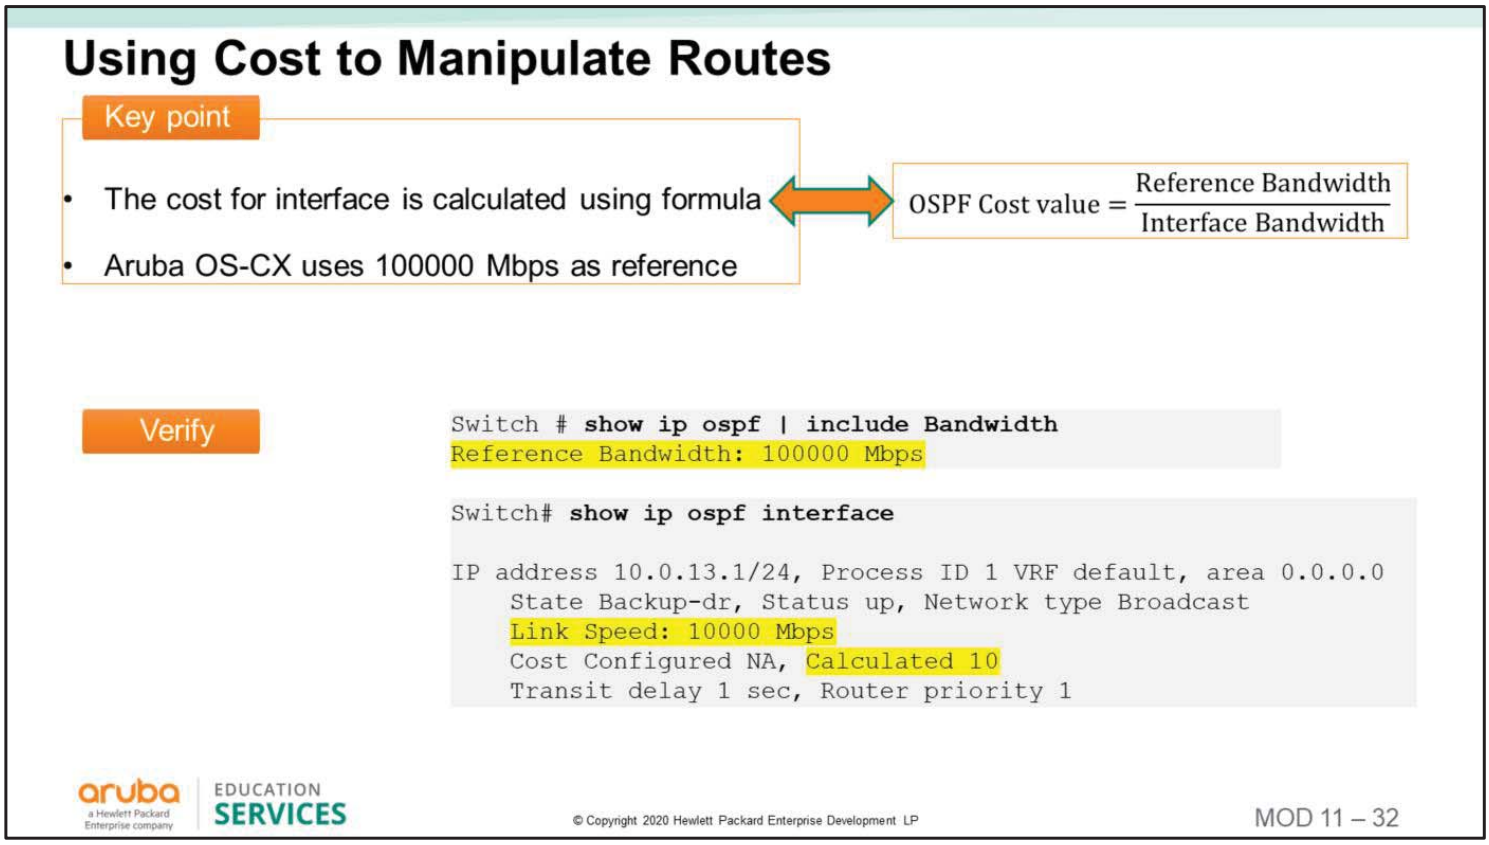

The cost associated for each interface is calculated using this formula:

$$\text{OSPF Cost value} = \frac{\text{Reference Bandwidth}}{\text{Interface Bandwidth}}$$

AOS-CX uses a default reference value of 100000 Mbps. You can verify this value using the **show ip ospf** command.

The result of the OSPF cost formula can be display using the **show ip ospf interface** command. Notice that a 10Gbps port has assigned the Cost of 10.

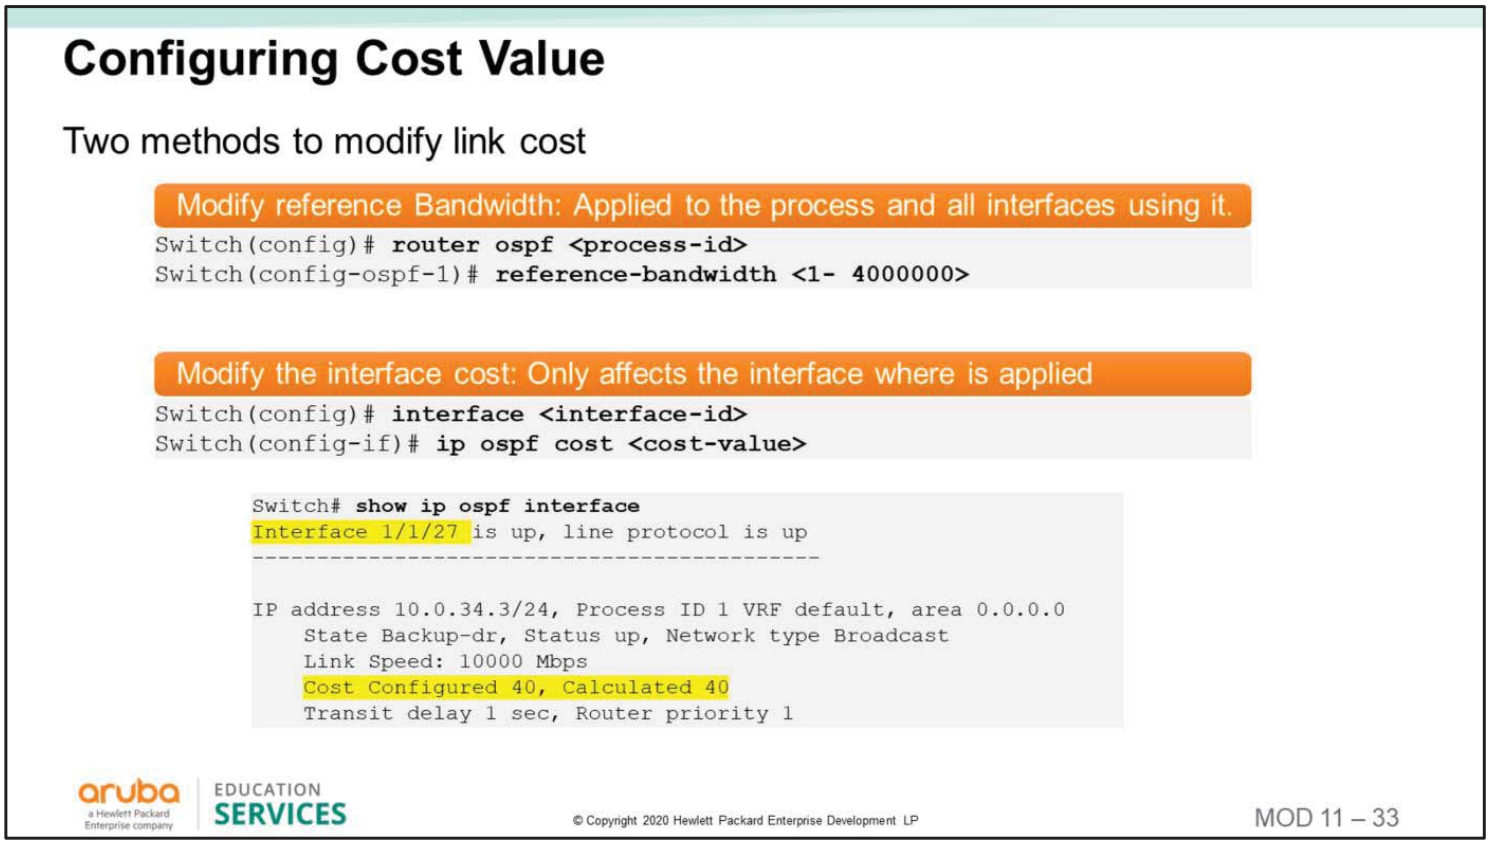

AOS-CX allows you to modify the cost of an interface in two different ways:

Modifying the reference bandwidth value. This method is applied to the entire OSPF process, which means that will affect all interfaces.

Switch(config)**# router ospf <process-id>**
Switch(config-ospf-1)**# reference-bandwidth <1- 4000000>**

Modifying the cost associated with an interface. This approach allows you to only modify the cost value to a specific interface.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf cost <cost-value>**

When this command is applied the router no longer considers the cost formula but instead simply uses the manual value that was entered.

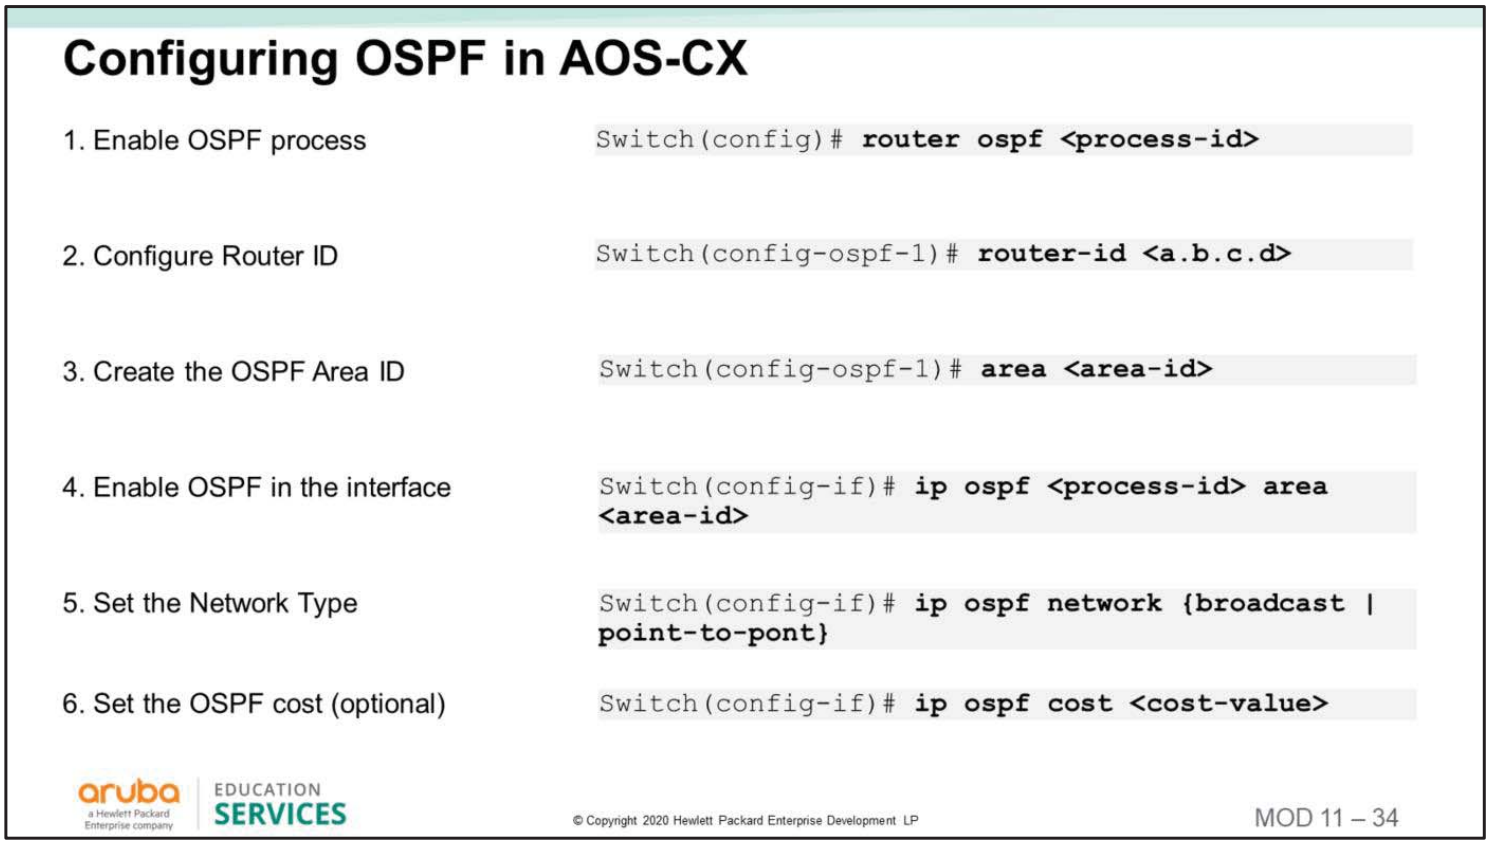

1. Enable OSPF process. AOS-CX is capable to run a single the OSPF process per VRF. The **router ospf <process-id>** command is used to enable the OSPF process and assigns a process value to it, this value is a number between 1 and 63.

Switch(config)**# router ospf <process-id>**

2. Configure the Router-ID.

Switch(config-ospf-1)**# router-id <a.b.c.d>**

3. Create the OSPF area. AOS-CX allows you to configure the Area ID in dot-decimal value or digital format. For example 0.0.0.0 is equivalent to 0.

Switch(config-ospf-1)**# area <area-id>**

4. Enable OSPF in the interface.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf <process-id> area <area-id>**

5. Set the network type.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf network {broadcast | point-to-point}**

6. Optional, set the OSPF cost associated to the interface.

Switch(config)**# interface <interface-id>**
Switch(config-if)**# ip ospf cost <cost-value>**

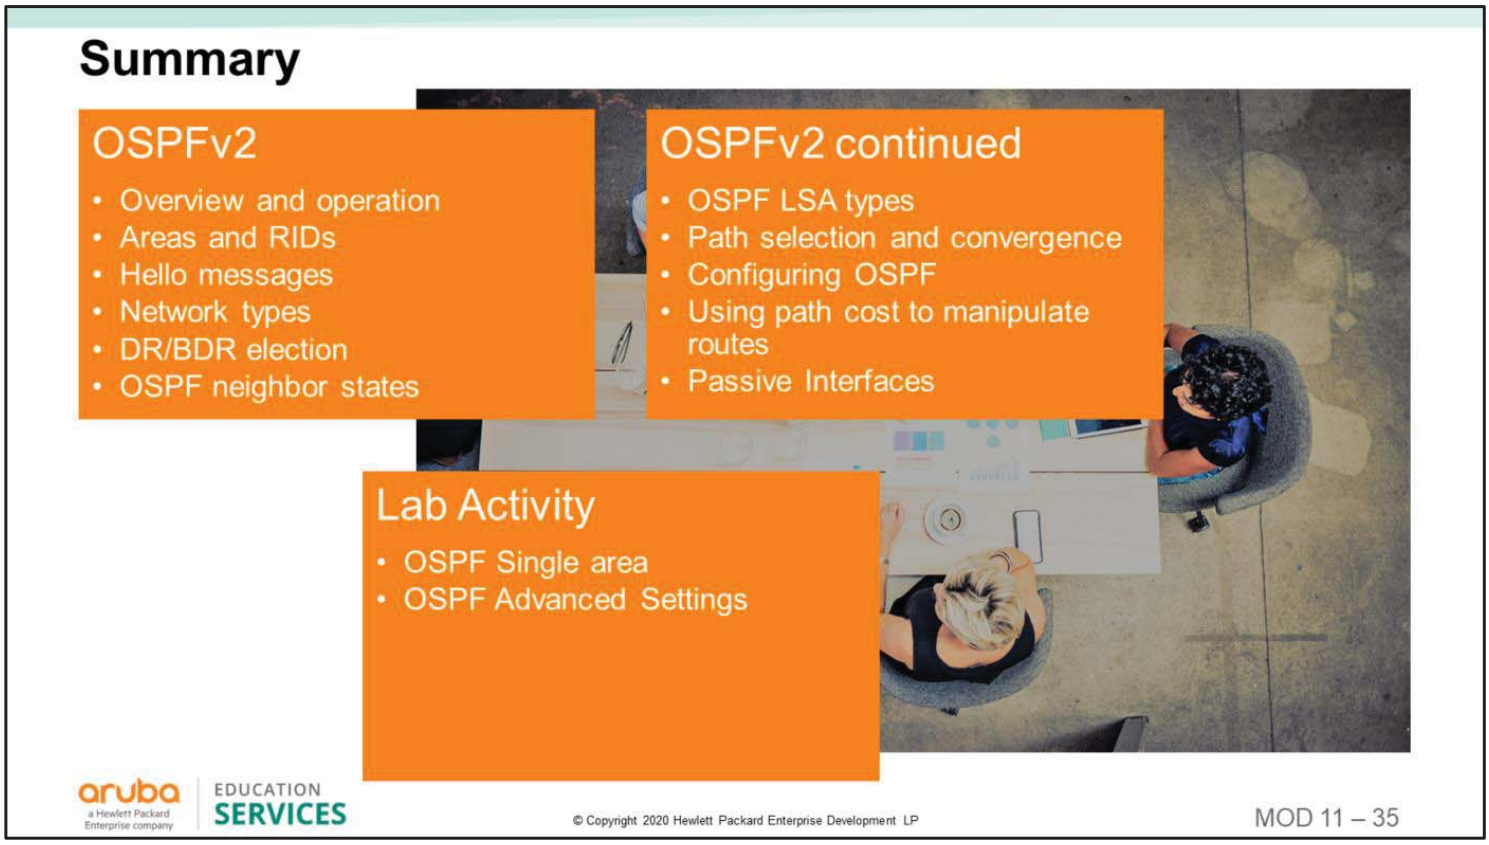

Congratulations! You should know have a good understanding of OSPF concepts and configuration. You explored OSPF operation, including OSPF areas, Router IDs, and various message types and neighbor states.

You learned how subnet information is advertised by various types of Link State Advertisements (LSA). Then you learned how these LSAs build a database of all paths, and how that database is then use to build a routing table – a list of best paths over which to forward end user data traffic.

Finally, you learned how to configure OSPF, and you applied that knowledge with a Hands-on lab activity.

aruba
a Hewlett Packard Enterprise company
Stacking
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 12 – 1

Objectives

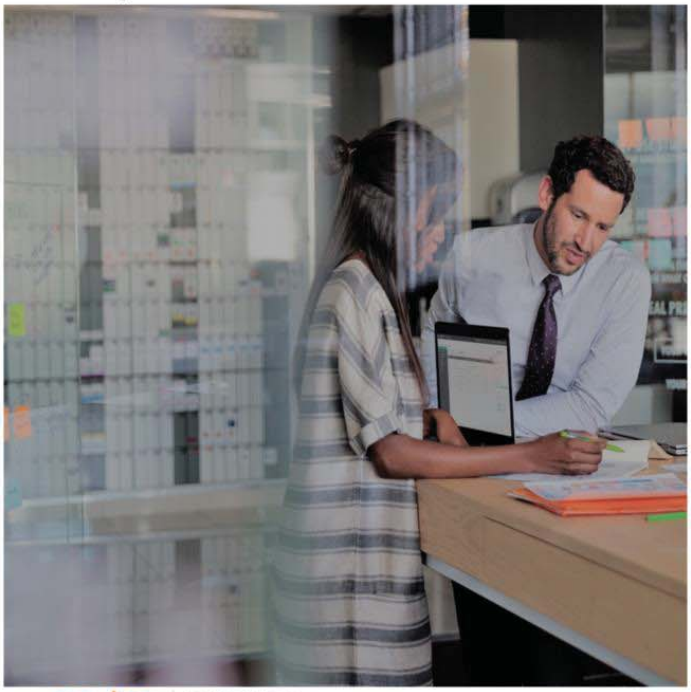

- Learning Objectives
  - Describe device operational planes
  - Describe stacking technologies and benefits
  - Explain and configure VSF
  - Trace Layer 2 traffic
  - Describe VSF failover mechanisms
  - Describe VSX

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 12 – 2

After completing this module, you will be familiar with stacking concepts and types. You will be able to describe device operational planes –the control, management, and data planes. You will also be able to describe stacking technologies and their benefits, and configure Aruba's Virtual Switching Framework (VSF), and to trace Layer 2 traffic.
You will be able to describe VSF failover mechanisms, and to describe Virtual Switching eXtension (VSX) at a high level.

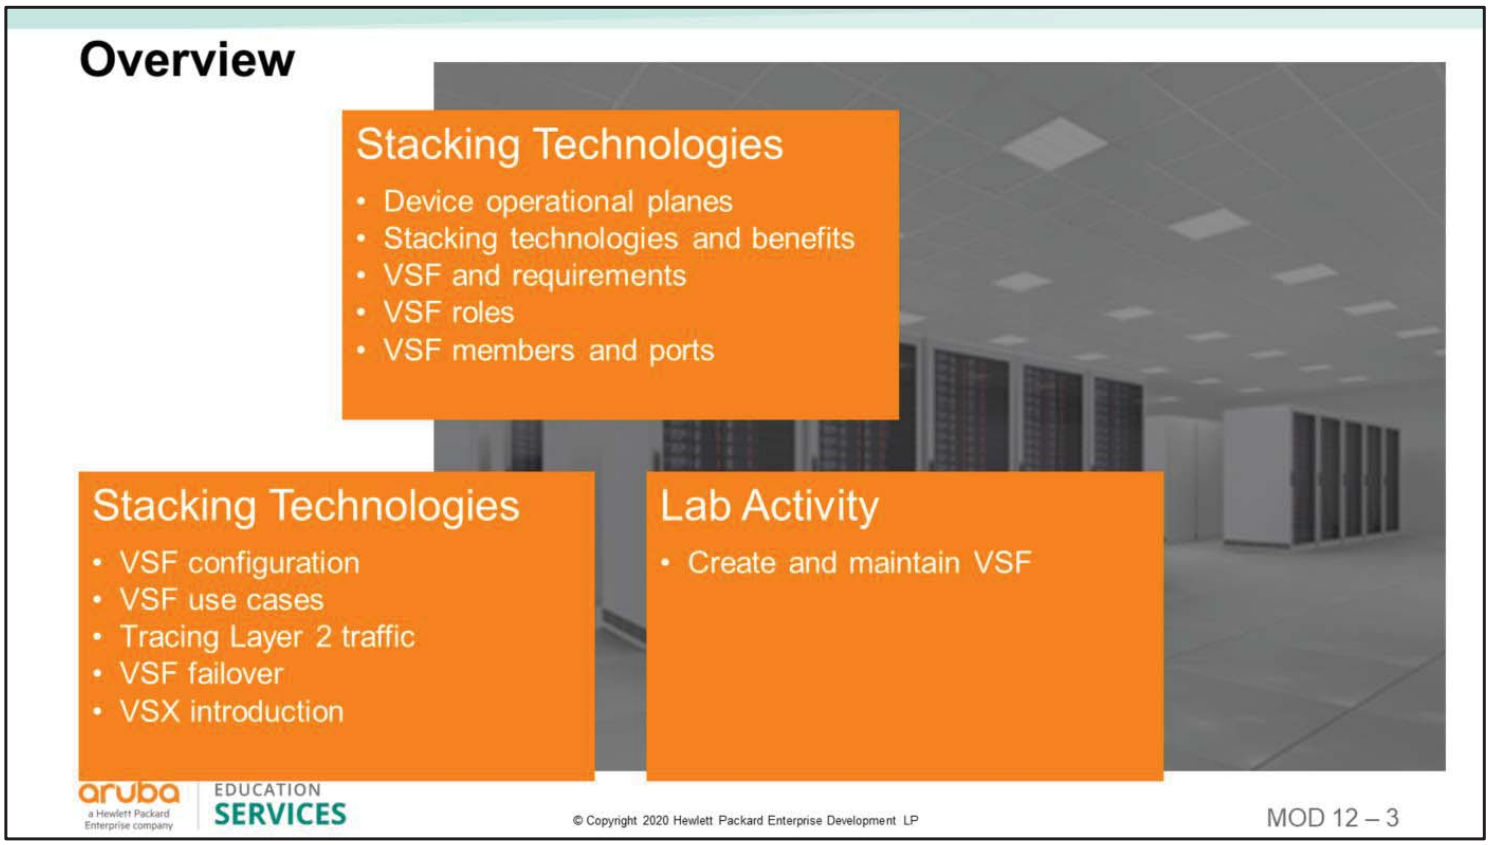

The knowledge you gain from this module about stacking technologies will help you to design, implement, and configure more resilient, reliable, high-performing networks. You first explore device operational planes – the control, management, and data plane, and the relationship between them. You will learn how this relates to stacking technologies – features that let you group multiple physical switches into a single virtual switch.
Then you learn about Aruba's stacking technology – Virtual Switching Framework (VSF). You will explore VSF operation, requirements, roles, members, and ports.
Then you will learn how to configure VSF, along with VSF use cases and about tracing Layer 2 traffic in a VSF scenario.
You will look at VSF failover scenarios, and how to improve upon them with Split Detection.
Then you will get a brief introduction to Aruba's Virtual Switching eXtension (VSX) technology.

Stacking Technologies

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

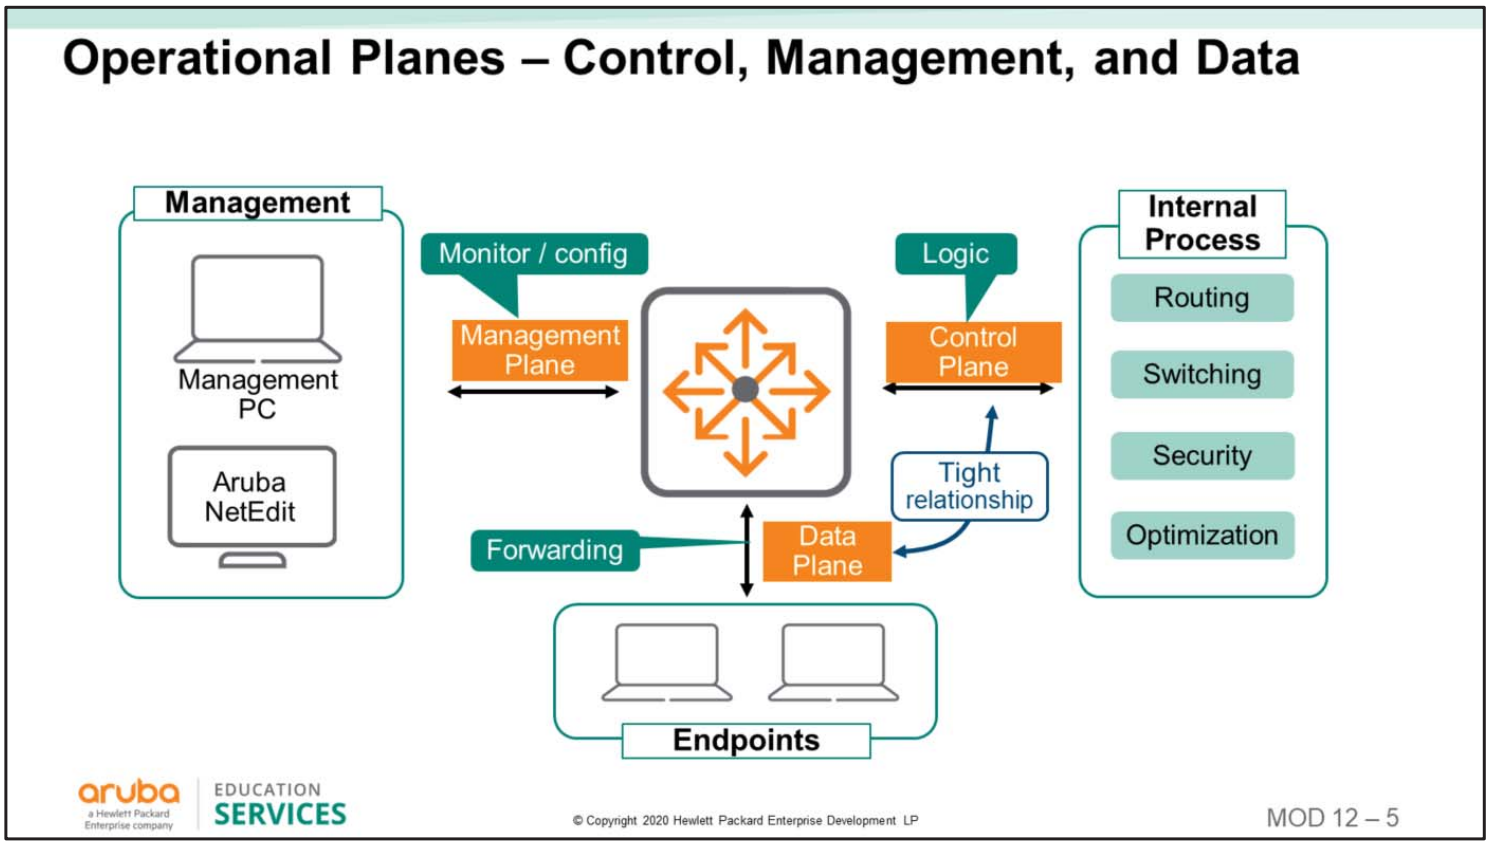

A network device is logically composed of three operational planes, and each plane performs specific tasks.

Data Plane. The data plane receives and sends frames using specialized hardware called Application-Specific Integrated Circuits (**ASIC**), which is much faster than using software. ASICs modulate and demodulate data, and handle other functions related to frame transmission and receipt.

Control Plane. The control plane logic determines what to do with the data that has been received. These decisions are made with internal processes - routing, switching, security, and flow optimization. Data plane and Control planes have a tight relationship, to process any data as fast as possible.

Management Plane. You use the management plane to monitor and configure the device. This plane must be separate from the data plane, for security and accessibility reasons. You do not want your access to the device to be completely reliant on things like VLANs or VRF. You must be able to access the device even if the control and data planes fail. Also, you don’t want end users to gain access to the management plane , this could be an egregious security issue.

Note. ArubaOS-CX devices have a specific Interface and VRF that is used for Out-of-Band Management which maintains a total separation from the data plane.

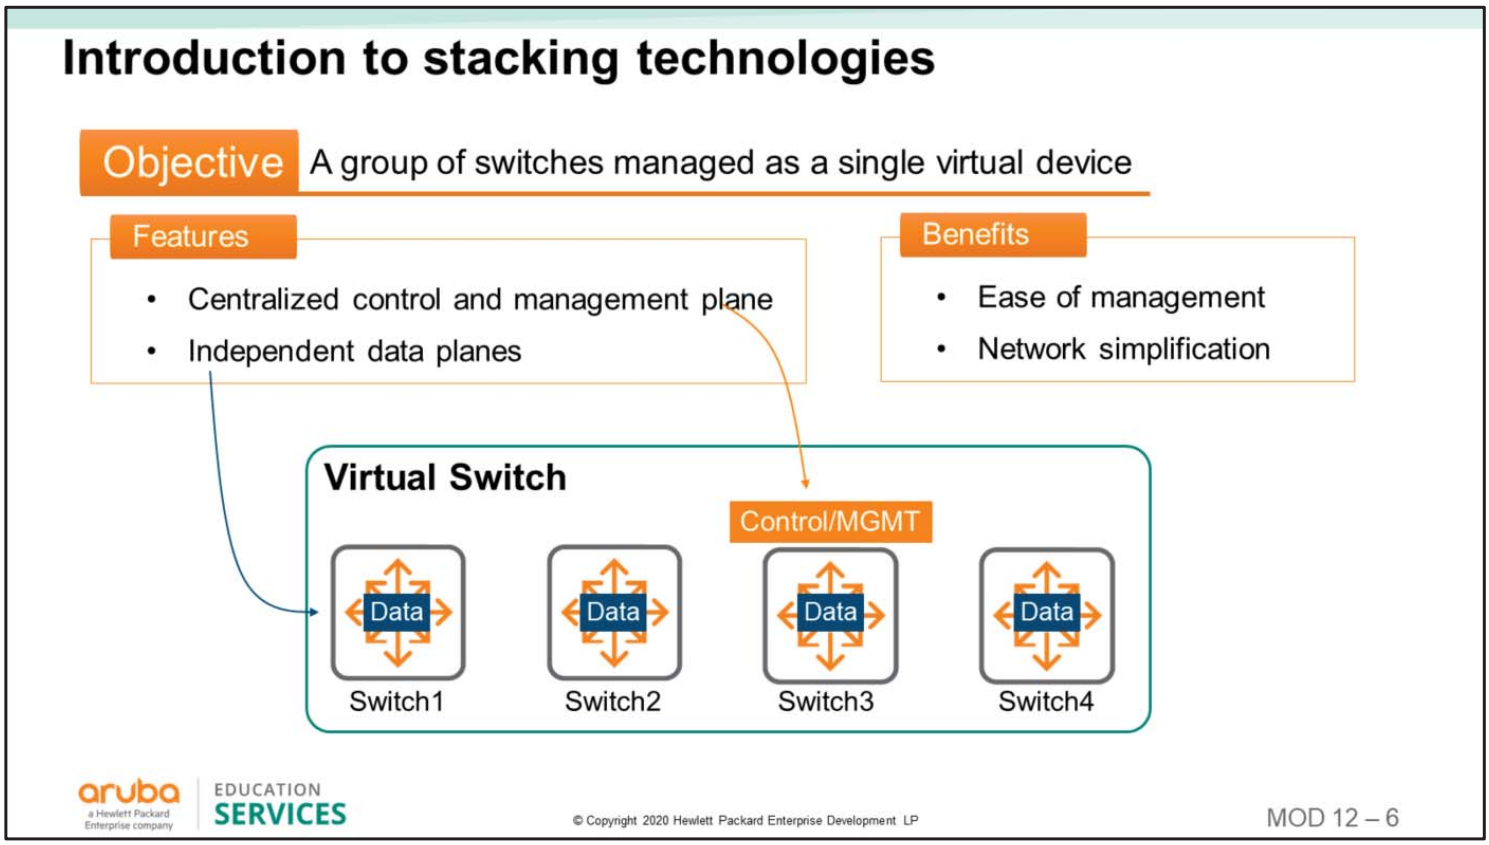

Stacking technology allows you to manage a group of switches as a single device – a Virtual Switch. Control and Management Plane functions are centralized in one group member, but each member runs its own independent data plane. The tight relationship between the control and data planes is maintained, it just happens on an inter-switch basis.

Stacking benefits include:

- **Ease of management:** You no longer need to connect to, configure, and manage each individual switch. You simply configure the primary switch. That configuration is then automatically distributed to other virtual switch members. This simplifies network setup, operation, and maintenance.
- **Network simplification:** Since multiple devices share a common control plane, routing protocols and Spanning-Tree are no longer needed inside the stacking group. Connected devices perceive the group as a single device.

Aruba switching families support two primary stacking technologies: Virtual Switching Framework **(VSF)** and Virtual Switching Extension **(VSX).** You will be introduced to VSX, but this training is focused on VSF.

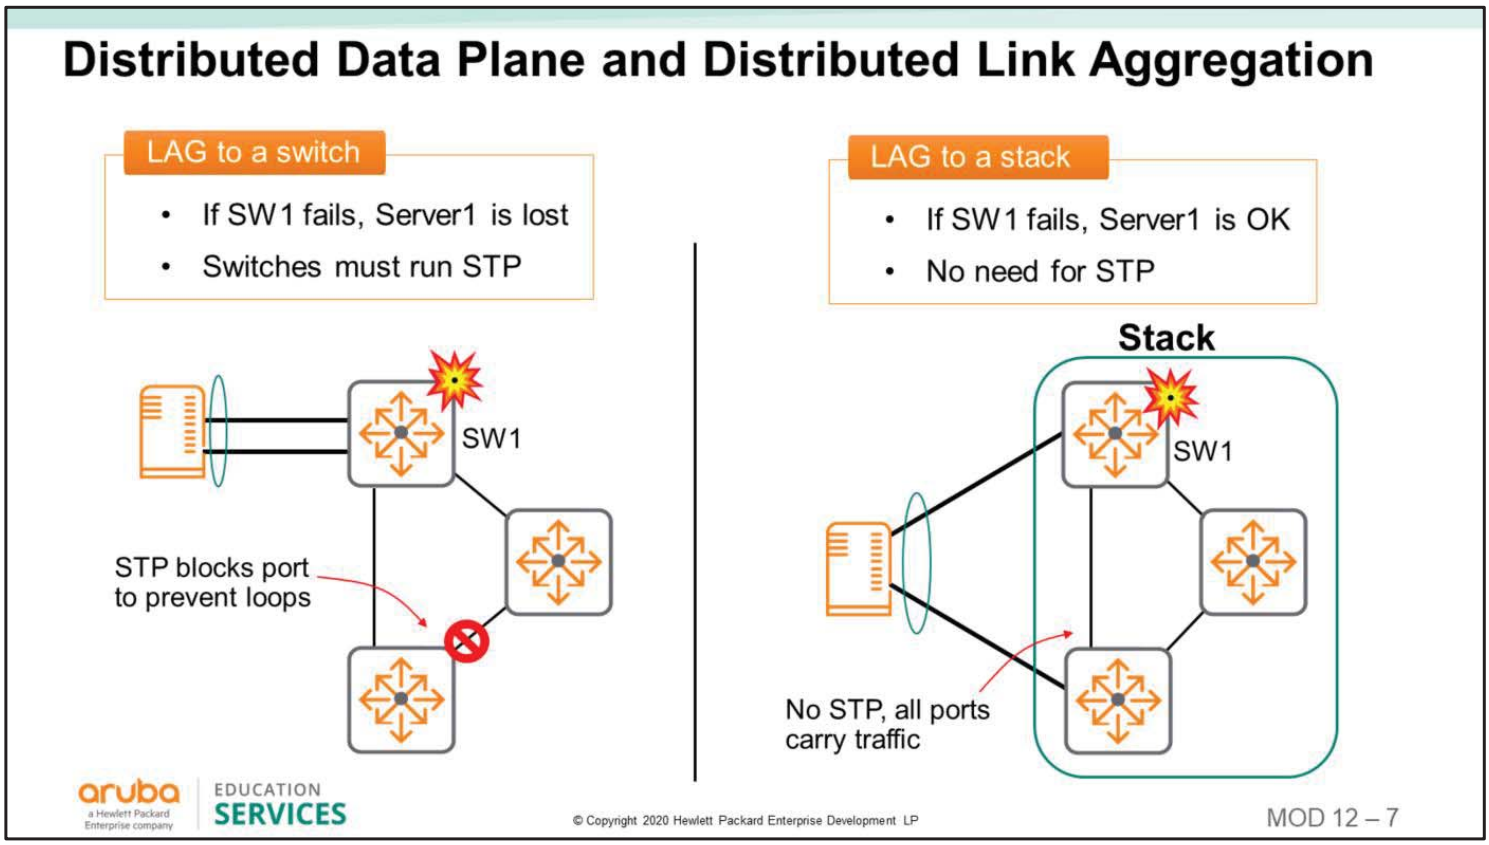

Stacking technologies maintain the data plane distributed across all the members, this implies each physical device individually creates and populates forwarding tables like ARP and MAC. These tables are then shared across all members using the control plane.

Mobility Controllers, Firewalls and servers can benefit from the stacking with LAG enable switches, since they perceive a LAG connection to a single device. However physical links terminate in different Stack members. If SW1 fails in the example shown, the traffic can still use other LAG links.

LAG and Stacking features enable the network to fully use all available links at the same time. Notice that Spanning Tree is not needed because the stack operates from a single control plane. Aruba highly recommends this implementation.

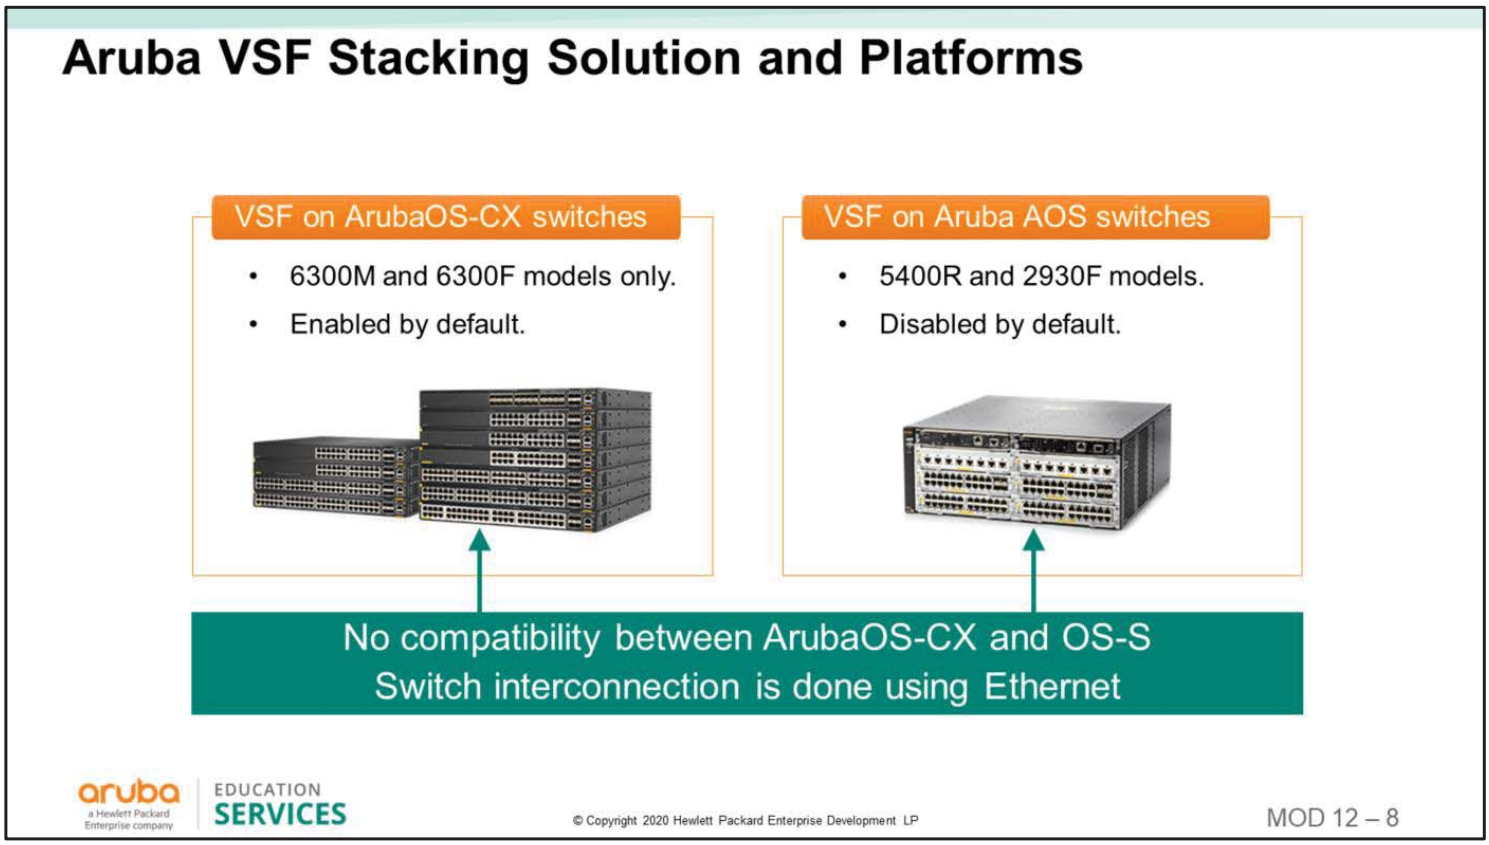

Aruba Virtual Switching Framework (VSF) defines a **virtual switch**, which is composed by individual physical switches interconnected using Ethernet links. All member switches share a **single control plane**. Devices connected to the VSF perceive a single device. As you recently learned, this virtual switch behavior has significant benefits related simplified management and improved connectivity.

In ArubaOS-CX, VSF is supported on the 6300M and 6300F models only - other models use a different technology, called VSX. You can configure a maximum of 10 members in the stack. This feature is **enabled by default**.

The VSF feature is also available in the Aruba AOS switching family, including the 5400 and 2930 series, but the feature is disabled by default. Understand that the VSF feature is not compatible between OS-CX and AOS platforms. You cannot form a VSF stack with switches from different OS families. This means that VSF can only be formed using any combination of 6300 ArubaOS-CX switching series.

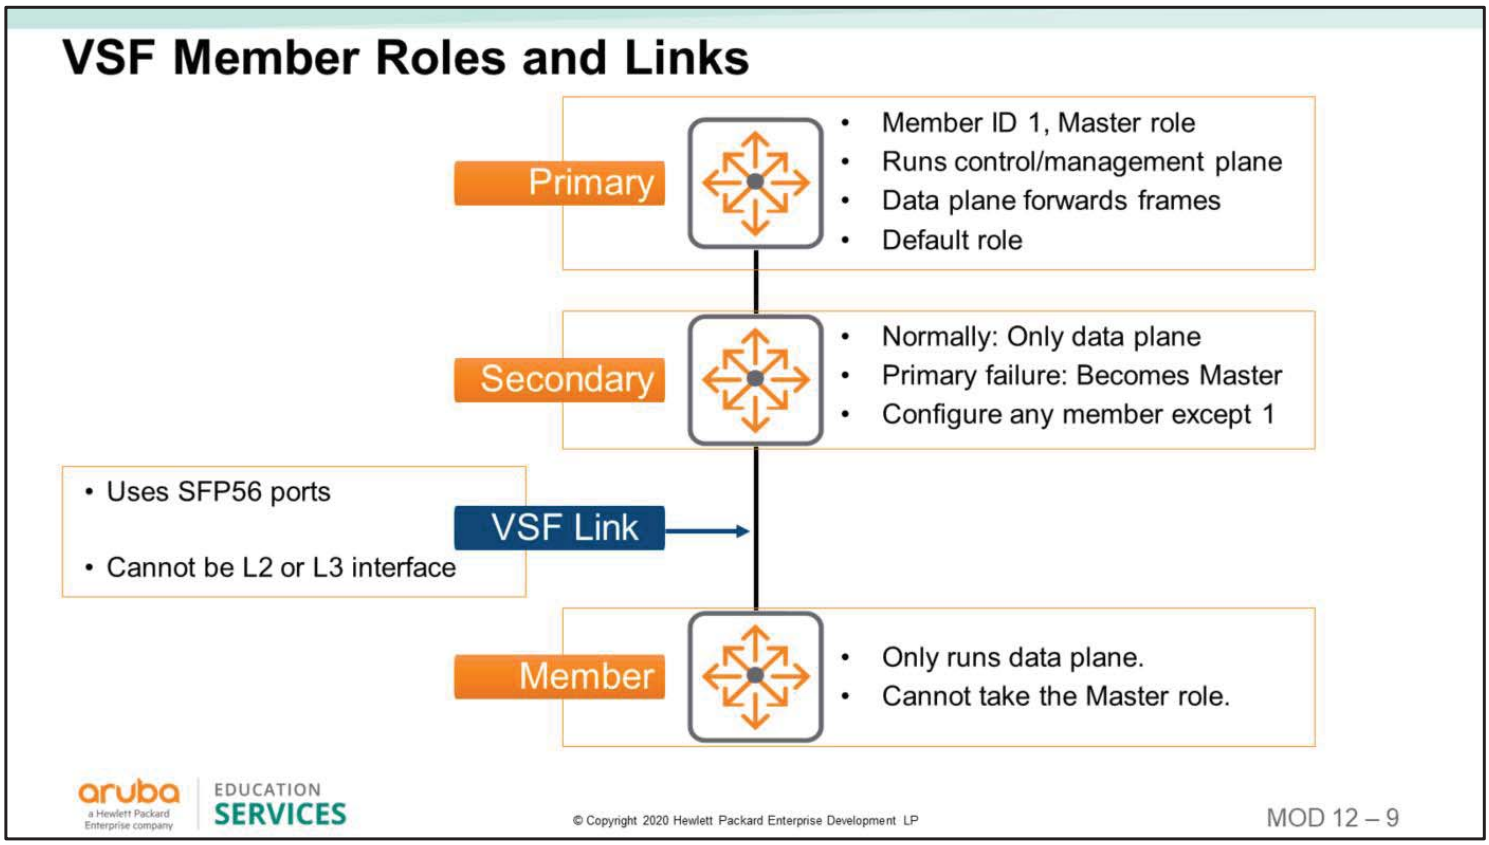

Aruba VSF creates a single control plane that runs on a single VSF member. This is the **Primary** member and always uses member ID 1. This device assumes the **Master** role in the stack.

Note. A factory-default Aruba 6300 switch boots up as VSF enabled switch with member ID 1. This implies that the switch behaves as the Primary member.

A **Secondary** member provides for high availability in case Primary failure. You can choose any member in the stack to take the Secondary role, but it must be explicitly defined. Just configure any member ID except 1.

The other devices in the VSF are **Members,** which only run the data plane, and cannot assume the Master role.

VSF switches are interconnected using SFP56 ports. When you configure a port for VSF it can no longer be used as a Layer 2 or Layer 3 interface. In other words, the port does not belong to the switch's Data plane.

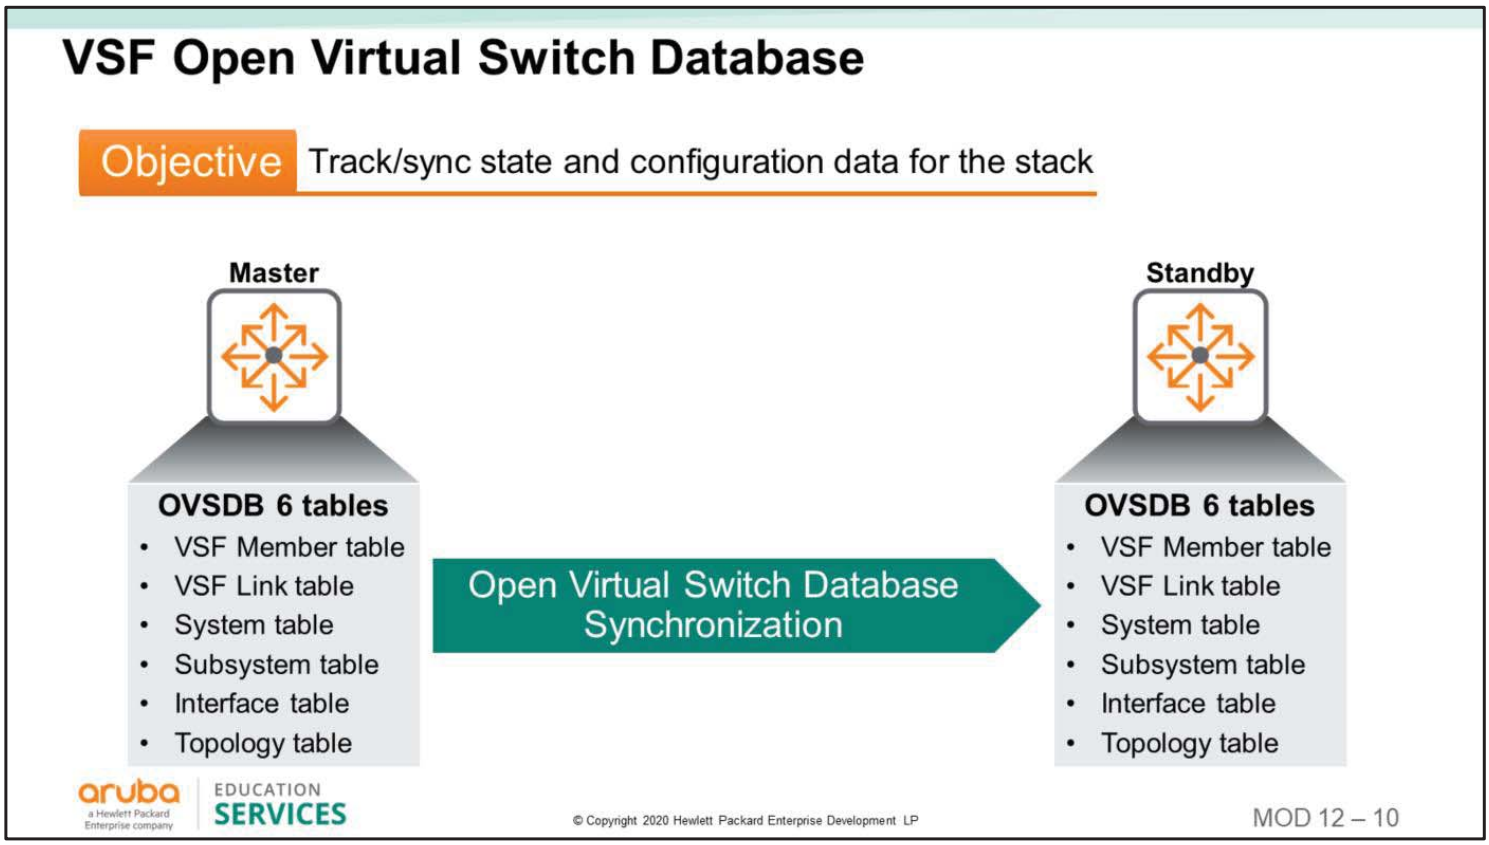

In module 2 you learned about the **Current State Database.** As the most important element of the ArubaOS-CX architecture, it is contacted by all control and management protocols. In a VSF stack this database run on the Master switch.

When VSF runs on a group of switches the Open Virtual Switch Database **(OVSDB)** is also created. This new database runs in the Master switch and contains state and configuration data for the VSF Stack itself.

The Master switch **synchronizes** OVSDB content with the standby, to ensure that it can quickly take over the master role without interruption.

The OVSDB database includes six tables:

- VSF member table
- VSF link table
- System table: Includes the number of members, MAD status, fragment status, and topology type.
- Subsystem table: Boot time for each member
- Interface table
- Topology table

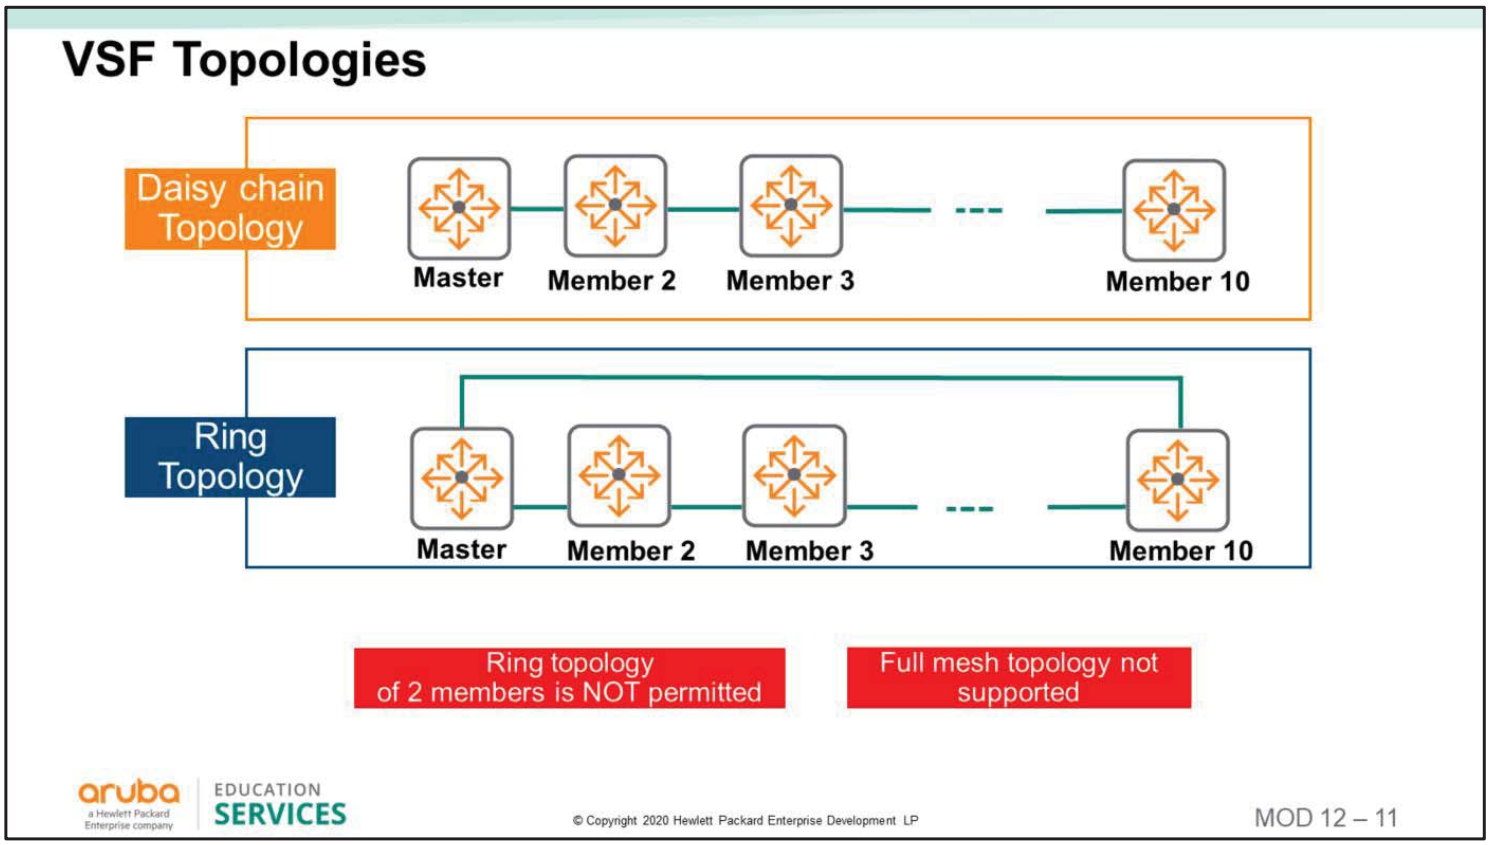

Of course, VSF members must be physically connected to form the switching stack, in one of two topologies: Daisy chain or Ring.

- **Daisy chain.** As the name implies, this topology interconnects VSF members with a single chain of Ethernet connections. From the figure you can easily see that a switch or link failure cause the stack to be split. This means that part of the stack is unable to provide endpoint connectivity.
- **Ring topology.** This topology is recommended since it offers a backup path in case of a switch or link failure.

Note. A ring topology created with a couple of switches is not permitted, since a unique link between two members most exist, this last rule takes precedence.
Note. Full mesh topology is not supported

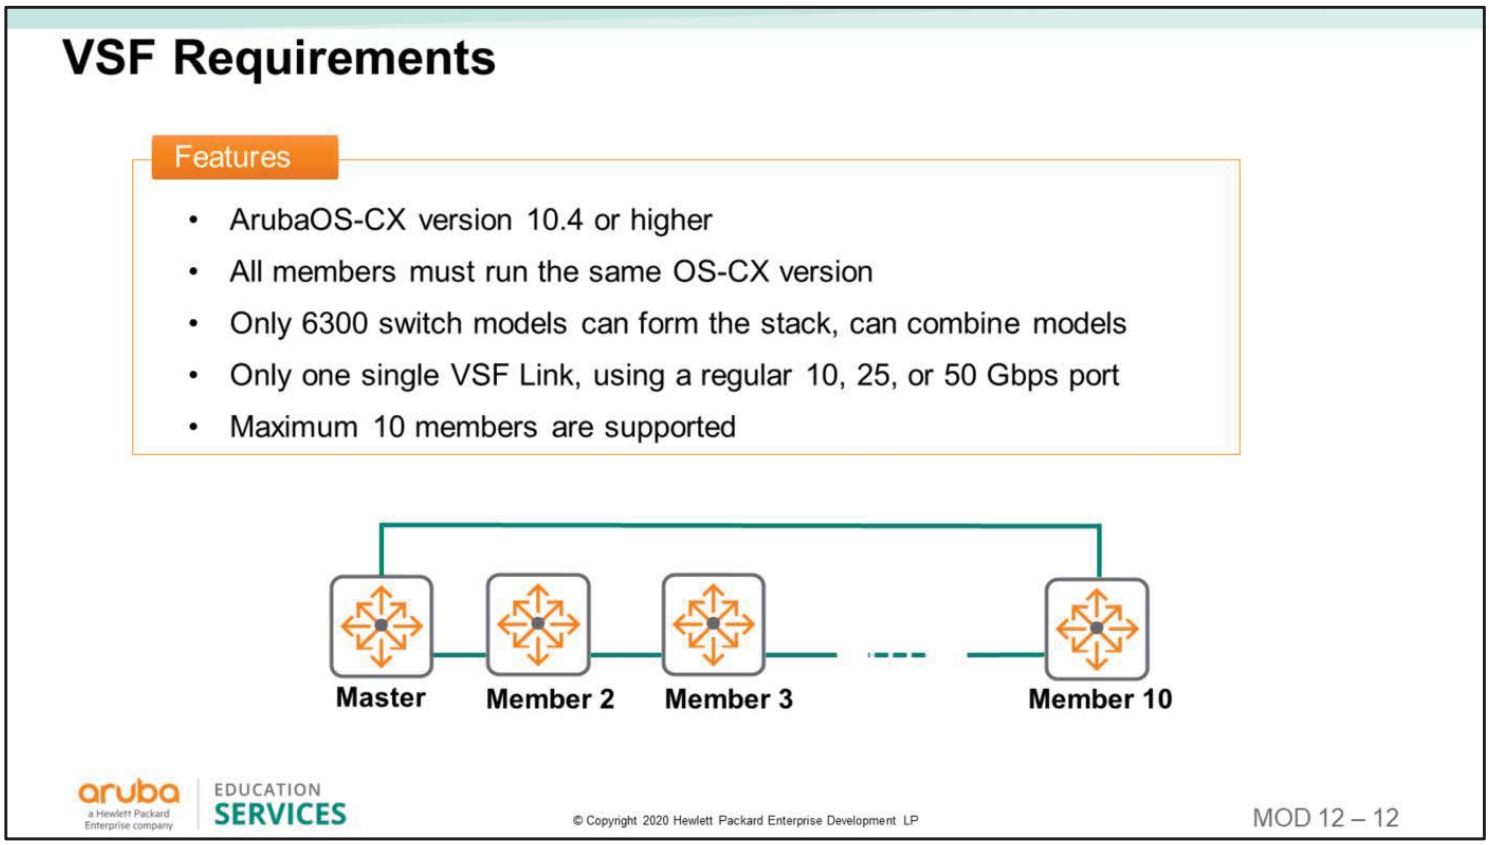

VSF has some requirements that you must meet:

- Use ArubaOS-CX version 10.4 or higher.
- All members must run the same OS-CX version.
- Only 6300 switch models can form the stack, but model combination is allowed.
- A single link is only allowed for VSF Link.
- VSF Link uses regular Ethernet port.
- Recommended to use the uplink ports which support 10, 25 or 50 Gbps.
- Maximum 10 members are supported.

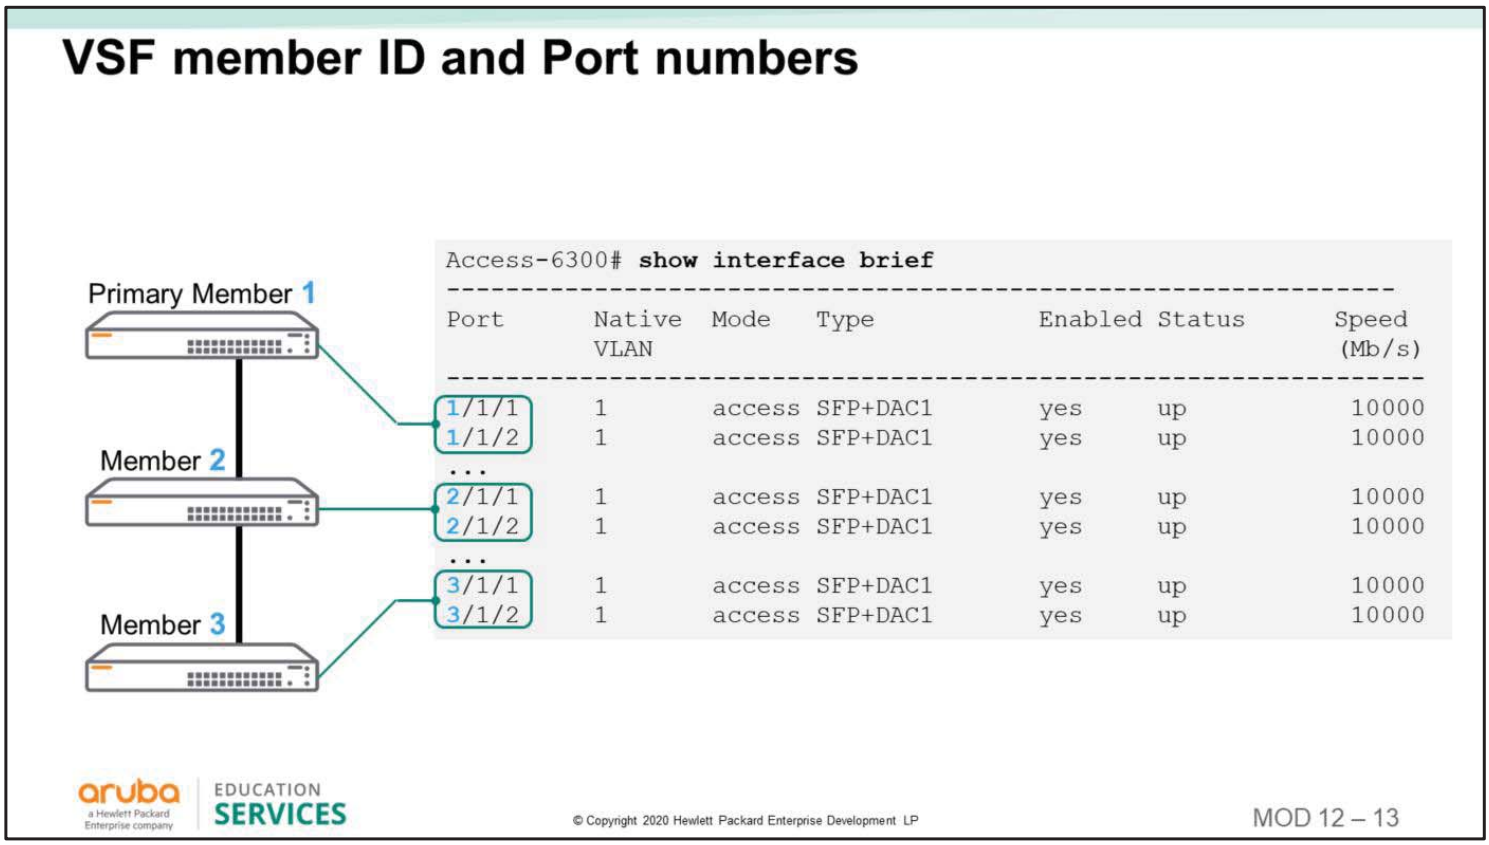

When the VSF stack forms, all physical devices use a single Control Plane. This means that all switch interfaces are be available for configuration, using the standard ArubaOS-CX **Member/Slot/Port** notation, as shown in the figure. Use the command **show interface brief** to see all available ports for all switches in the stack.

Note. Since all 6300 switch series models are fix switches, the slot number is always 1.

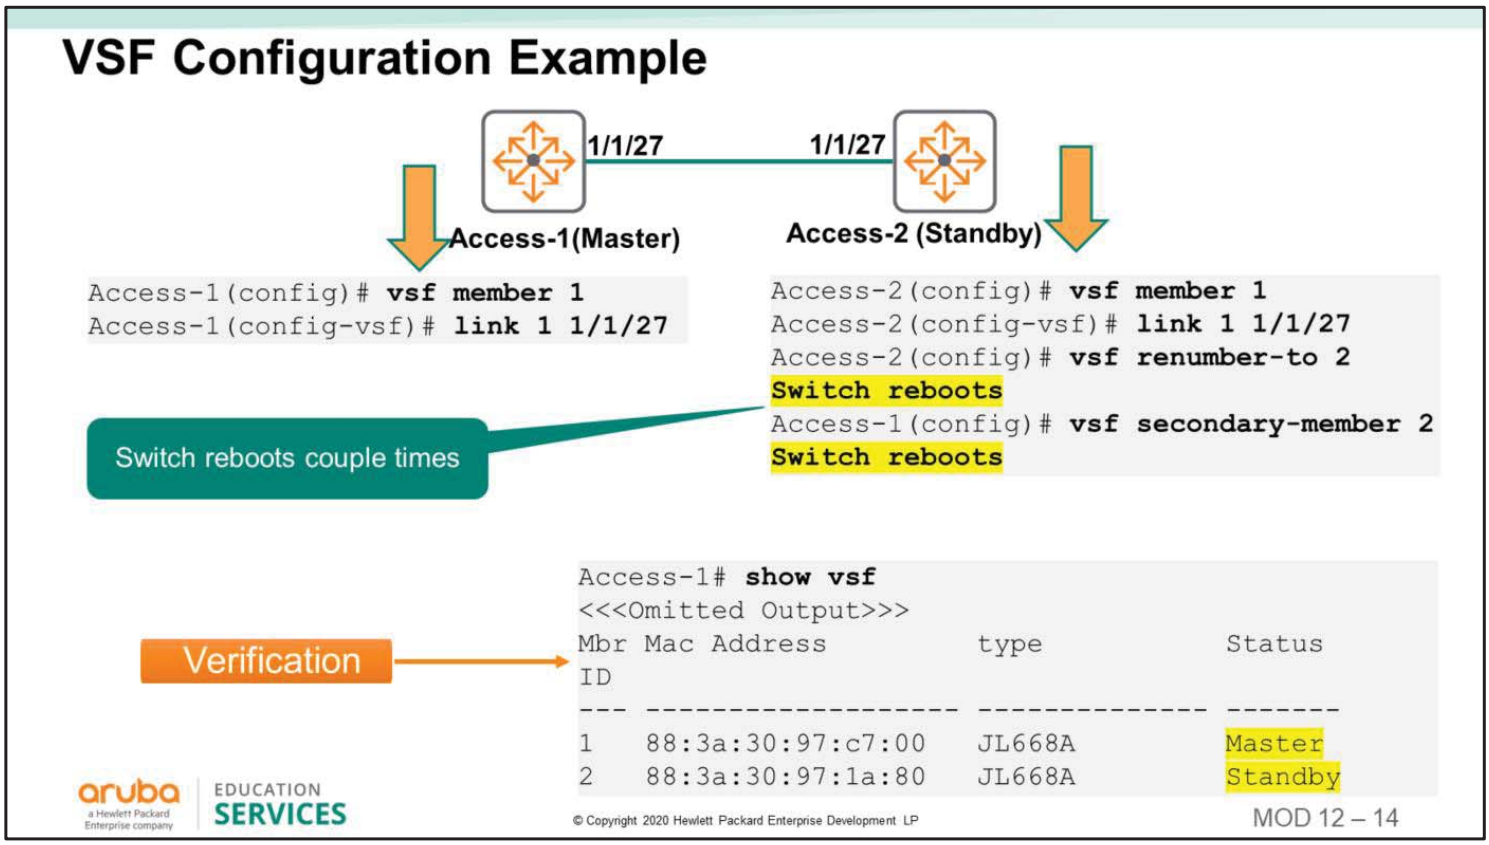

Example. Consider the scenario where a couple of 6300 switches are assigned to the same VSF stack.

Acces-1 switch

Access-1(config)# **vsf member 1**
Access-1(config-vsf)# **link 1 1/1/27**

Access-2 switch

Access-2(config)# **vsf member 1**
Access-2(config-vsf)# **link 1 1/1/27**
Access-2(config)# **vsf renumber-to 2**
Access-2(config)#**! Switch reboots**

Adding a secondary member.

Access-1(config)# **vsf secondary-member 2**
Access-1(config)#**! Switch reboots**

Use the **show vsf** command to determine member status and role.

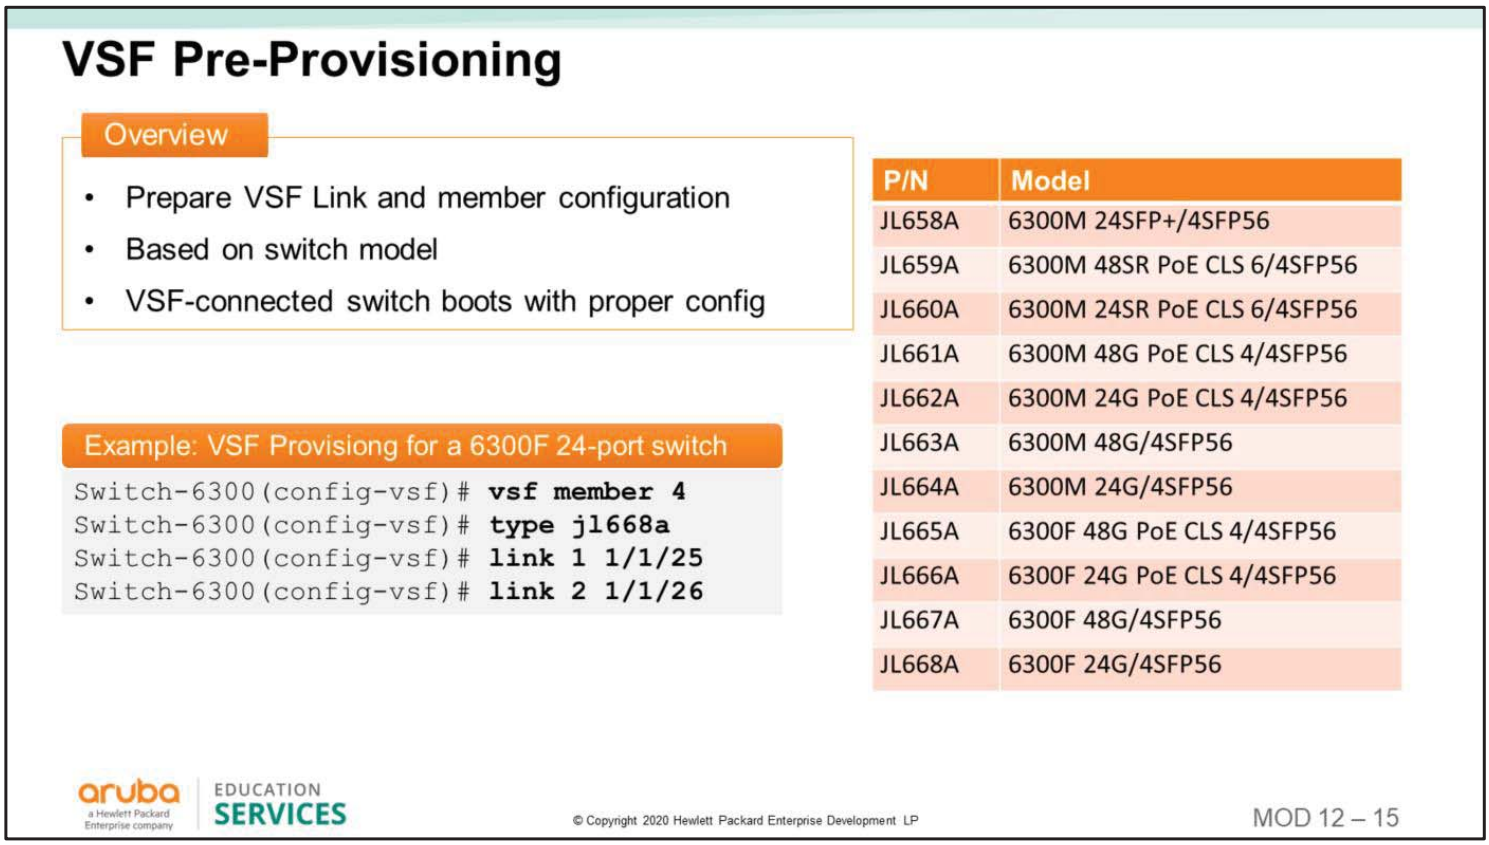

ArubaOS-CX supports member pre-provisioning. This enables you to prepare the VSF Link and member configuration for a specific 6300 switch model before the switch is connected to the stack. When the switch joins the stack, it boots up with the proper configuration.

Switch-6300(config-vsf)# **vsf member 4**
Switch-6300(config-vsf)# **type jl658a**
Switch-6300(config-vsf)# **link 1 1/1/25**
Switch-6300(config-vsf)# **link 2 1/1/26**

The table on the slide indicates the available options.

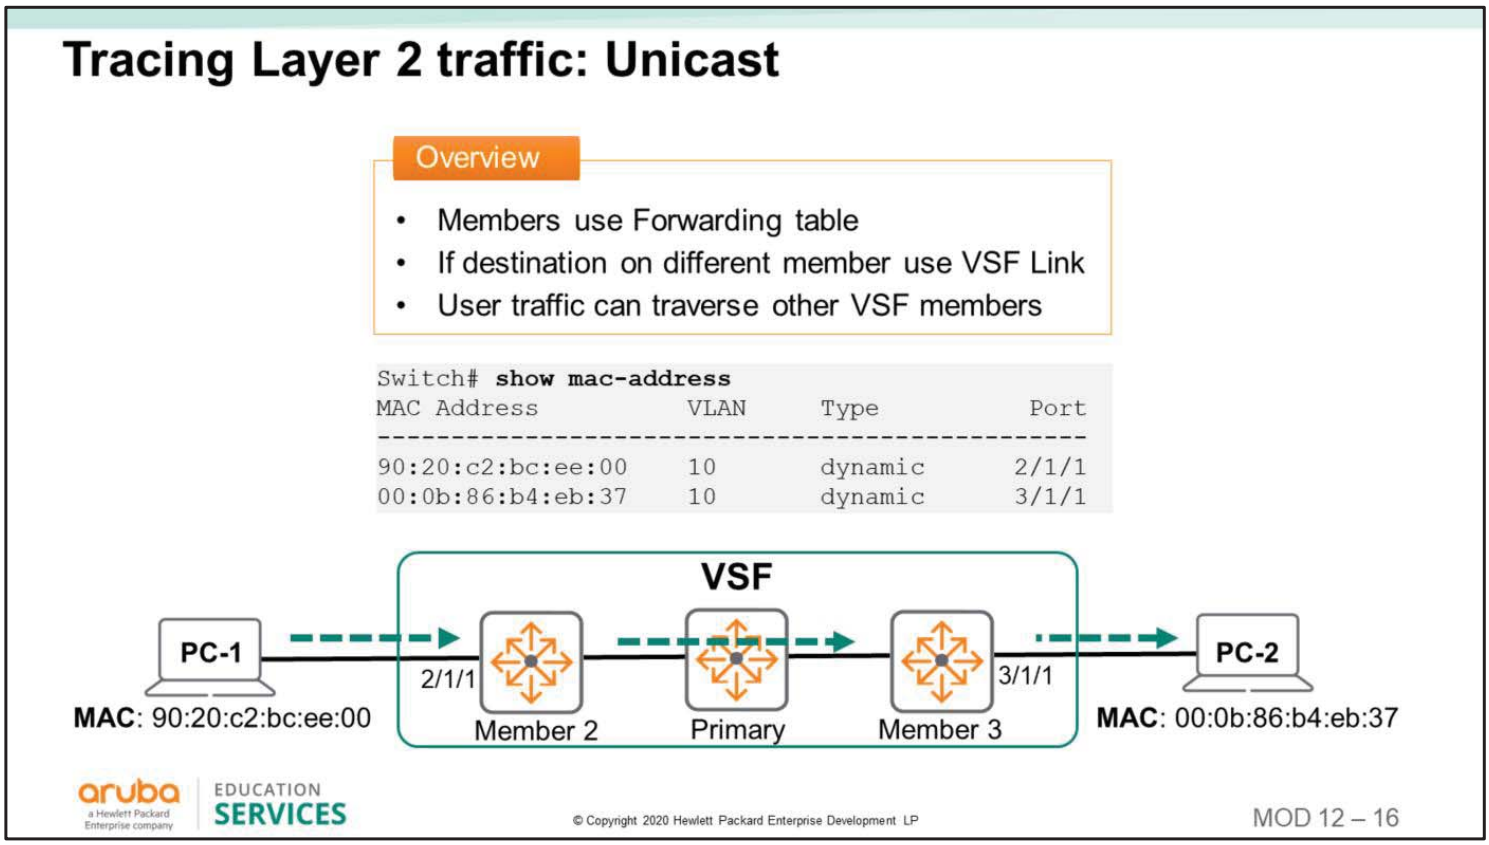

When a VSF member receives a frame for Layer 2 forwarding, it consults the L2 forwarding table to determine the egress interface. It forwards the frame out this interface, which might be local or on another member. If the egress interface is on another member, the source member forwards the packet on the VSF link.

Tracing Layer 2 traffic: Unicast

VSF Members

- Do not use the LAG hash
- Prefers to use local links
- Avoids VSF links

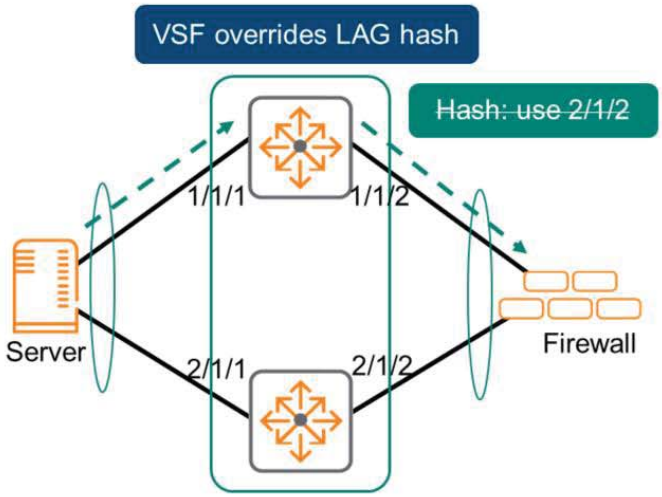

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 12 – 17

A VSF fabric is similar to any switch configured with link aggregation. It learns MAC addresses on a logical LAG entity, as opposed to the physical interfaces. It selects one LAG member link for forwarding each conversation.

However, VSF overrides the typical LAG hash function used for physical interface selection. VSF member prefers to use their own local links, and to avoid using VSF links. If the member has multiple local links in the aggregation, then it uses the typical hashing mechanism to choose between those.

VSF Failover and OSPF Graceful-Restart

VSF Master handles OSPF processes

- Losing this device affects OSPF neighbors
- Secondary member is recommended

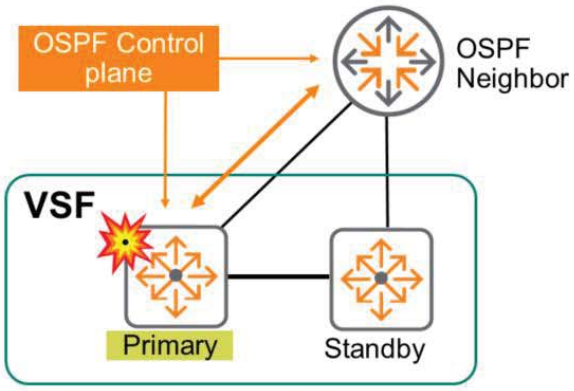

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 12 – 18

The Primary member is the most important device in the stack - it runs the control and management plane. External devices exchange routing and switching protocols directly to this member. If this device fails, the entire stack is down. Therefore you should configure a VSF Secondary member for redundancy.

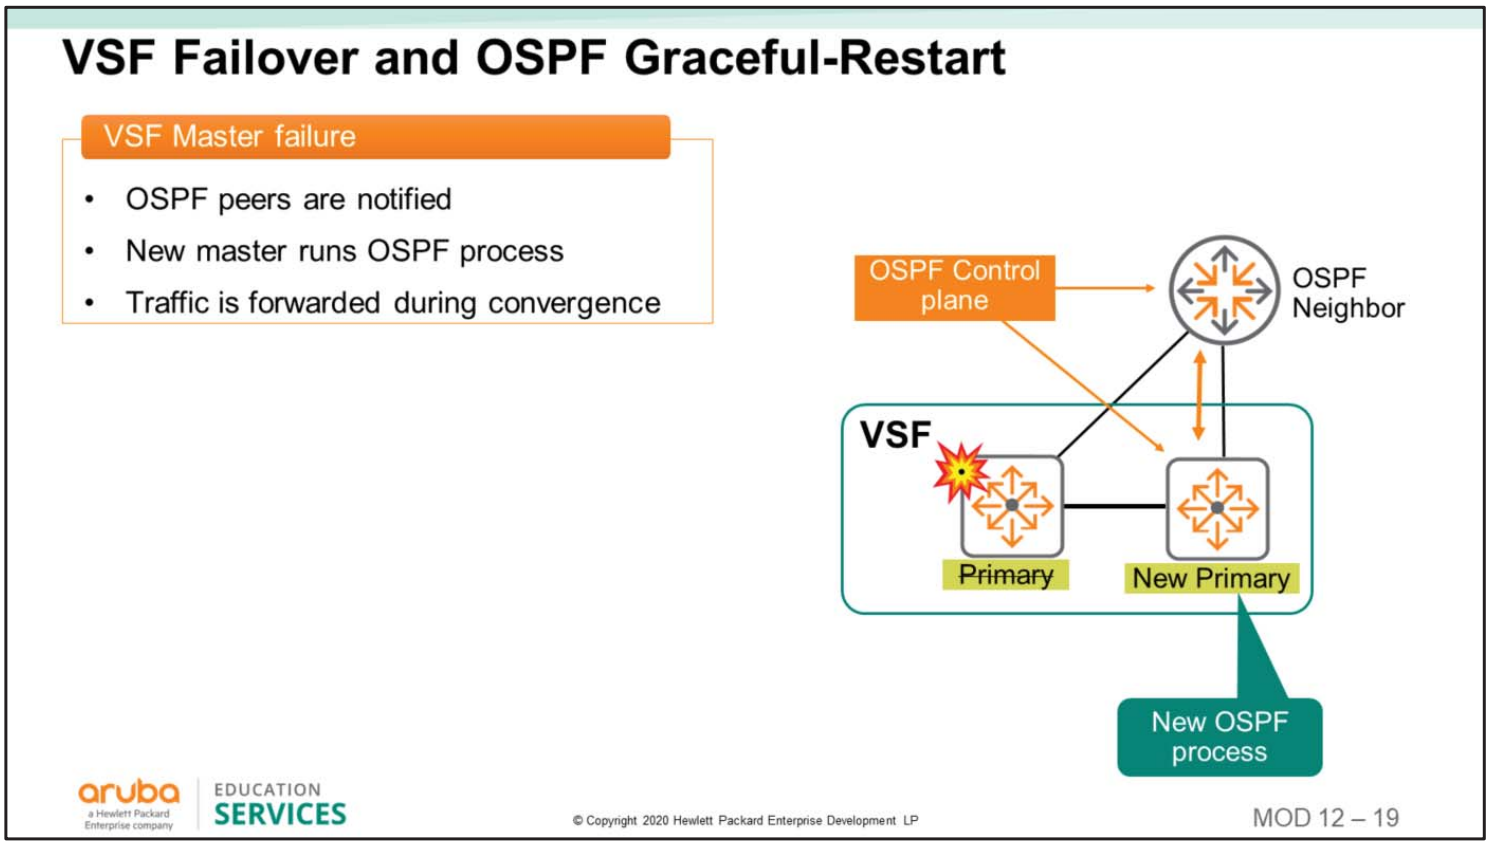

The Secondary assumes the Master role upon Primary failure. The new master runs all Control plane protocols, uses the configuration databases, and responds to management sessions.

When the Primary member fails, Layer 2 traffic continues to be forwarded without disruption. Layer 3 protocols such as OSPF perform graceful restart by notifying peers of the failover event. This triggers a rebuild of OSPF adjacencies. During convergence, the switch continues to route traffic based on the last-known routing information, prior the failure. This typically only lasts a few seconds. After OSPF is fully operational, routing uses the new information.

With VSF there is **no preemption**, this means that when the failed member re-joins the stack it will not replace the current Master of the stack, instead it takes the Standby role.

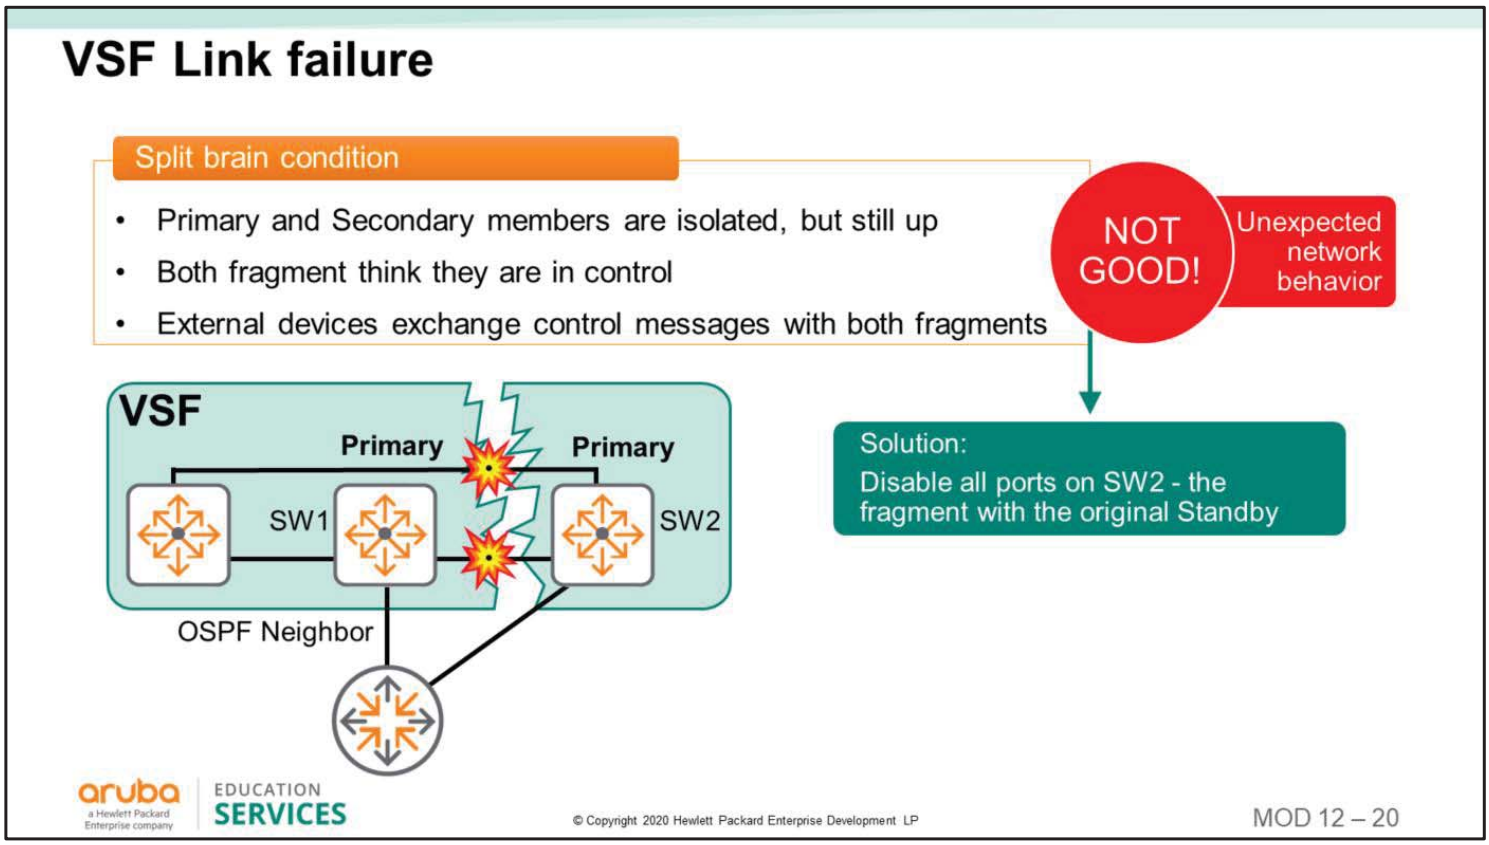

VSF Link failures could cause a fragmented stack.

The figure shows a fragmented VSF stack, due to link failures. Originally, SW1was Primary and SW2 was Secondary. After fragmentation, there is no direct connection between the two switches. SW1 continues to think that it is Primary. SW2 no longer hears from SW1, and so now thinks that it is the Primary. This is known as Split-brain condition - both fragments continue to function although there is no communication between them.

Split-brain can cause unstable, unexpected network behavior. A packet received in one fragment and destined for the other fragment is discarded. Even worse, both segments use the same IP address and the same routing information - external devices could start populating duplicated data. This can cause very strange network behavior.

Note. Split-brain situation could also occur if a VSF member fails in a Daisy chain topology.

The best way to solve a split-brain situation is to disable the ports one of the segments.

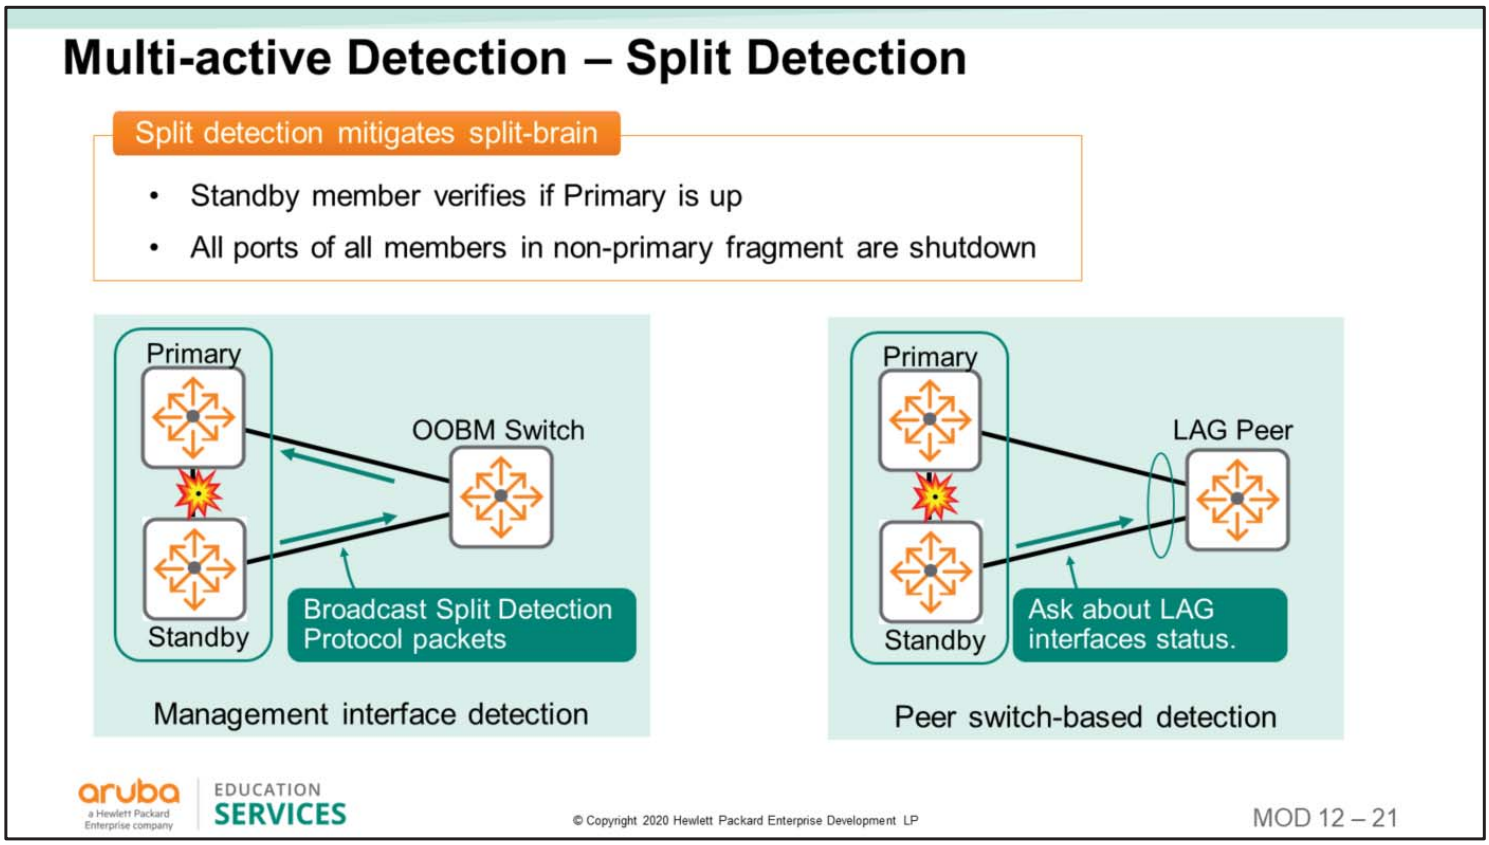

You can use Multi-Active Detection (MAD) to avoid split-brain situations. Thus, if a VSF link failure occurs, the segment that includes the Standby member verifies Primary member status. If the original Primary is up, then all members in the fragment that does not include the Primary member will disable all their ports.

VSF uses two mechanisms to detect and verify the status of the Primary member.

- **Management interface split detection.** This method requires you to connect Out-Of-Band Management (OOBM) interfaces to primary and secondary stack members. These interfaces must be in the same Layer 2 broadcast domain (VLAN). This network is used to identify active stack fragments. Each member broadcast Split Detection Protocol Packets to identify stack fragments that are currently operational.

- **Peer switch-based detection.** This method does not require additional connections and relies on the Link Aggregation Group (LAG) implementation. Switches ask the LAG Peer about its interface states – those interfaces connected to primary and secondary stack fragments. If the LAG peer indicates that its interfaces toward the Primary member are up, then the Standby member has detected a Split-brain situation, and shuts down its interfaces.

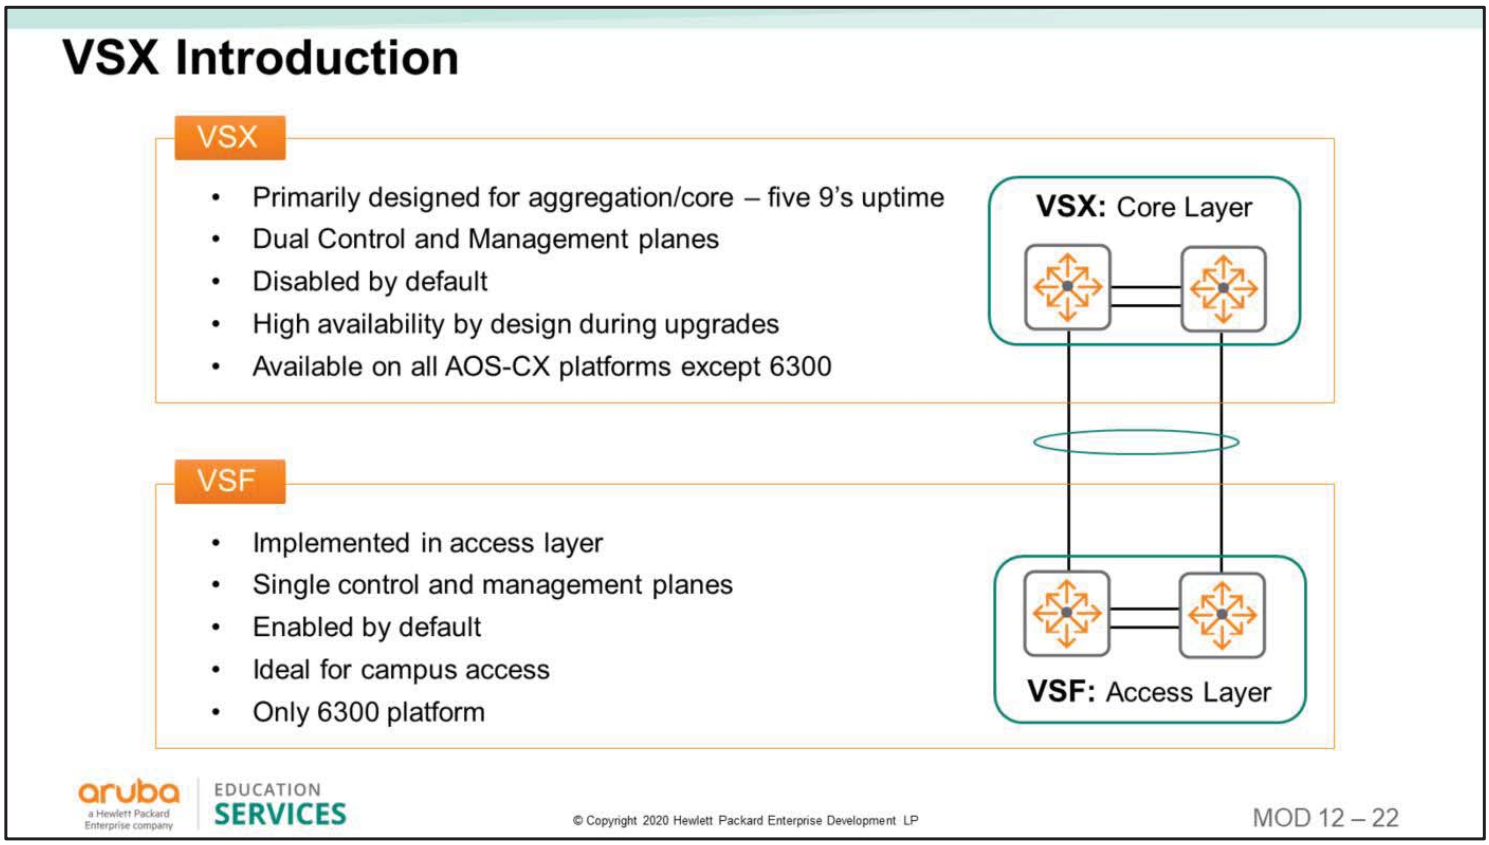

Aruba Virtual Switching Extension (VSX) is a virtualization technology for ArubaOS-CX switches. This technology can be run in all the ArubaOS-CX portfolio models except the 6300 series, which only runs VSF. VSX is commonly implemented in core devices and at data centers. VSF is suitable for the Access Layer.

VSX improves data plain performance. With VSF the control plane only can be run in the Primary member. Thus, some time is wasted when non-primary members ask the Control plane how to handle packets. With VSX, each member runs its own Control plane, allowing for faster decisions, reduced latency, and better performance.

Although VSX switches run separate control planes, they still maintain database synchronization for the configuration. Unlike VSF, each switch can modify and populate the Control plane, while presenting themselves as one virtualized switch to other devices.

VSX allows you to upgrade members with near zero downtime and with continuous packet forwarding.

Question #1

Operational Planes – Control, Management, and Data

Which of the statements below accurately describe network devices operational planes?

a. The control plane is used to control the switch configuration.

b. The management plane manages internal Layer 2 and Layer 3 switch processes.

c. The data plane moves data from ingress to egress port.

d. The control and data planes are tightly integrated.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 12 – 23

Correct answers:
c.
d.

Question #2

Aruba VSF Stacking Solution and Platforms, VSF Member Roles and Links

Which of the options below describe a valid VSF scenario?

a. An ArubaOS-CX switch acts as the primary and an Aruba AOS switch has a backup role.

b. Configure two switches one as Primary and one as member using a single VSF Link.

c. Configure VSF on 5 8325 ArubaOS-CX switches.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 12 – 24

Correct answer:
b.

Question #3

VSF Topologies, VSF Requirements, VSF Configuration

Which of the statements describe valid VSF requirements and specifications?

a. You can daisy-chain up to 10 VSF members.

b. You can connect 10 VSF members in a ring topology.

c. You can mesh multiple members together for redundancy.

d. If you use an Aruba 6300 series as the master, you can connect Aruba 5300's as members.

e. The configuration of VSF may cause members to reboot.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 12 – 25

Correct answers:
a.
b.
e.

Lab Activity

12.1 Create a Virtual Switching Framework Fabric
12.2 Maintain the VSF Stack

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

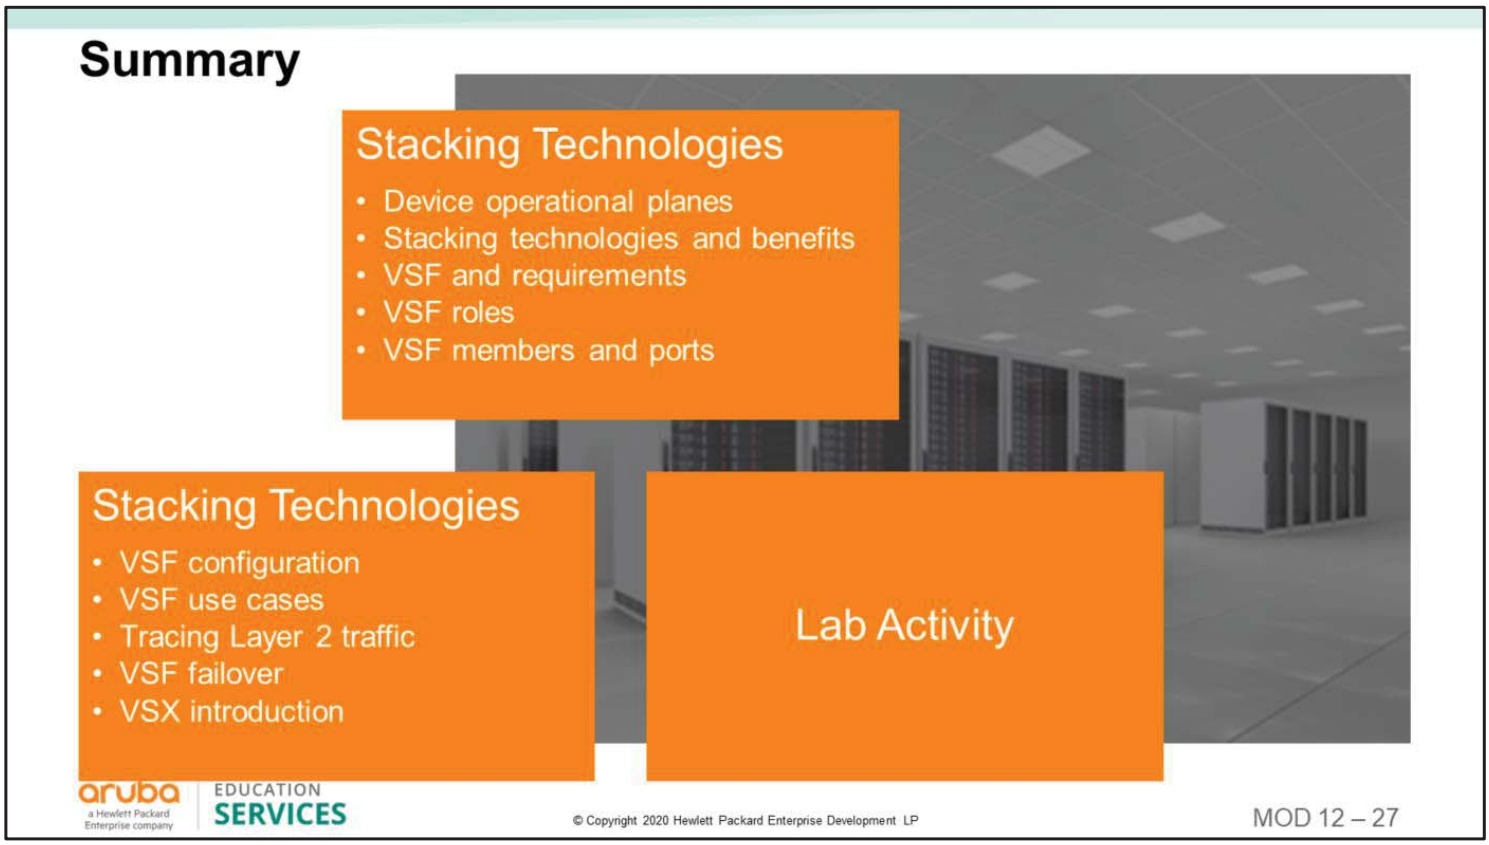

Congratulations! The knowledge you gained from this module about stacking technologies will help you to design, implement, and configure more resilient, reliable, high-performing networks. You first explored device operational planes – the control, management, and data plane, and the relationship between them. You learned how this relates to stacking technologies – features that let you group multiple physical switches into a single virtual switch.

Then you learned about Aruba's stacking technology – Virtual Switching Framework (VSF). You explored VSF operation, requirements, roles, members, and ports.
Then you learned how to configure VSF, along with VSF use cases and about tracing Layer 2 traffic in a VSF scenario.
You examined VSF failover scenarios, and how to improve upon them with Multi-Active Detection (MAD). Then you experienced a brief introduction to Aruba's Virtual Switching eXtension (VSX) technology.

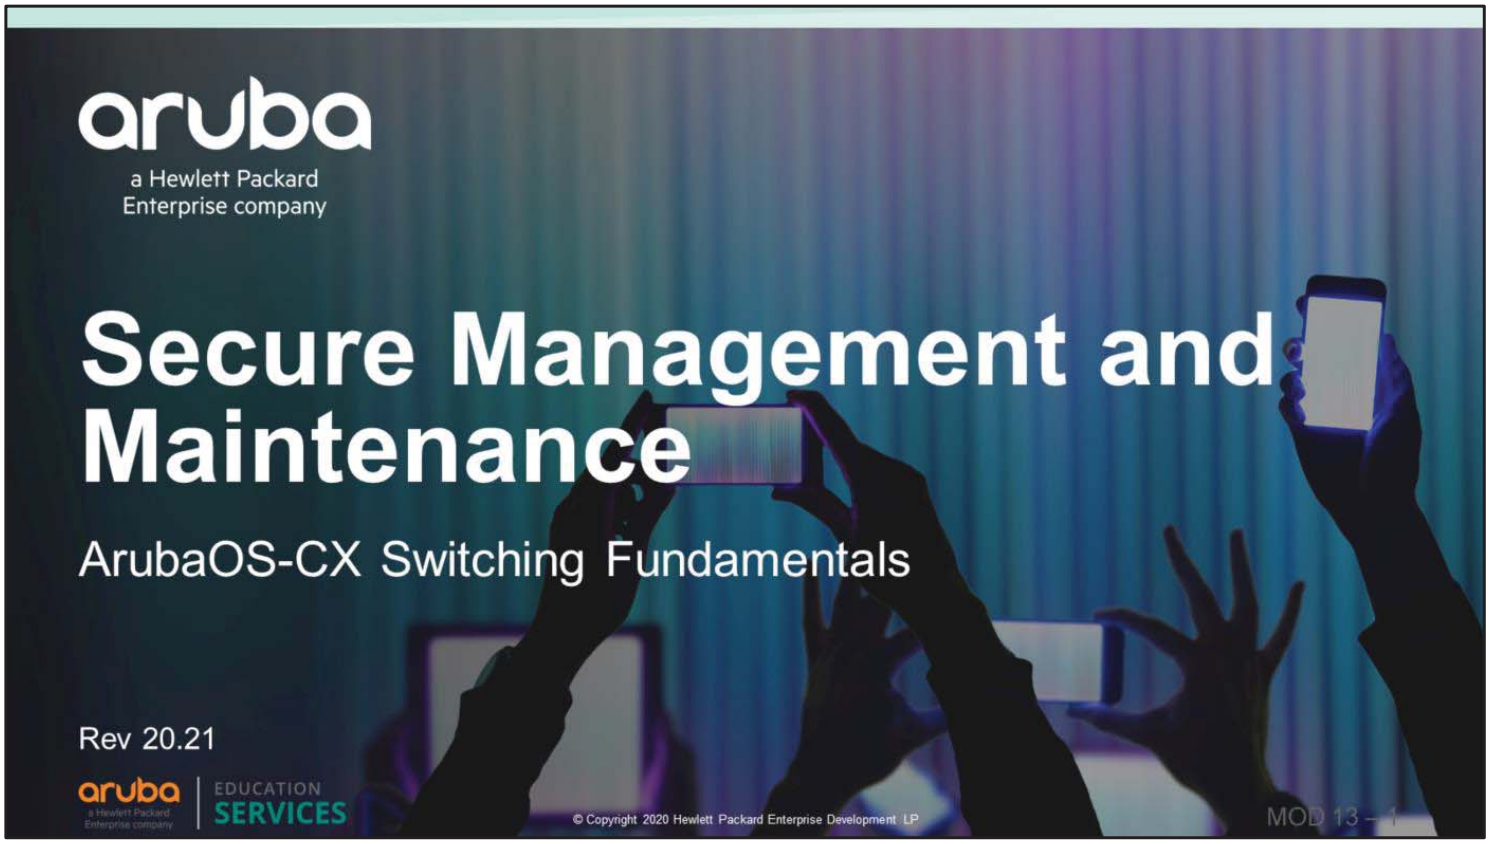
aruba
a Hewlett Packard Enterprise company
Secure Management and Maintenance
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba
EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 13 – 1

Objectives

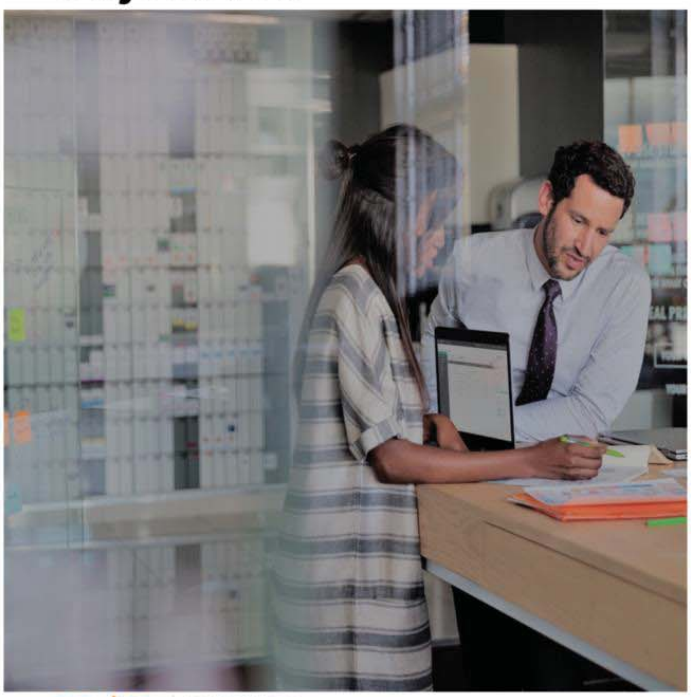

–Learning Objectives

- Describe the OOBM port and management VRF
- Explain secure management protocols
- Describe RADIUS-based management authentication
- Explain SNMP
- Describe configuration file management
- Describe OS image management
- Restore and Aruba OS-CX device to factory default and perform password recovery

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 2

After completing this module, you will be familiar key management techniques, protocols, and configurations.

You will be able to describe the Out Of Band Management (OOBM) port and the management VRF.
You will be able to explain secure management protocols, including SSH, HTTPS, AAA, and RBAC.
You will be able to describe RADIUS-based management authentication, and how it can centralize AAA services for scalability.
You will be able to explain the Simple Network Management Protocol (SNMP) –operation, architecture, versions, and configuration
You will be able to describe configuration file management and OS image management.
Finally, you will be able to explain how to restore an ArubaOS-CX device to its factory default condition, and to perform password recovery

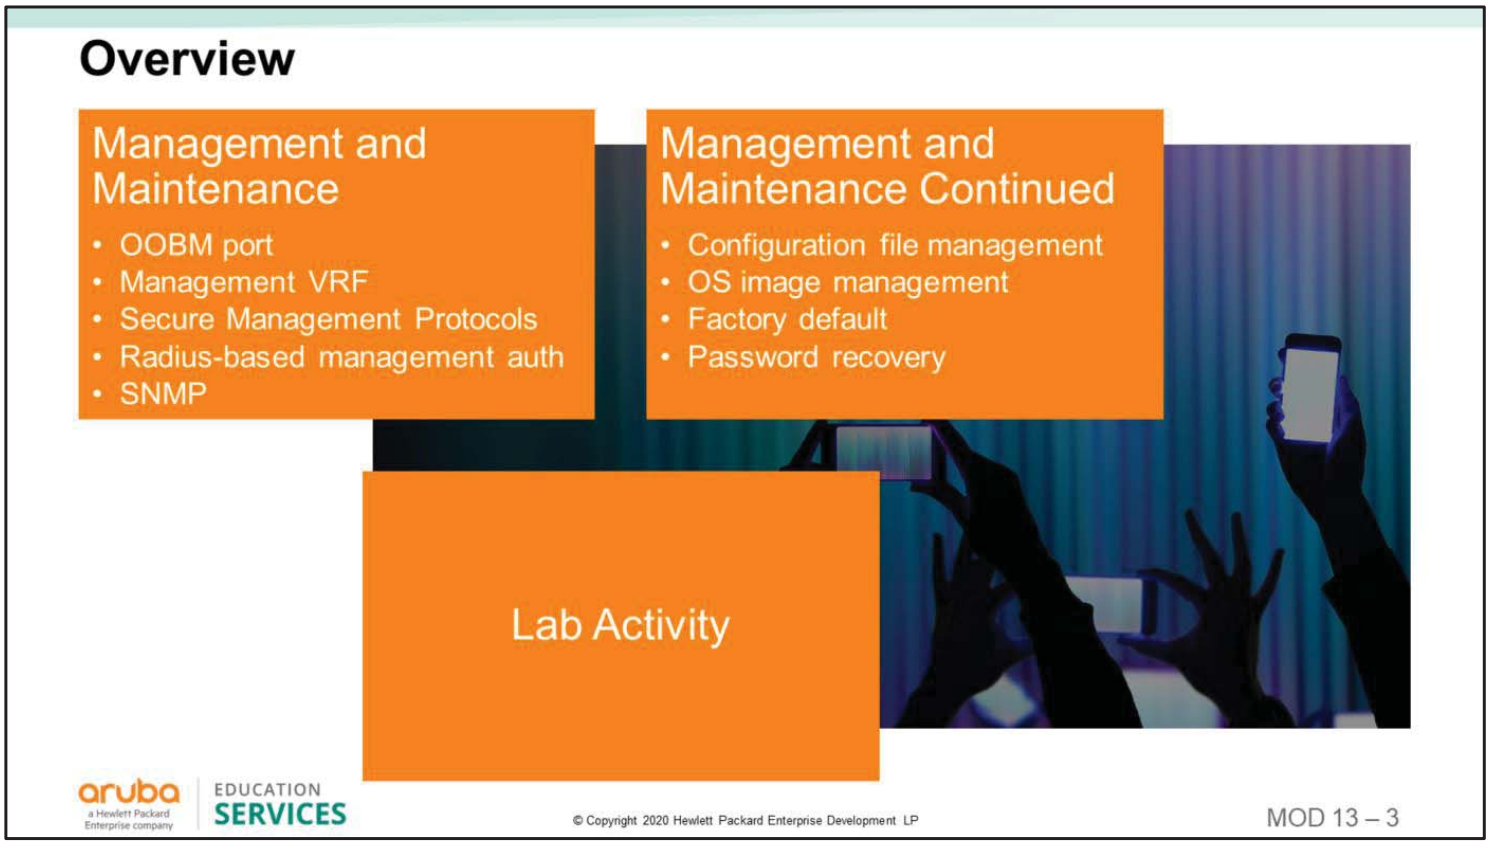

Network management is a vital skill for prospective network administrators and engineers. This module will give you the foundational knowledge to understand and perform the most important network management skills.

First you will learn about how ArubaOS-CX devices have isolated the management and data planes using Virtual Routing and Forwarding (VRF). You see how a separate management VRF supports the OOBM interface – which is purely for management operations. These devices also have a default VRF for typical data plane operations, to support connectivity for end users and other network devices.

You explore secure management protocols, including SSH for CLI access, HTTPS for GUI access, AAA and RADIUS for centralized security, RBAC for efficient, role-based security. You also learn about the SNMP – operation, architecture, and configuration.

The journey continues as you dive into configuration file management and Operating System image management. Then you learn how to restore an ArubaOS-CX device to a factory default state,and perform password recovery.

Management and Maintenance

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

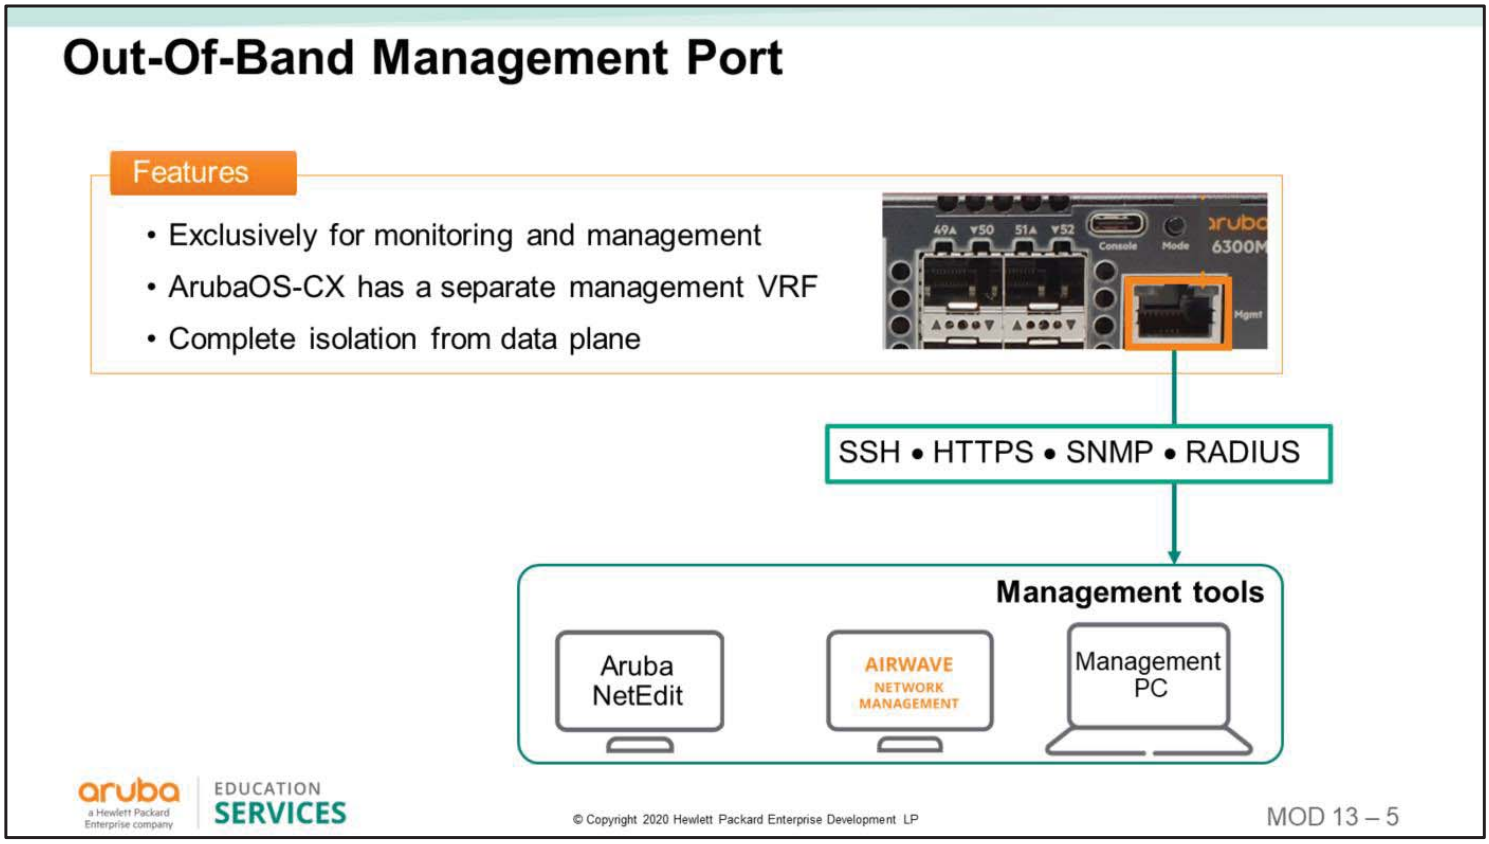

The ArubaOS-CX switch families include an Out Of Band Management (OOBM) port, used exclusively to monitor and manage the switch. Because ArubaOS-CX has a separate management VRF, there is complete isolation between the management and data planes. Data traffic can never see or use this port. Due to this complete isolation, management traffic does not use any bandwidth in the data plane. This is commonly referred as Out-of-Band Management or OOBM.

You do have the option of remotely connecting to and managing the switch via a normal Ethernet port in the data plane. This is referred to as In-Band Management. This can seem convenient but is not recommended. A misconfiguration of switch routing or switching protocol could cause you to lose administrative access to the device. There is also a security concern. Without manually configured security controls, it is possible for hackers to access your switch configuration and take control of the device.

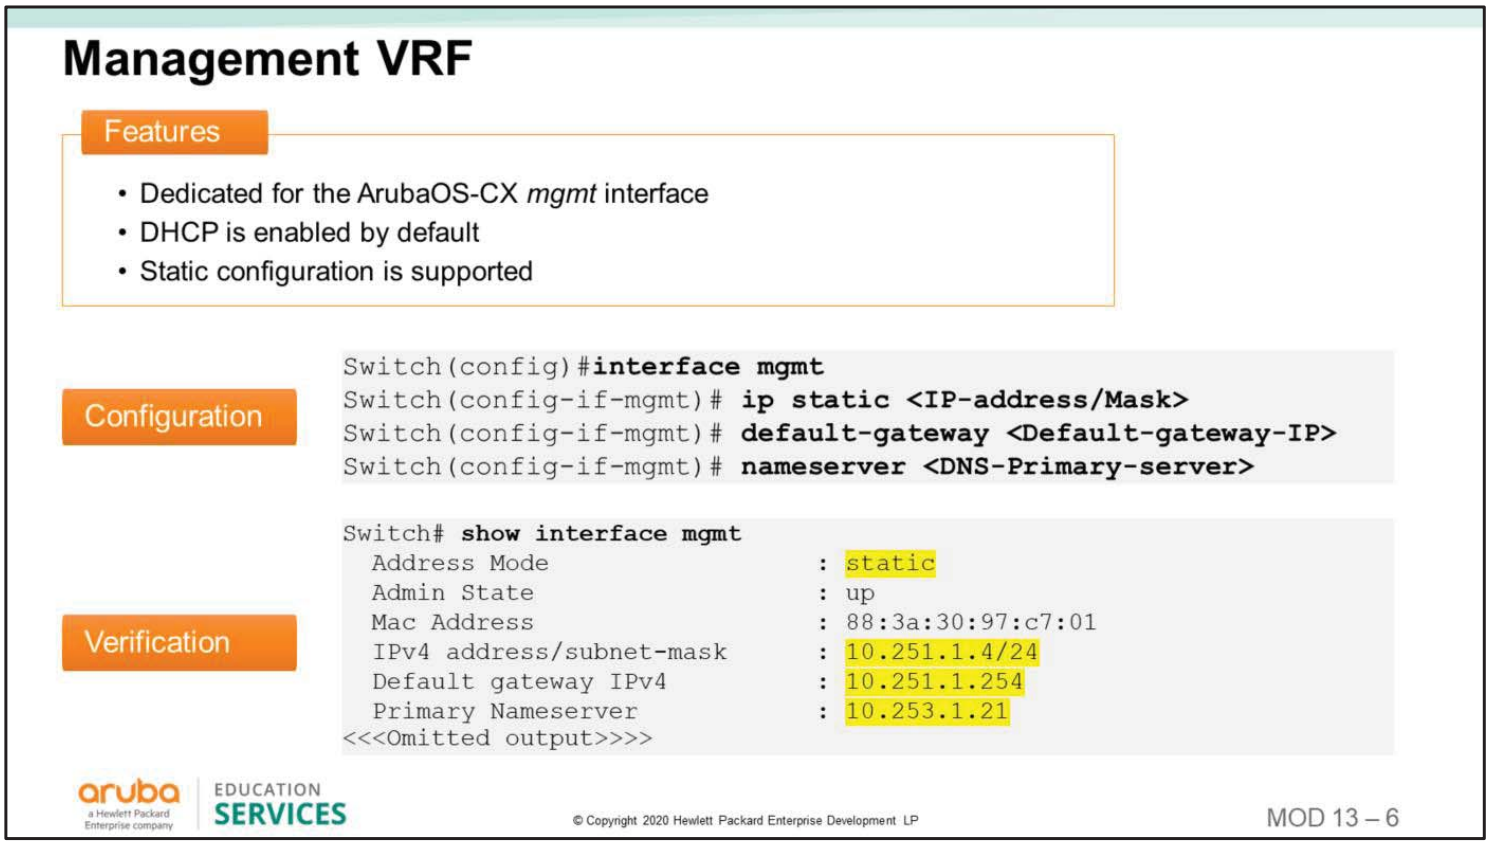

You learned about Virtual Routing and Forwarding (VRF) in a previous module of this course. VRF creates separate virtual routers inside a physical router, with separate routing tables. ArubaOS-CX devices have a **default VRF** for the data plane, and a separate **mgmt VRF** for the Management port to handle OOBM traffic.

The ArubaOS-CX ***mgmt interface*** is enabled by default. It is configured to receive its IP parameters using DHCP. Use the **show interface mgmt** command to verify the state of this interface.

If preferred, you can manually configure IP parameters on the management interface, as shown in the figure and below.

Switch(config)#**interface mgmt**
Switch(config-if-mgmt)# **ip static <IP-address/Mask>**
Switch(config-if-mgmt)# **default-gateway <Default-gateway-IP>**
Switch(config-if-mgmt)# **nameserver <DNS-Primary-server>**

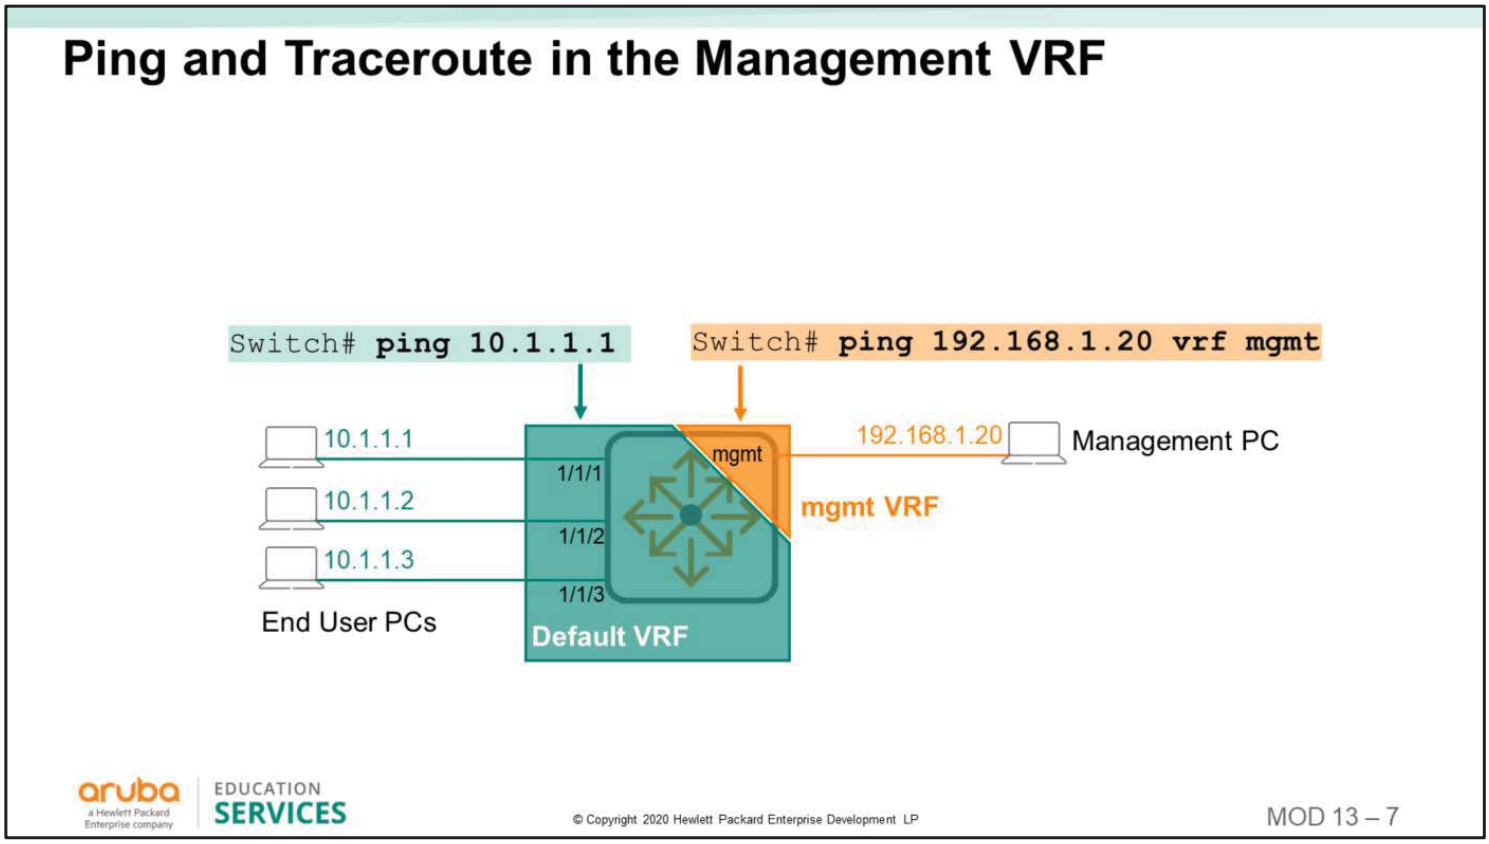

By default, ArubaOS-CX switch ping and traceroute commands are generated for the data plane's **default VRF**. To verify connectivity on the **mgmt VRF**, you must use special ping and traceroute syntax to specify the mgmt VRF, as shown in the figure.

SSH for ArubaOS-CX

Objective Secure protocol for switch CLI access

ArubaOS-CX 6300 series
- Enabled by default on and **mgmt** and **default** VRF

ArubaOS-CX non 6300 series
- Enabled by default on **mgmt** VRF

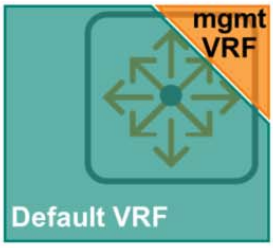


```
8325(config)# ssh server vrf default
8325# show ssh server vrf default

SSH server configuration on VRF default :

    IP Version         : IPv4 and IPv6    SSH Version           : 2.0
    TCP Port           : 22               Grace Timeout (sec)   : 60
    Max Auth Attempts  : 6
    Host-keys          : ECDSA, ED25519, RSA
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 8

You must use the Secure Shell (SSH) protocol to connect to the ArubaOS-CX switch CLI. SSH provides secure communications between the switch and your management PC. SSH is enabled by default in the data plane's default VRF and disabled in the management plane. Use the syntax shown in the figure to enable SSH for VRF mgmt.

Switch(config)# **ssh server vrf mgmt**
Use the **show ssh server vrf mgmt** command to check SSH status on the mgmt interface.

Once SSH is enabled in the management plane, many administrators prefer to disable SSH in the data plane. Although SSH is secure and requires proper authentication credentials, disabling SSH in the data plane ensures that SSH connectivity is not possible for end users and potential bad actors.

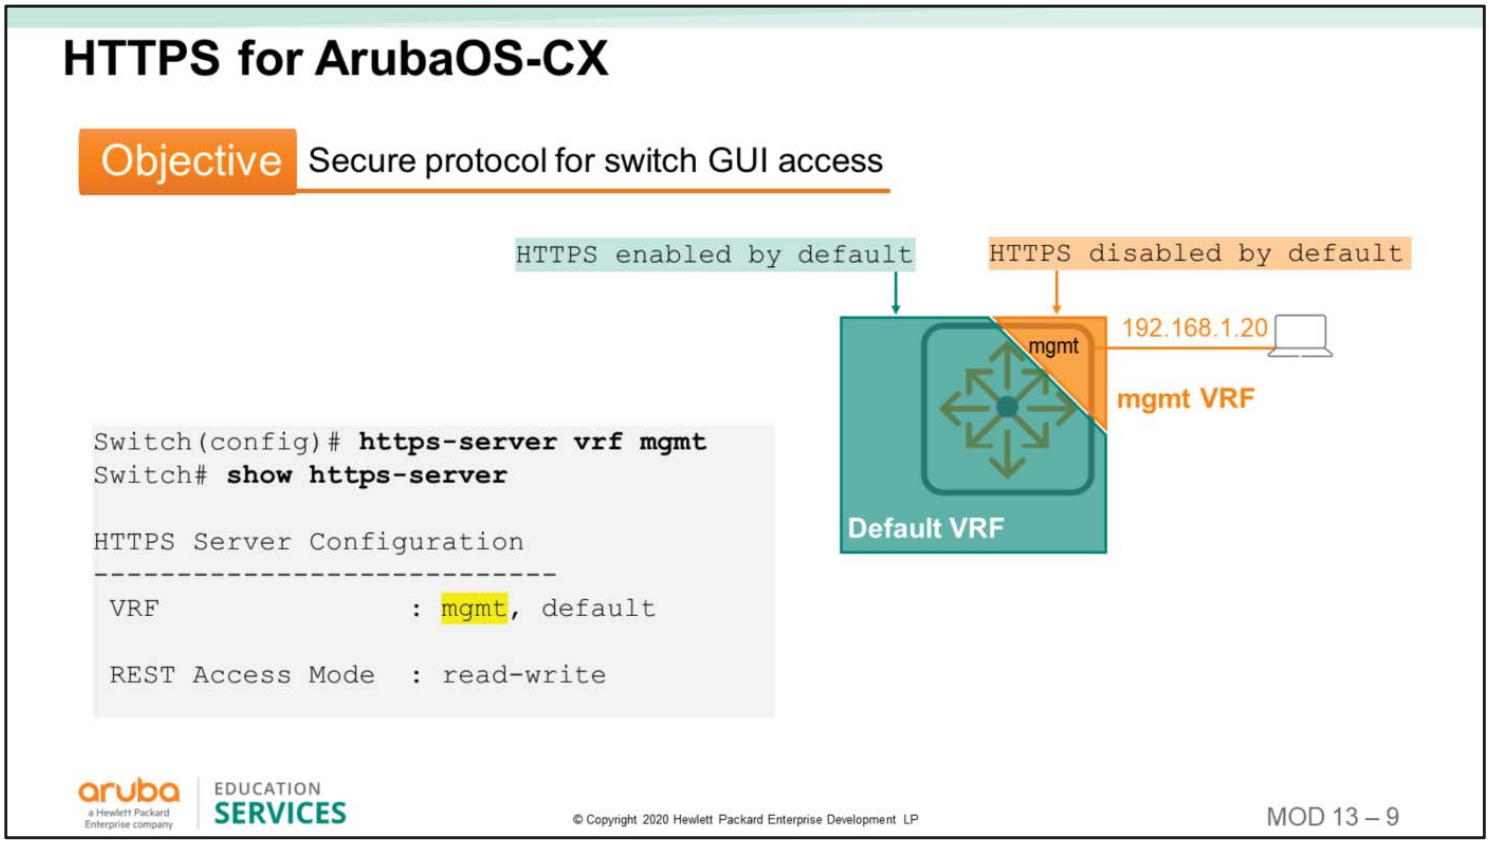

HTTPS provides secure GUI access to the switch. Like SSH, HTTPS is enabled by default in the **default VRF**, and disabled in the management plane. Treat HTTPS like SSH – enable it in the management plane using the syntax shown in the figure. Then disable it in the data plane's default VRF to maximize security.

Switch(config)# **https-server vrf mgmt**

Use the **show https-server** command to verify HTTPS status.

To disable the HTTPS service in the VRF default:

Switch(config)# **no https-server vrf default**

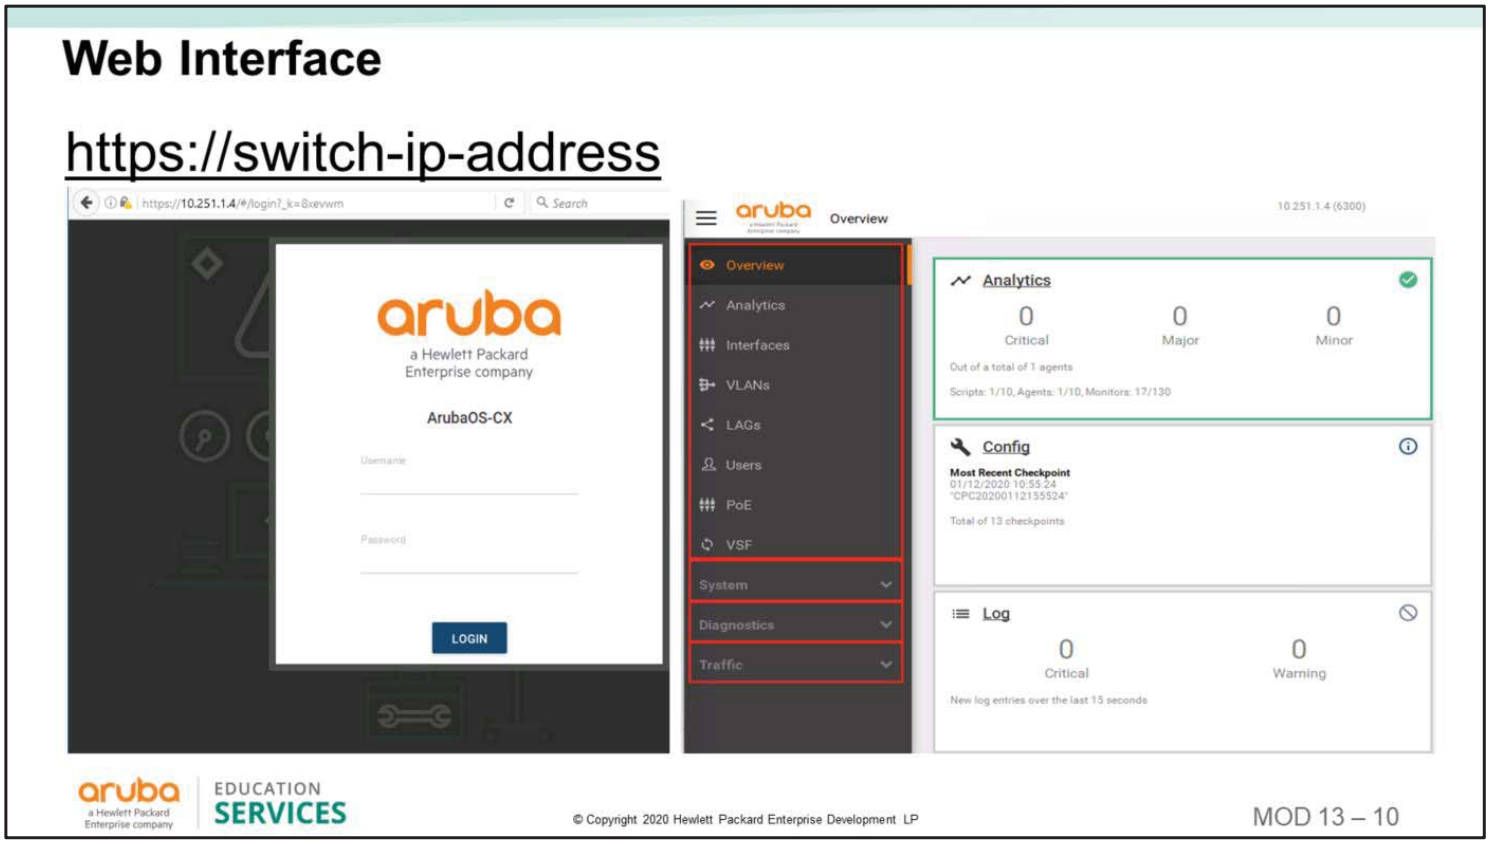

ArubaOS-CX switches integrate a GUI web interface for device monitoring and management. To use this interface, open a web browser and enter the switch IP address: https://switch-ip-address

The menu on the left is divided into four different sections:

Overview. Displays basic protocols and feature information. Those include Analytics, Interfaces, VLANs, LAGs, Users, Power over Ethernet and VSF.
System. Displays Environmental information, Log, DNS servers, SNMP configuration and Configuration management.
Diagnostics. Include features like Ping and Traceroute.
Traffic. This menu includes information about Spanning-Tree

Remember that it is a best practice to permit HTTPS/GUI use only in the **mgmt** VRF.

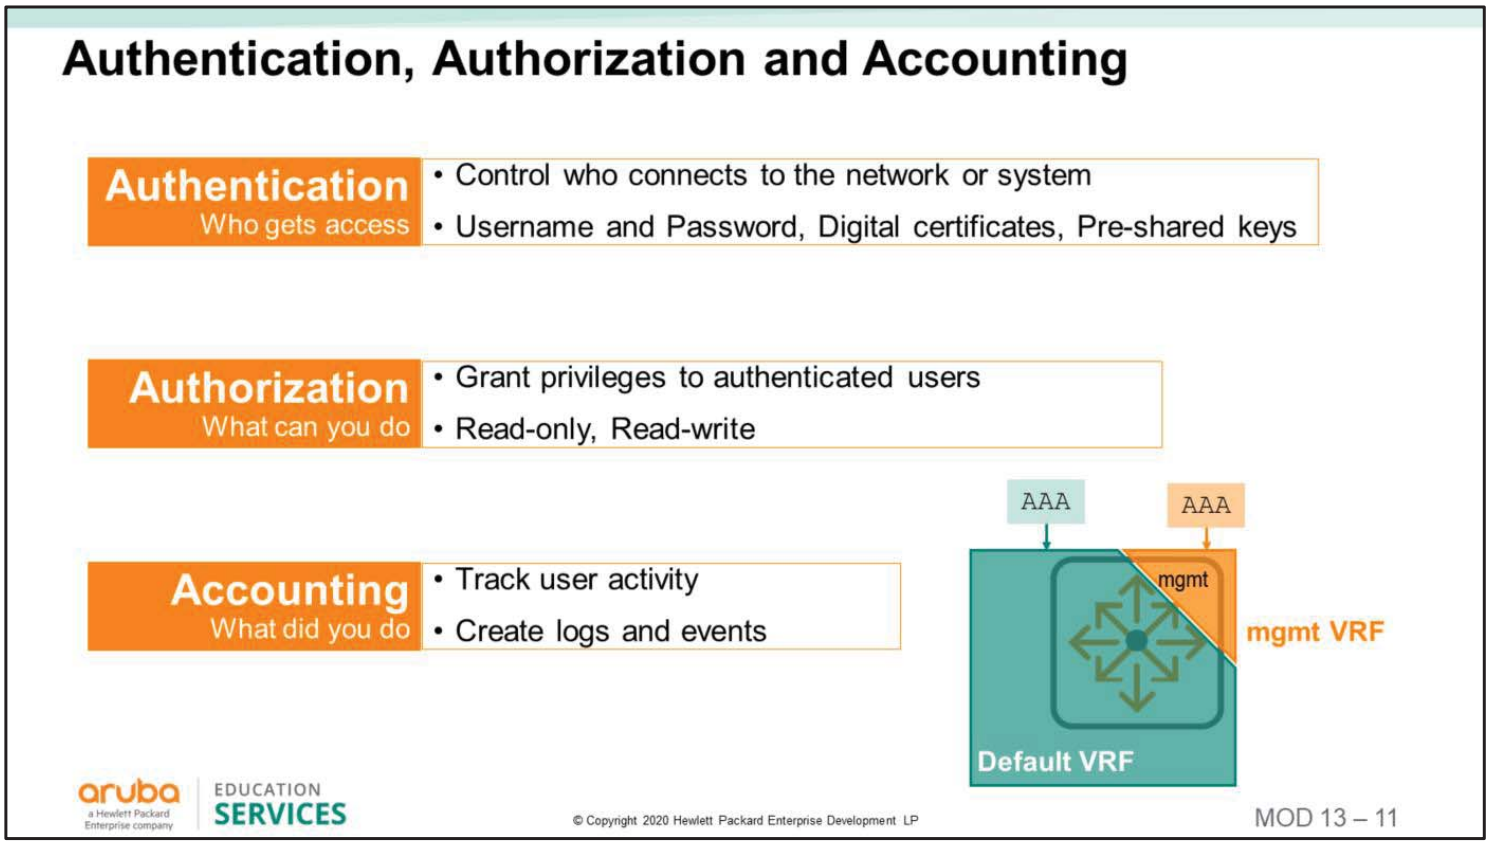

AAA is a security concept that involves three individual processes:

- **Authentication controls who gets access**. The system validates credentials as users attempt data plane connectivity to the default VRF, or management connectivity to the mgmt. VRF. The credentials can be in the form of a usernames and password, digital certificate, or Pre-Shared Keys (PSK).
- **Authorization control what you can do upon authentication** - this process grants privileges to authenticated users. Can you connect to the network and use resources? Can you have read-only access to the switch CLI or GUI? Can you have full read/write access to the switch CLI? This depends on your level of authorization.
- **Accounting tracks what you did by creating event logs –** Jose Rodriquez accessed his sales files at 3:30pm, John Smith deleted accounting files at 2:00AM, Mary Owens modified the switch CLI last Thursday at 12:45.

AAA is often implemented with a Remote Authentication Dial-In User Service (RADIUS) server, about which you will soon learn.

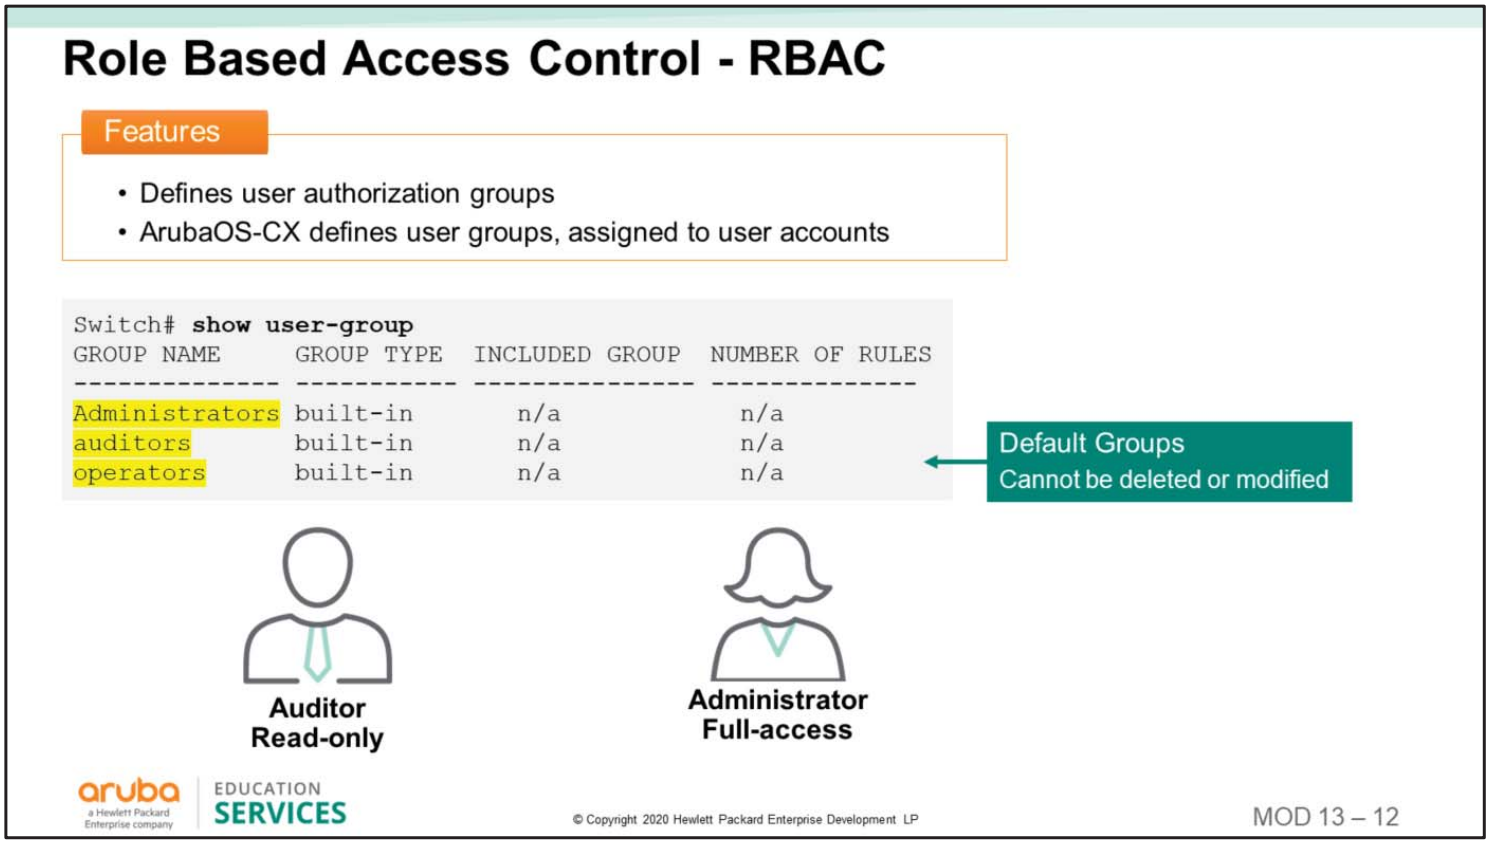

Administrators can access and configure the switch. They have full visibility of all switch processes. This is perfect for well-trained, trustworthy employees, but others should be restricted based on their expertise and job function. This concept is known as Role Based Access Control or RBAC.

ArubaOS-CX applies RBAC defining **user-groups**, where each group is manually configured with specifically allowed switch commands. ArubaOS-CX defines three default user-groups: administrators, auditors and operators, which can neither be deleted nor edited.

Use the **show user-group** command to verify this information.

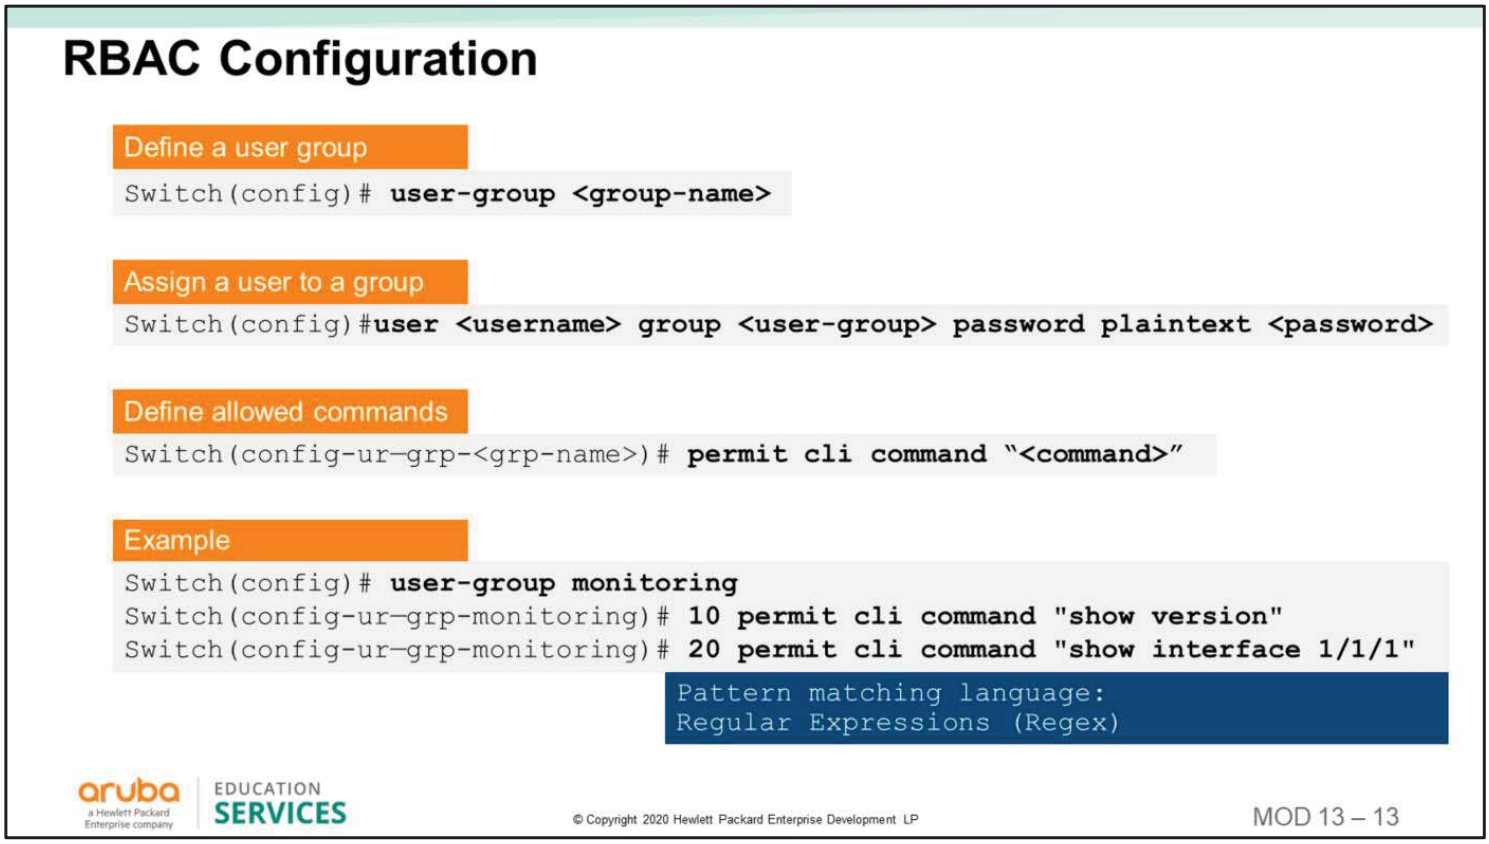

Define new user-groups as shown below:

Switch(config)# **user-group <group-name>**

Map a user account to a defined group as shown below:

Switch(config)# **user <username> group <user-group> password plaintext <password>**

Use rules to define the commands that will be available, as shown below:.

Switch(config)# **user-group monitoring**
Switch(config-usr-grp-monitoring)# **permit cli command "<command>"**

In the example below, a group named monitoring can only use the commands **show version** and **show interface 1/1/1.**

user-group monitoring
10 permit cli command "show version"
20 permit cli command "show interface 1/1/1"

ArubaOS-CX accepts complete command syntax, as shown above and accepts Regular Expression (Regex) syntax. This is a set of variables and logic symbols that let you specify a group of commands with less typing.

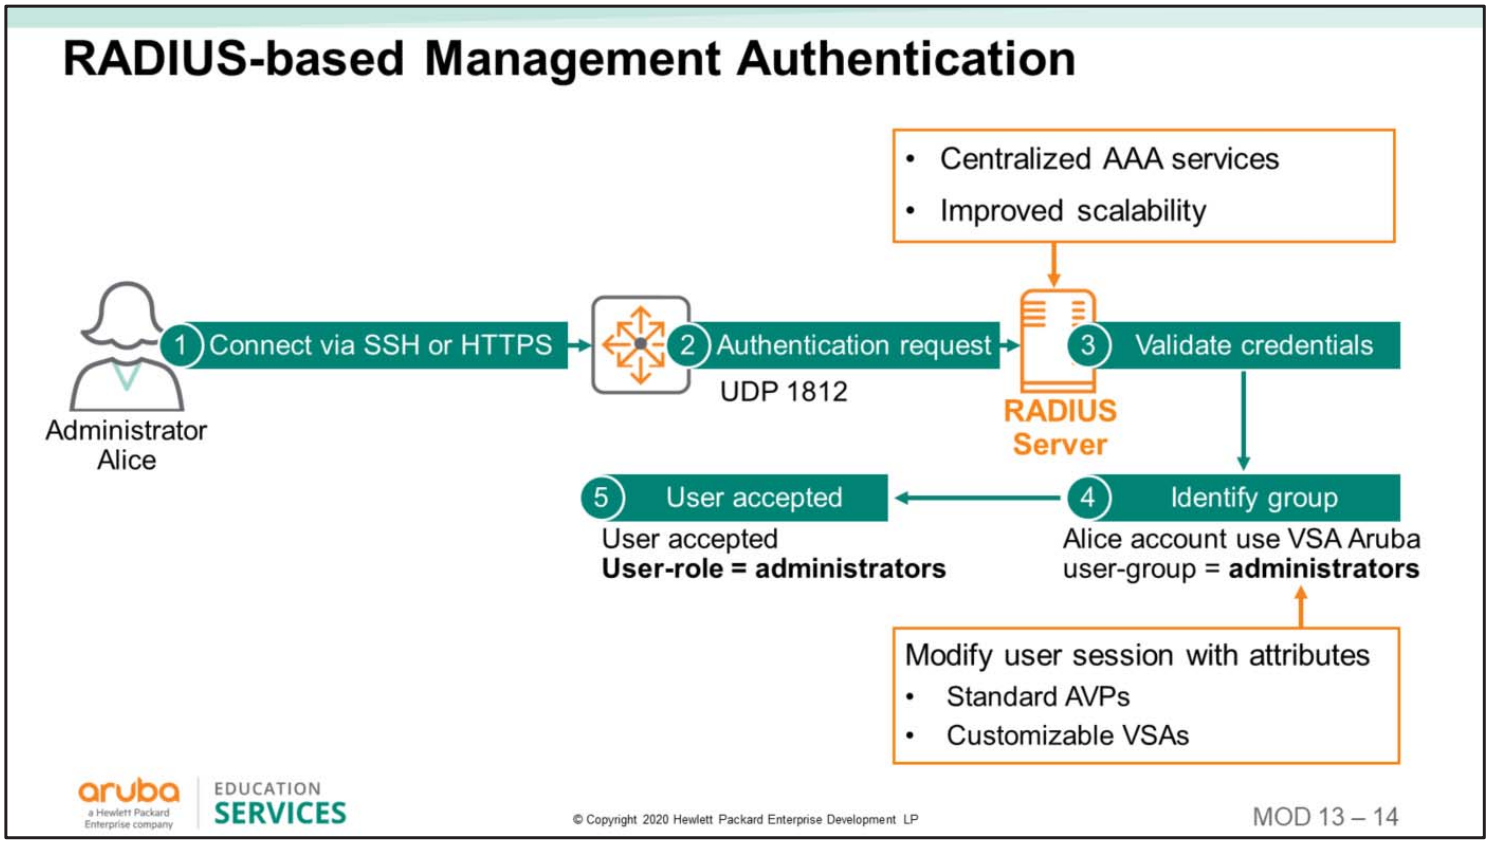

For authentication, some administrators define local credentials on each switch. This is good for smaller environments but becomes a burden as the network grows. If you have 100 network devices, you must maintain 100 separate credential databases, and try to keep them consistent.

Many people prefer to centralize AAA services on a Remote Authentication Dial-In User Service (RADIUS) server. Hundreds and thousands of devices can use the RADIUS protocol (UDP port 1812) to access a centralized RADIUS server – a centralized repository for credentials.

A user or administrator asks the switch if it can connect, and the switch in turn asks the RADIUS server, “Hey, I have user Alice, with password Secret123. Should I grant them access?”
The RADIUS server receives this request, validates credentials against its database, and responds to the switch with an ACCEPT message or a REJECT message. RADIUS messages can also include attributes. These are like “notes” that can modify how a connection session shall proceed.

For example, user Alice attempts to login into a switch, which sends Alice’s credentials to the RADIUS server. The server validates these credentials and sends an ACCEPT message, which can include an attribute that helps the ArubaOS-CX switch to place Alice in the Administrator user-group.

One reason for the RADIUS standard’s popularity is that it was designed to be customizable by equipment manufacturers and vendors, using Vendor Specific Attributes (VSA). Aruba has created VSAs that are specific for its products, to enhance usability and

features. By default RADIUS servers use standard attributes called Attribute Value Pairs (AVP). Custom VSAs must be manually added to the RADIUS server. Aruba ClearPass is a powerful RADIUS server that uses Aruba VSAs by default, without need for manual addition.

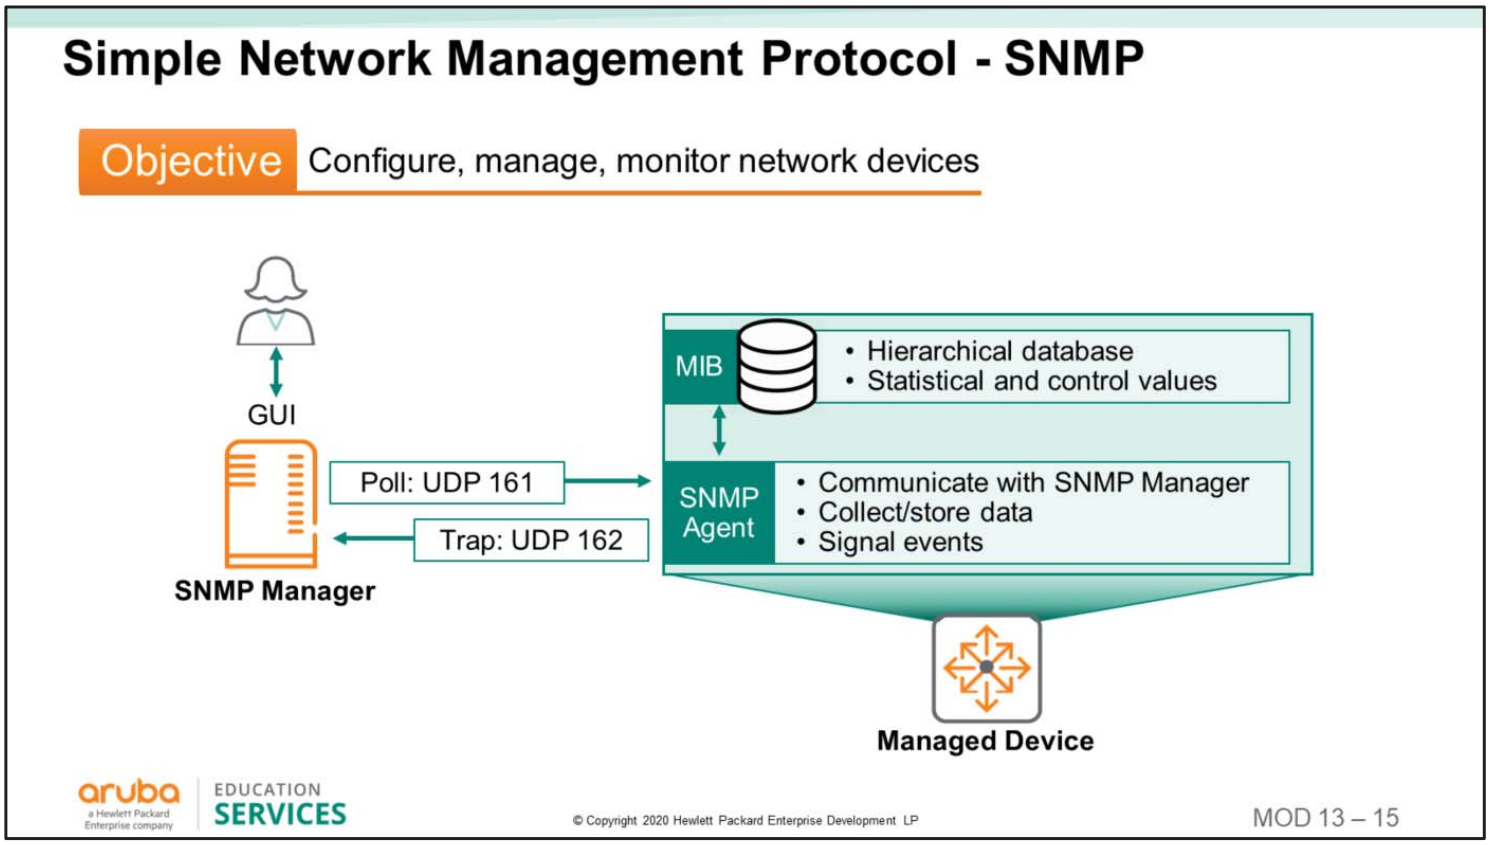

SNMP is an application-layer protocol defined in RFC 1157 for exchanging management information between network devices. This protocol is a popular method to configure, manage, and monitor network elements from any vendor. The figure shows the SNMP architecture, as described below:

SNMP Manager. Is typically a dedicated server that communicates with SNMP agent. It often has a GUI interface, so administrators can view and change parameters for many network device from a single interface. The manager uses several message types to gather and modify network device parameters, Some of its key functions include:

- Queries agents
- Gets responses from agents
- Sets variables in agents
- Acknowledges asynchronous events from agents

•

Managed devices. The routers, switches, firewalls, wireless controllers, access points and any other network devices that should be centrally monitored, managed, and configured. Internally, these managed devices have two main SNMP components.

1. The **SNMP Agent** collects information and communicates with the SNMP Manager about device configuration and status. It stores and retrieves this information in a Management Information Base (MIB).

2. The **MIB** is a hierarchical database that stores device status, statistics, and configuration information. This is the information that is exchanged between the SNMP Manager and the managed device.

Communication and ports

Usually, SNMP Manager to Agent communication is a “pull” mechanism. When this device requires to get information from the managed device it requests the information using simple commands such as GET, GET NEXT, GET BULK. Changes in the device are done using the SET command. This communication uses UDP 161.

The exception to the “pull” mechanism is when the managed device has important information that cannot wait for the manager’s period polling. “Newsflash!! My CPU utilization is dangerously high”. In cases like this, the managed device “pushes” information to the manager without being asked, using an SNMP Trap. SNMP Traps use UDP 162.

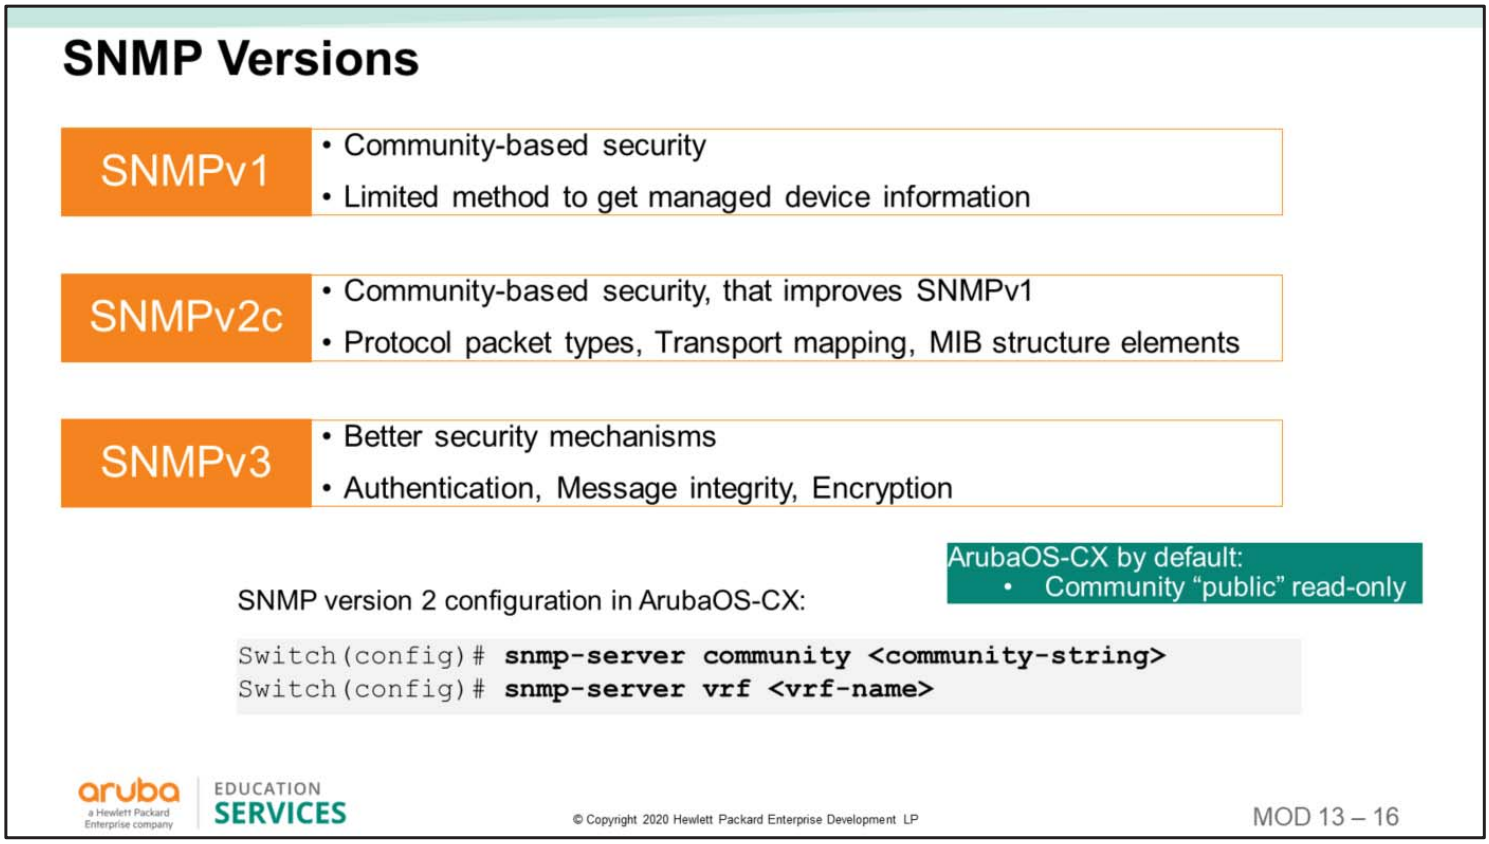

The figure summarizes SNMP versions, as described below.

SNMP version 1. First version of the SNMP protocol and uses a community-based security mechanism. The "community string" is like a simple pre-shared passcode. If the Managed device has a community string of **public,** then any SNMP manager that can reach the device can access the agent's MIB, as long as **community string = public** is used.

SNMP version 2c. This is the revised protocol, which includes enhancements of SNMPv1 in the areas of protocol packet types, transport mapping, MIB structure elements. This version also uses community-based security mechanism.

SNMP version 3. Provides a higher level of security, including message integrity, authentication, and encryption.

A simple configuration using SNMP version 2 in ArubaOS-CX is presented in the following lines:

Switch(config)# **snmp-server community <community-string>**
Switch(config)# **snmp-server vrf <vrf-name>**

Note. By default SNMP read-only community is set to **public** in ArubaOS-CX. Write-only community is not set, therefore not accessible – ArubaOS-CX switches are read-only by default.

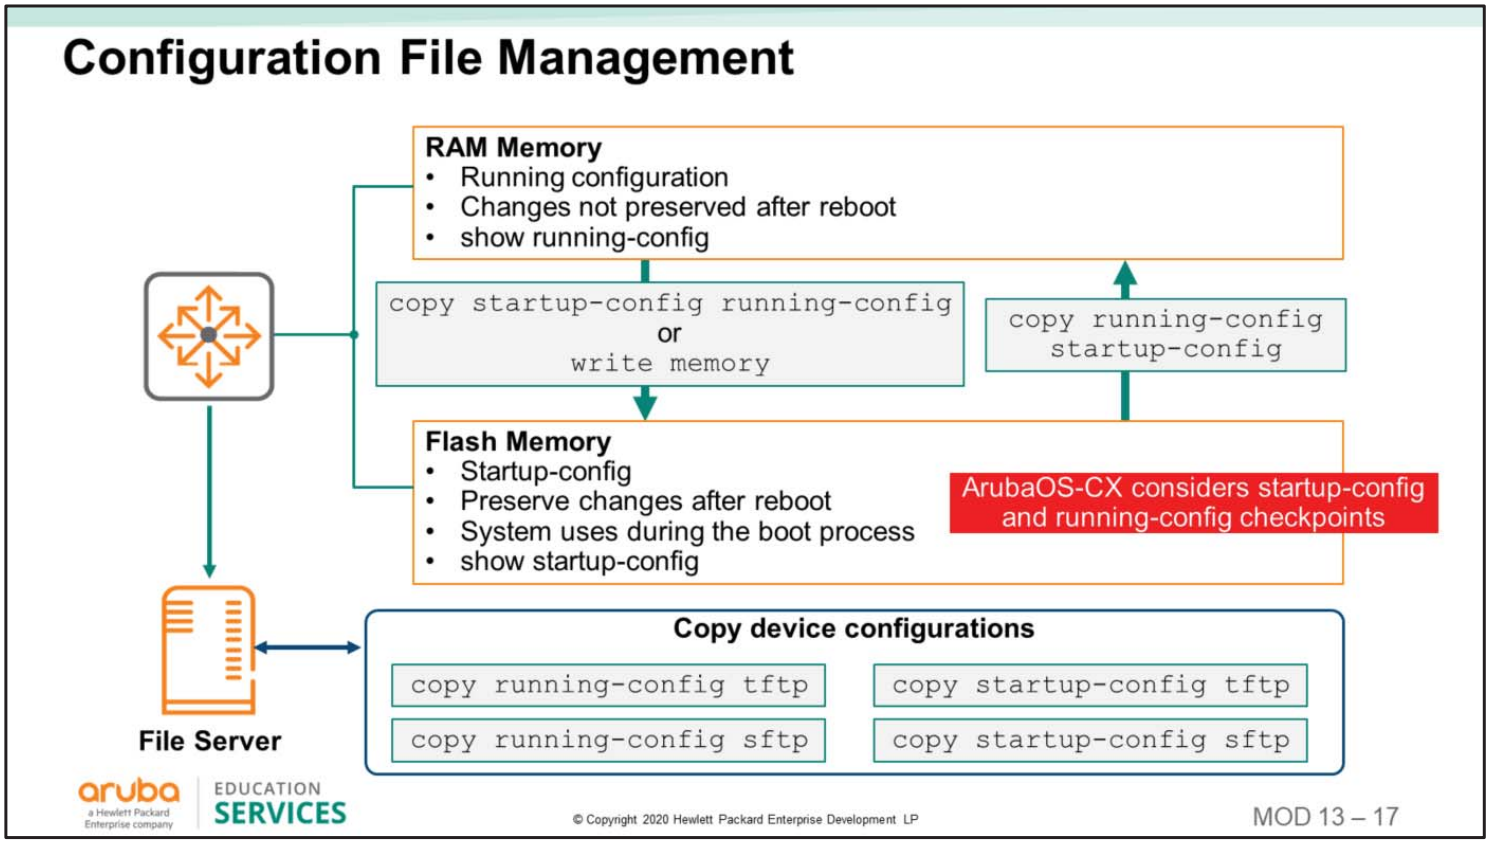

An ArubaOS-CX switch uses RAM and Flash memory to save configuration files. You should be aware of RAM and Flash memory characteristics, as they relate to the storage of device configuration files.

RAM memory is used to save the configuration that is currently running on the device – the running-config. Remember that RAM is not permanent storage – the contents of RAM are lost when a device reboots or powers down. If you modify the running-config and the switch reboots, your changes are lost. Use the command **show running-config** to view this file.

Flash memory is used to permanently save a configuration file. The contents of flash memory survives during reboot or power off. Use the **show startup-config** command to verify the contents on this file.

Of course, any time you make changes to the running-config, you should copy those changes to the startup-config using the copy command. To permanently save the configuration that is running in the device, use the command **copy running-config startup-config.**

Note. Alternatively you can use the command **write memory** which does the exact same task.

If you need to restore a configuration from the startup-config use the command **copy startup-config running-config.**

As a best practice, you should export the running and the startup configuration files to an external file server. This will help you to recover from catastrophic device failures. To

accomplish this task you can use the **copy** command, as shown below.

Switch**# copy {startup-config | running-config} {tftp:///|sftp:///USER@}{IP|HOST}[:PORT][;blocksize=VAL]/FILE**

For example to copy the running-config to a TFTP server with IP address 10.253.1.21, use the command **copy running-config tftp://10.253.1.21/switch_config.cfg**

ArubaOS-CX switches also allow you to copy files to a USB flash drive, which can be directly connected to the switch. To copy your running or startup-config to a USB device, use the following syntax:

Switch# **copy {running-config | startup-config} usb:/file**

ArubaOS-CX considers startup and running config checkpoints, you will learn about checkpoints in the next slide.

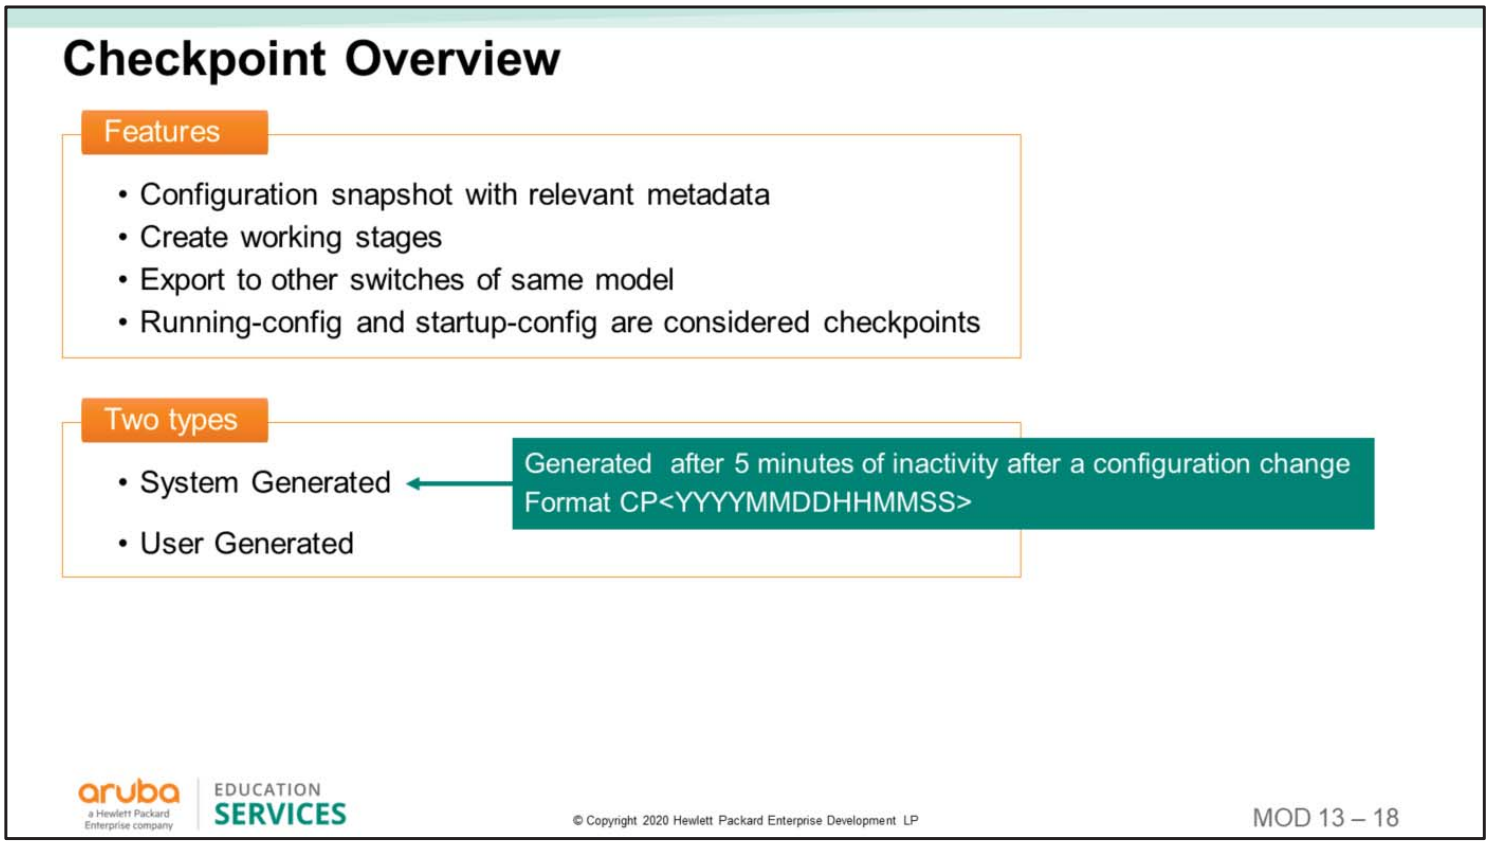

Using the copy commands you just learned allows you to store configuration files and recover lost configurations in case of a power outage or administrative mistake. This is a good thing, but these files do not include any additional data about the state of networking processes. For a true recovery, a new approach is needed.

A checkpoint is a snapshot of the switch running configuration and its relevant metadata during the time of creation. You can use checkpoints to apply the switch configuration stored within a checkpoint whenever needed, such as revert to a previous working configuration. A checkpoint is flexible since it can be applied to other switches of the same model.

There are two types of checkpoints:

- System generated checkpoints are generated after 5 minutes of inactivity after a configuration change. By default a checkpoint is taken every 300 seconds or 5 minutes.
  - The format that this checkpoint type uses is **CP<YYYYMMDDHHMMSS>.** For example **CPC20200111000038**
- User generated checkpoints are manually generated by the administrator.

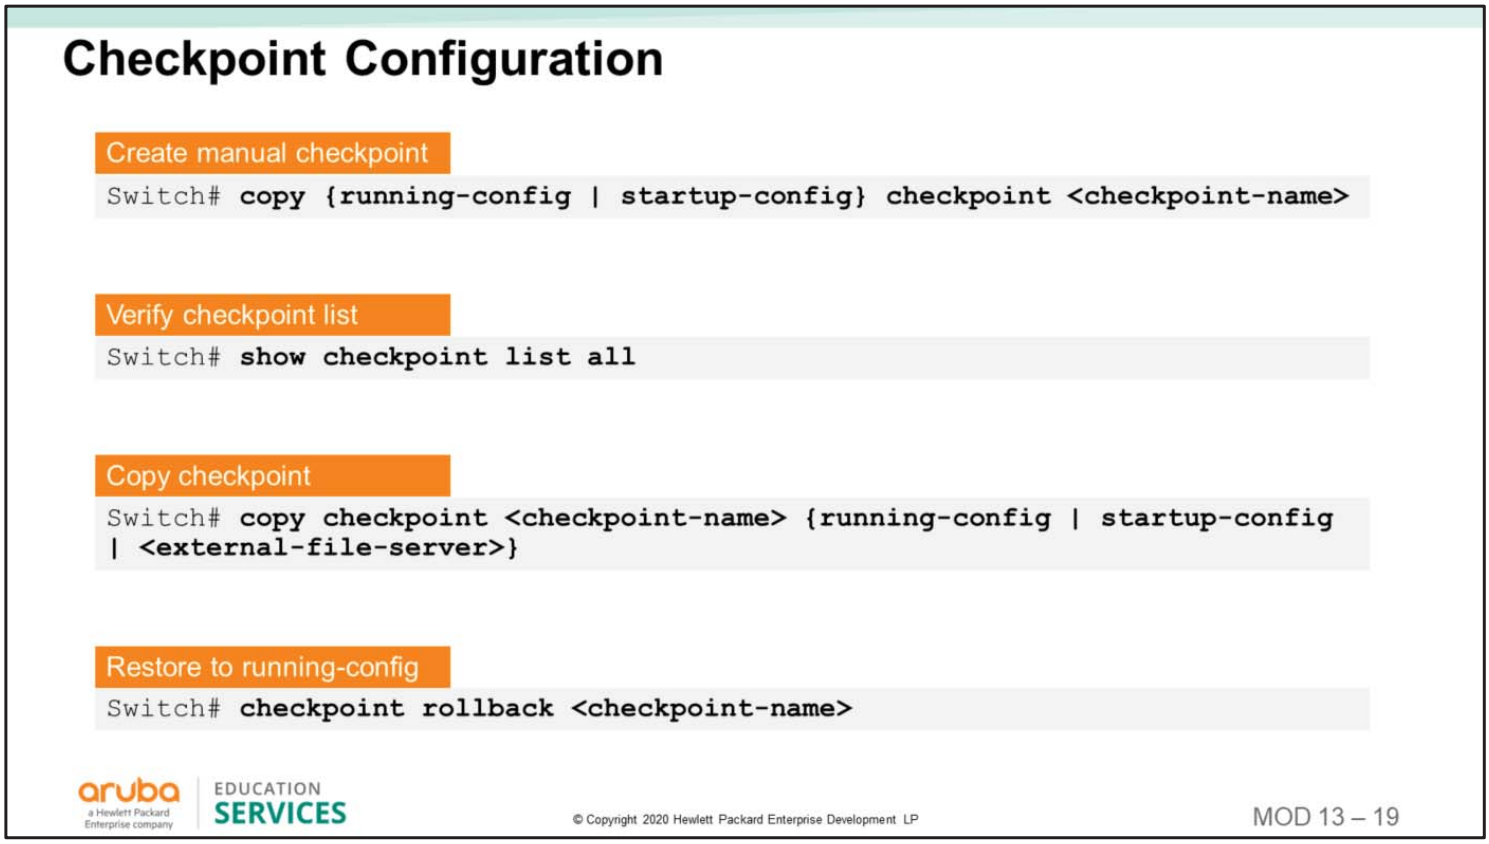

To create a checkpoint use the copy command.

Switch# **copy {running-config | startup-config} checkpoint <checkpoint-name>**

To verify the checkpoint list, use the command show checkpoint list all.

To revert the switch from a checkpoint use the following command.

Switch# **copy checkpoint <checkpoint-name> {running-config | startup-config}**

Note. You can export a checkpoint file to an external server.

Alternatively, you can use the **checkpoint rollback** command. This command only restores the configuration to the running-config file.

Switch# **checkpoint rollback <checkpoint-name>**

Note. checkpoint rollback command does the same task as **copy checkpoint running-config** command.

Checkpoint Auto Mode

1. Administrator enables checkpoint auto mode.

```
Switch# checkpoint auto <minutes>
```

2. The system generates a checkpoint and starts a timer.
3. Administrator configures the device.
4. Administrator must confirm the changes to be permanent.

```
Switch# checkpoint auto confirm
```

If no confirmation received, system rolls back to checkpoint.

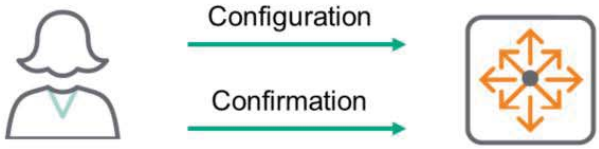

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 20

There are situations when applying a configuration could lock out your access to the switch. This leaves no other option but to locally connect to the device using the console port and remove the offending configuration lines.

The checkpoint auto mode feature addresses this problem in an elegant way. The solution includes a timer that you set. The system then creates a checkpoint with the current state, You then make your desired configuration changes, which the system applies, after a confirmation. If the configuration lines did not affect your access, you will be able to confirm the change. If your changes locked you out of the device, confirmation does not occur. The switch thus automatically reverts to the latest snapshot after the time expires.

To use this mode issue the following command.

Switch# **checkpoint auto <minutes>**
Switch# **checkpoint auto confirm**

Operating System Image Management Introduction

2 partitions:
- **Primary**
- **Secondary**

```
Switch# show images
---------------------------------------------------------------
ArubaOS-CX Primary Image
---------------------------------------------------------------
Version : GL.10.04.0003
Size    : 415 MB
Date    : 2019-10-31 12:33:52 PDT
SHA-256 : 794336ee0b1f941aac421bec5333f16b4ebbcd1f4035a388d3b79aa
---------------------------------------------------------------
ArubaOS-CX Secondary Image
---------------------------------------------------------------
<<<<<Omitted output>>>>>>>

Default Image : primary
-------------------------------------------------------
Management Module 1/1 (Active)
-------------------------------------------------------
Active Image       : primary
Service OS Version : GL.01.05.0002
BIOS Version       : GL-01-0013
```

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 21

ArubaOS-CX switches have two flash memory modules or partitions for storing switch software image files. AOS-CX refers to those modules as primary and secondary images.

Having two flash memory locations helps to manage the switch more easily, and to plan firmware upgrades. For example, you can download a new firmware code in the flash memory that is not currently in use and then instruct the switch to boot from it during a maintenance window.

Show images. This command shows information about the software in the primary and secondary images.

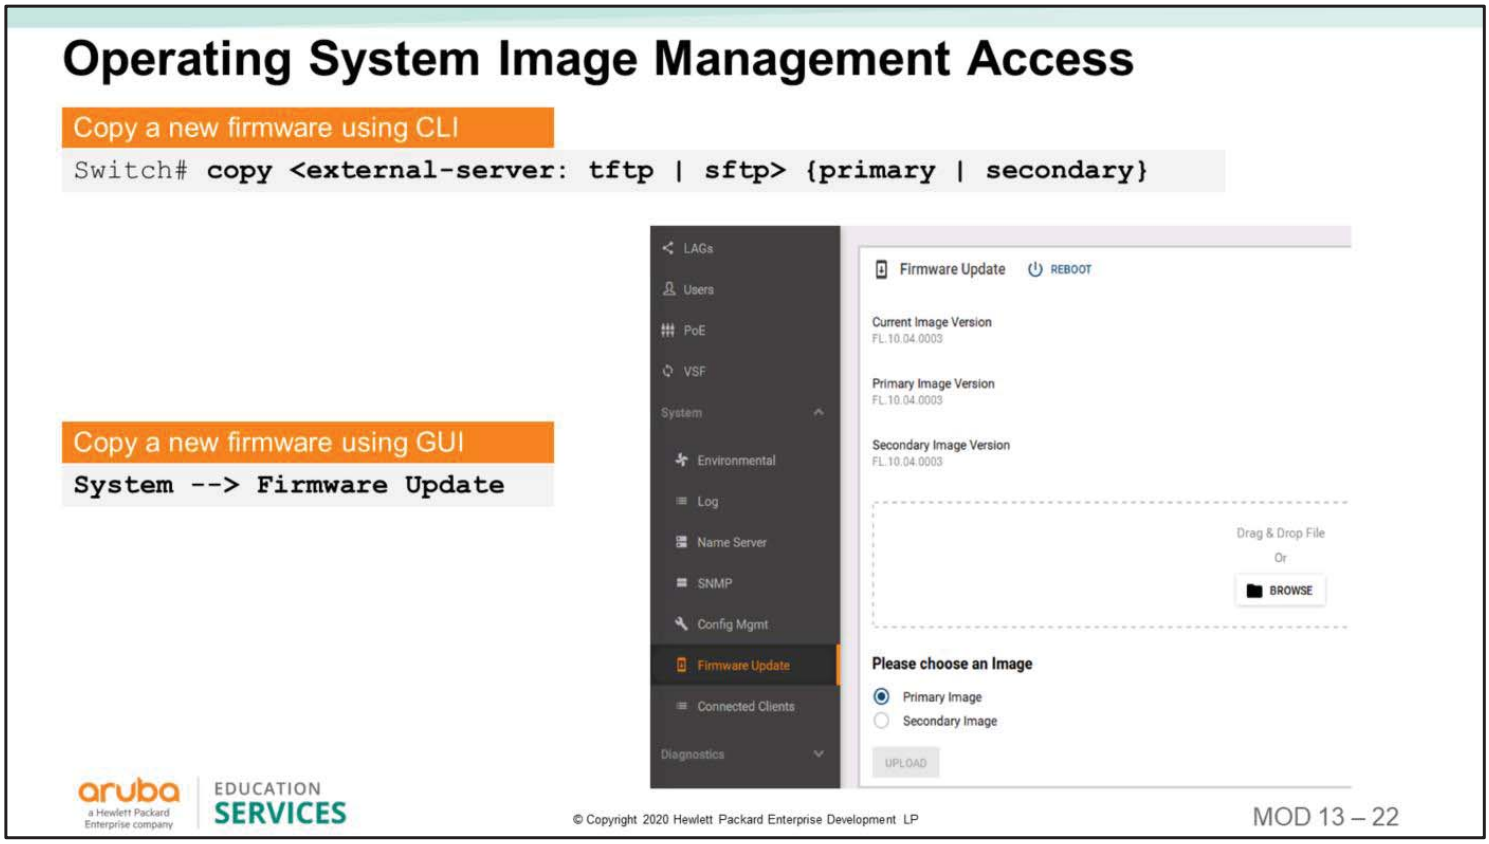

You can update the switch using the CLI or the GUI:

Using the CLI issue the **copy** command.
Switch# **copy {tftp:///|sftp://USER@}{IP|HOST}[:PORT][;blocksize=VAL]/FILE {primary | secondary}**

Suppose that you want to copy the image file GL.10.04.0003.swi from the TFTP server with IP address 10.253.1.21 to the Primary partition. Use the following command.

Switch# **copy tftp://10.253.1.23/GL.10.04.0003.swi primary**

To use Secure FTP with the username admin, use the following command.

Switch# **copy sftp://admin@10.253.1.21/GL.10.04.0003.swi primary**

Using the GUI for this task is simple. Just navigate to **System → Firmware Update** submenu. Then browse for the file in your local machine, select the flash partition, and click Upload.

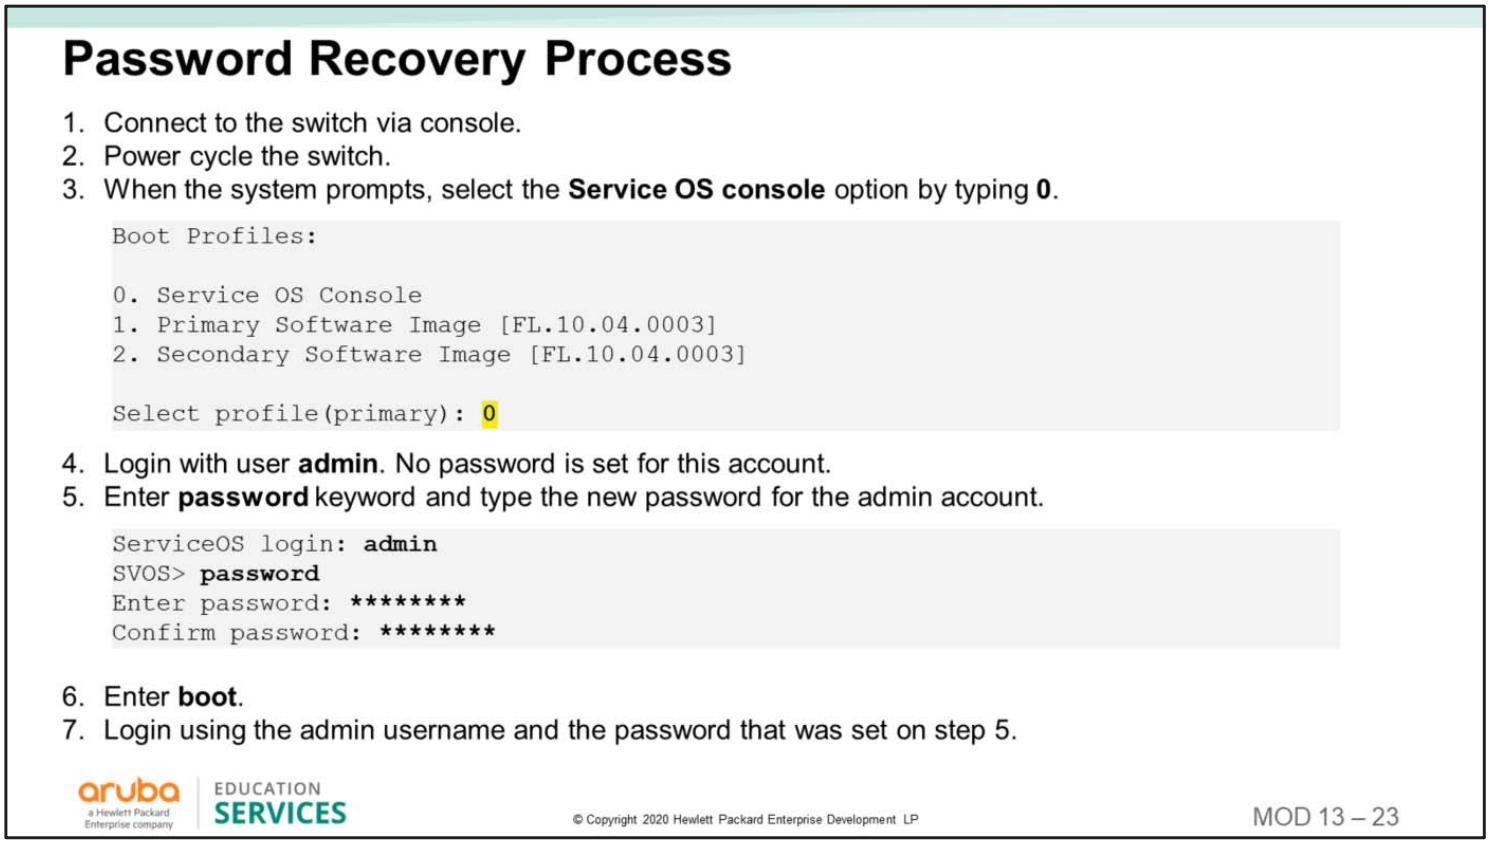

In case that the credentials do not work to login into the switch a password recovery process has to be performed.

1. Connect to the switch using the console port.
2. Power cycle the switch.
3. When the system prompts, select the **Service OS console** option by typing **0**.
4. Login with the user **admin**. No password is set for this account.
5. Enter password keyword and type the new password for the admin account.
6. Enter **boot.**
7. Login using the admin username and the password that was set on step 5.

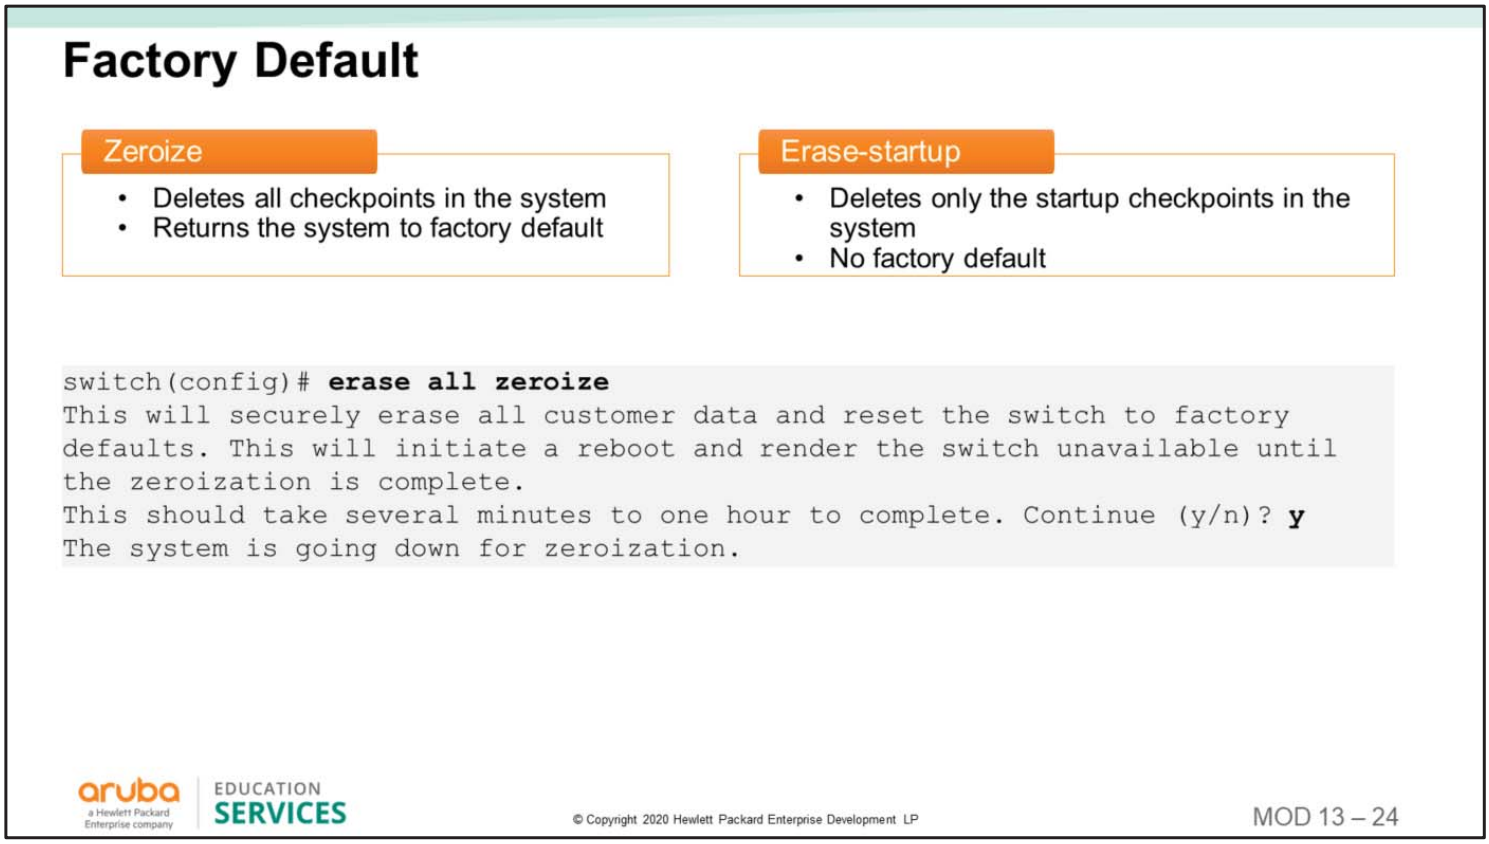

The command **erase all zeroize** restores the switch to its factory default configuration. You will be prompted before the procedure starts. Once complete, the switch will restart from the primary image with factory default settings.

The command **erase-startup** erases the startup checkpoint. Using the zeroize command shown here erases ALL checkpoints.

Question #1

OOBM Port, Management VRF, Ping and Traceroute in the Management VRF

What is true about managing ArubaOS-CX devices?

a. You can connect using SSH, HTTPS, and SNMP on the management interface and on any typical access interface on the switch.

b. The control plane uses a special management VRF.

c. The management interface can acquire an IP address via DHCP or by manual configuration.

d. You can use the command “ping 10.1.1.1” to ping a PC at that address which has been connected to a switch OOBM port.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 25

Correct Answers:
a.
c.

Question #2

SSH for ArubaOS-CX, HTTPS for ArubaOS-CX, Web Interface

Which of the options below accurately describe accessing a 6300 ArubaOS-CX device for CLI or GUI access?

a. By default, SSH is enabled in the management and default VRF.

b. By default HTTPS is enabled in the default VRF.

c. SSH is the preferred method for GUI access.

d. It is a best practice to only enable GUI access on the management VRF.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 26

Correct answer:
d.

Question #3

AAA, RBAC, RADIUS-based Management Authentication

Which of the statements describe valid ArubaOS-CX secure management access concepts?

a. AAA Accounting controls what you can do once you login to the device.

b. AAA services can be used in both the management VRF and the default VRF.

c. Role-Based Access Control (RBAC) is based on a user's group assignment.

d. Advantages of using a RADIUS server include simplicity and faster authentication operations.

e. An easy way to modify several user's access is to modify the "operators" group.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 27

Correct answers:
b.
c.

Question #4

Configuration File Management, Checkpoint Overview, Password Recovery

What is true about ArubaOS-CX configuration file management?

a. It is a good idea to copy the running configuration to the startup configuration after you have made changes to the configuration file that is currently in use.

b. The configuration file that is used while the switch is in operation is typically stored in a flash file system.

c. The advantage of a "Checkpoint" is that it stores operational metadata, in addition to the device configuration.

d. Checkpoints are files – a snapshot of the configuration at a given time

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 13 – 28

Correct answers:
a.
c.
d.

Lab Activity

13.1 Secure Management

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

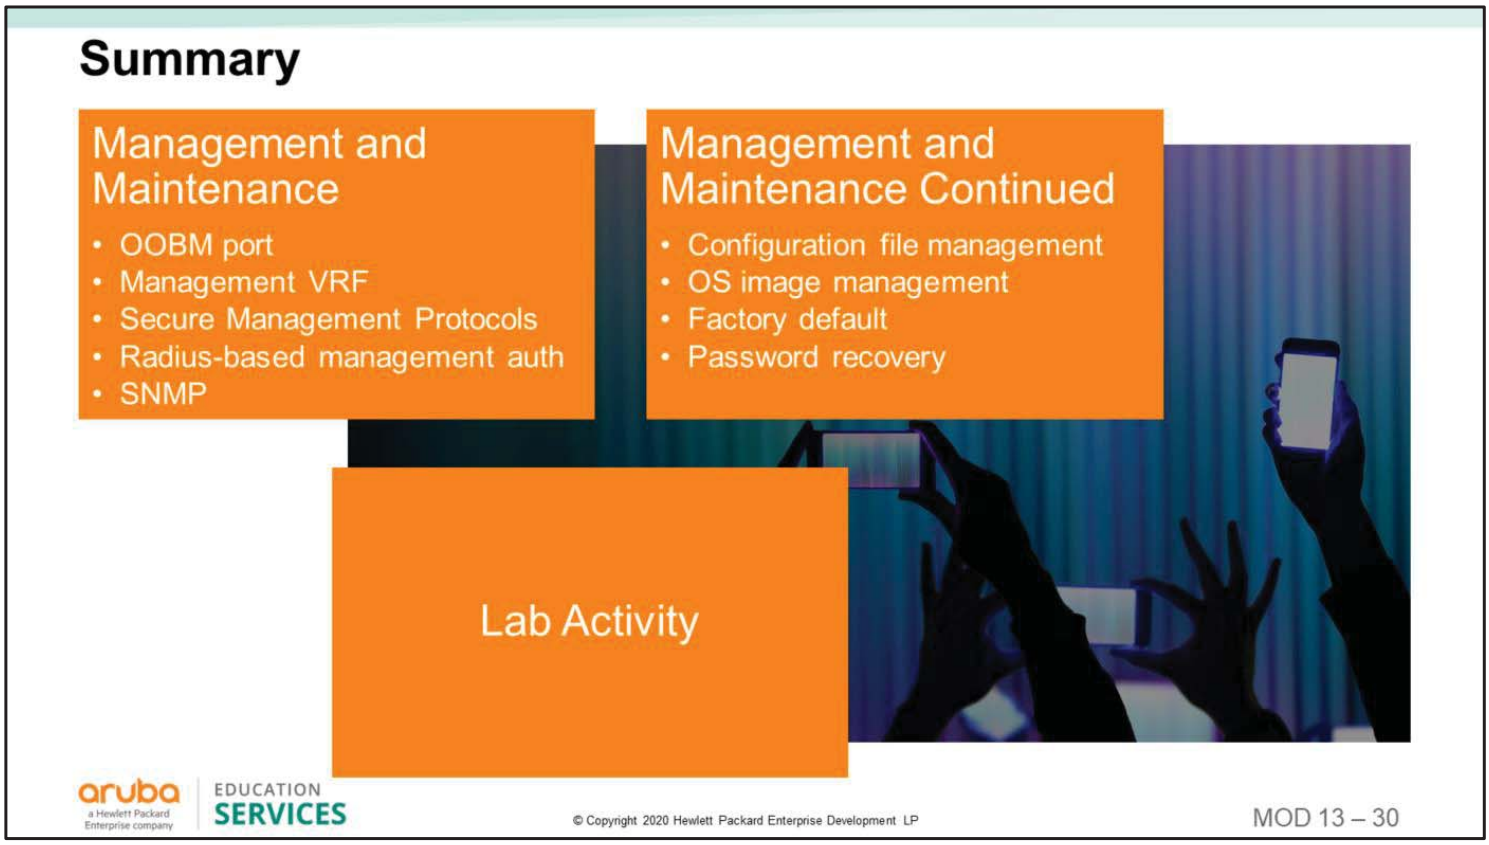

Congratulations! You now understand network management - a vital skill for prospective network administrators and engineers. You acquired the foundational knowledge to understand and perform the most important network management skills.

First you learned about how ArubaOS-CX devices have isolated the management and data planes using Virtual Routing and Forwarding (VRF). You saw how a separate management VRF supports the OOBM interface – which is purely for management operations. These devices also have a default VRF for typical data plane operations, to support connectivity for end users and other network devices.

You explored secure management protocols, including SSH for CLI access, HTTPS for GUI access, AAA and RADIUS for centralized security, RBAC for efficient, role-based security. You also learn about the SNMP – operation, architecture, and configuration.

The journey continued as you learned about configuration file management and Operating System image management. Then you learned how to restore an ArubaOS-CX device to a factory default state and to perform password recovery.

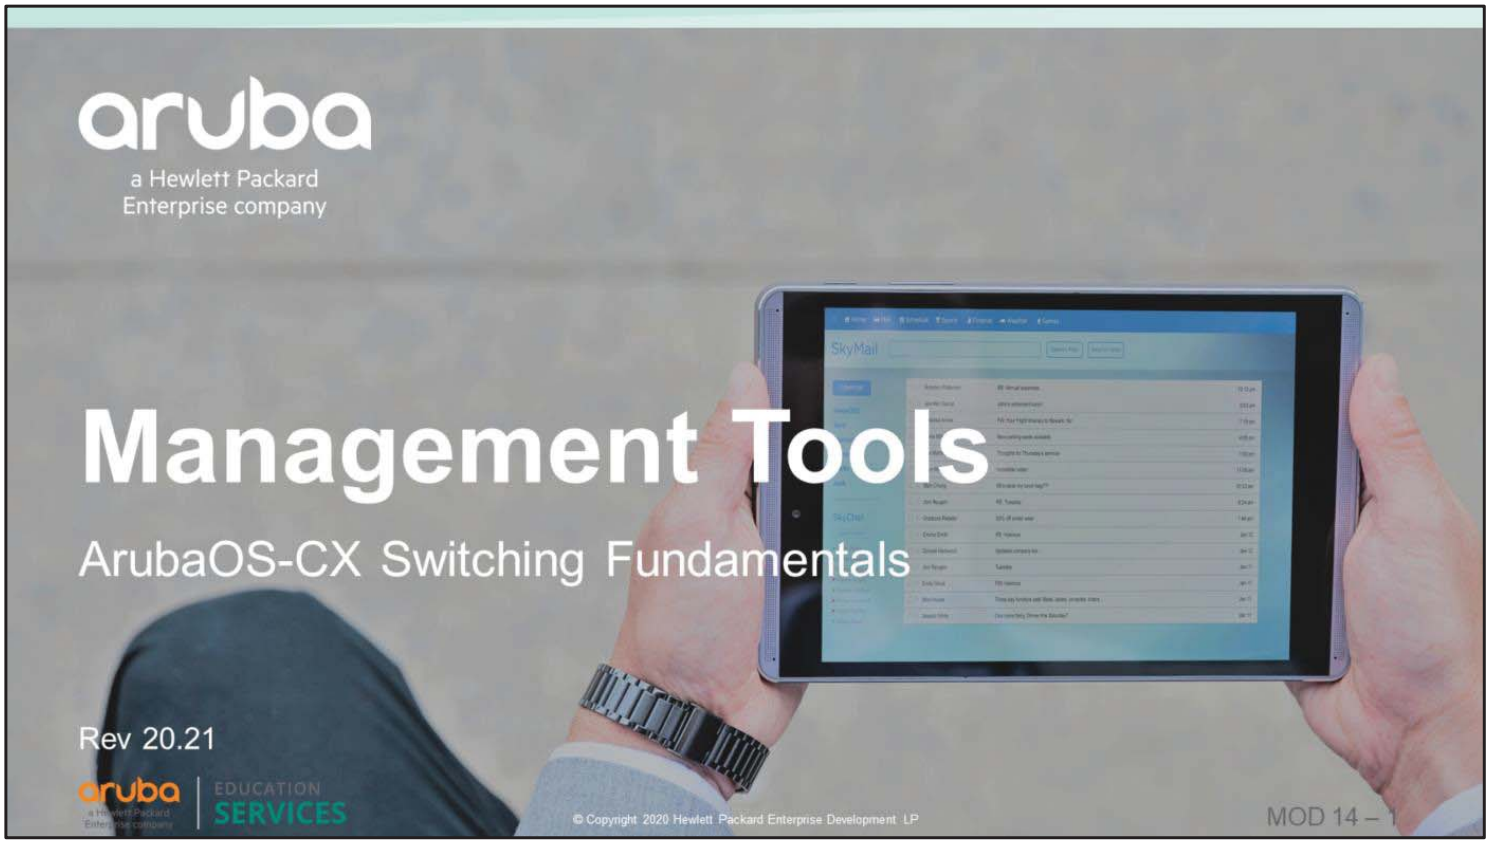
aruba
a Hewlett Packard Enterprise company
Management Tools
ArubaOS-CX Switching Fundamentals
Rev 20.21
aruba EDUCATION SERVICES
© Copyright 2020 Hewlett Packard Enterprise Development LP
MOD 14 – 1

Objectives

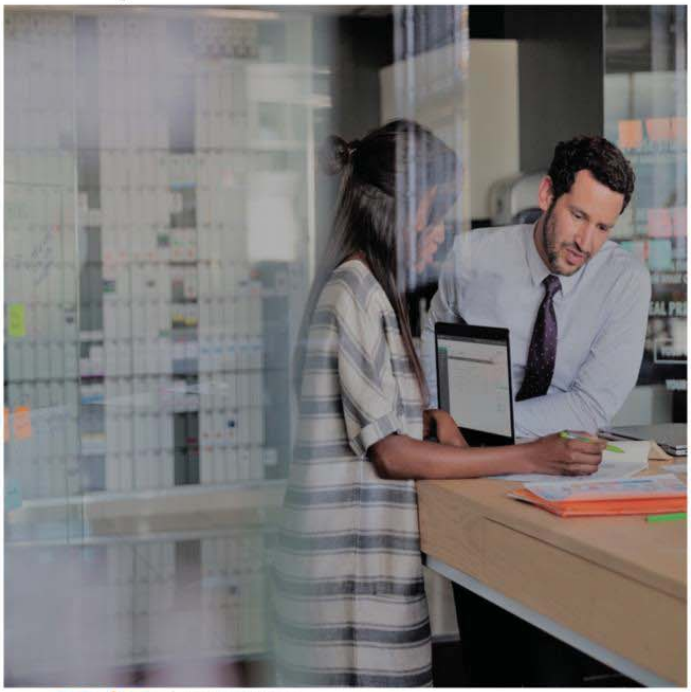

–Learning Objectives

- Describe NetEdit features, configuration, access
- Configure switches to support NetEdit
- Describe the Aruba CX Mobile app

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 2

After completing this module, you will be familiar with two key ArubaOS-CX management tools – NetEdit and the Aruba CX Mobile App.

|You will be able to describe NetEdit Features, configuration, and access.
|You will be able to configure ArubaOS-CX switches to support NetEdit.
|You will be able to describe the Aruba CX Mobile App.

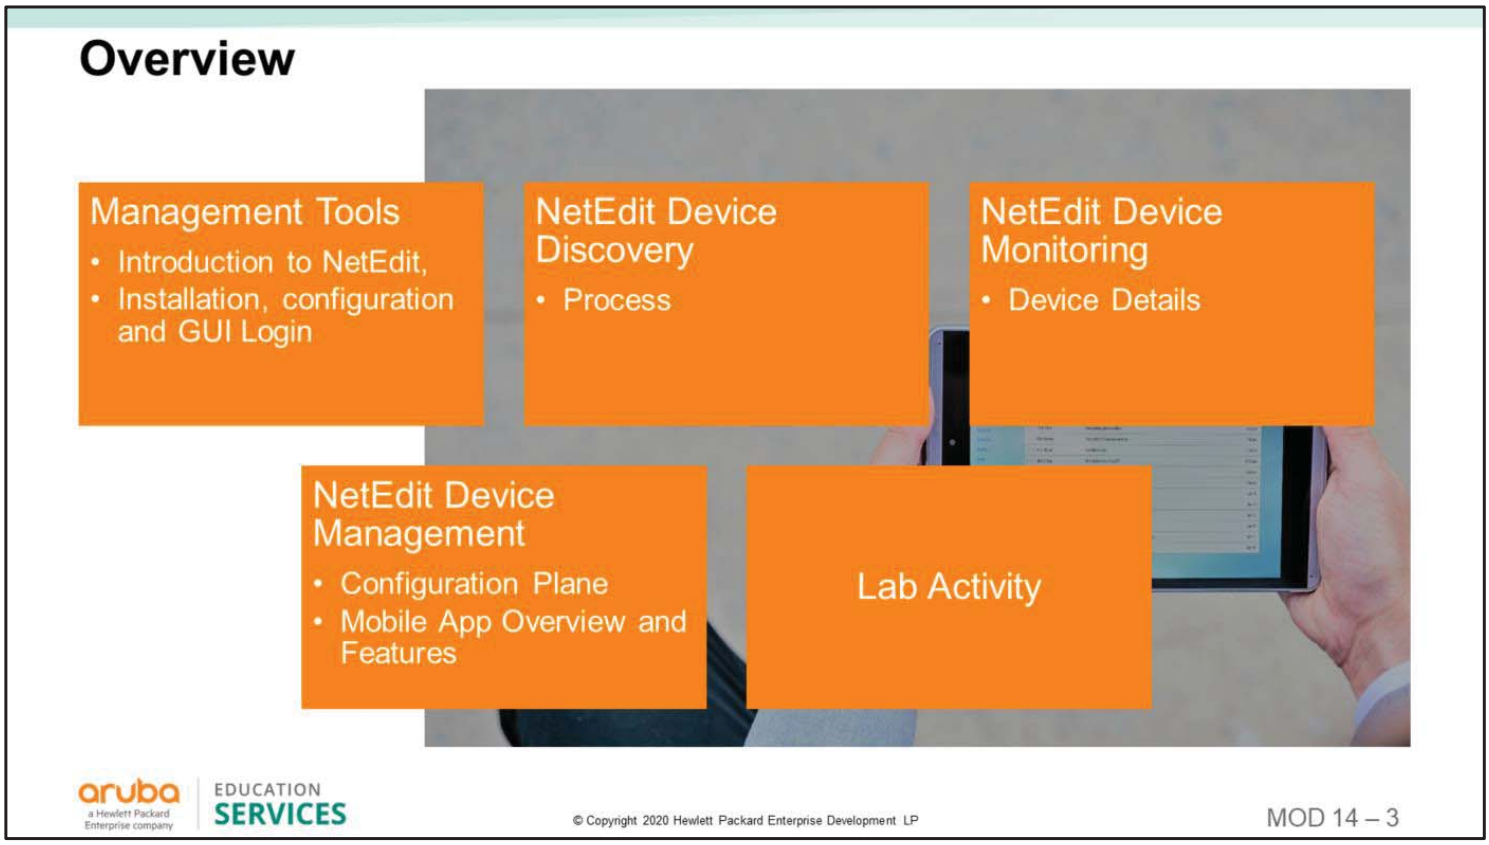

This module covers some very useful ArubaOS-CX management tools – to improve your network administration efficiency, and the ease with which you perform many configuration and monitoring tasks.

You begin by exploring NetEdit features, installation, configuration, and access, before learning how to configure ArubaOS-CX switches to support NetEdit.
Then you will explore the Aruba CX Mobile App.

Management Tools

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

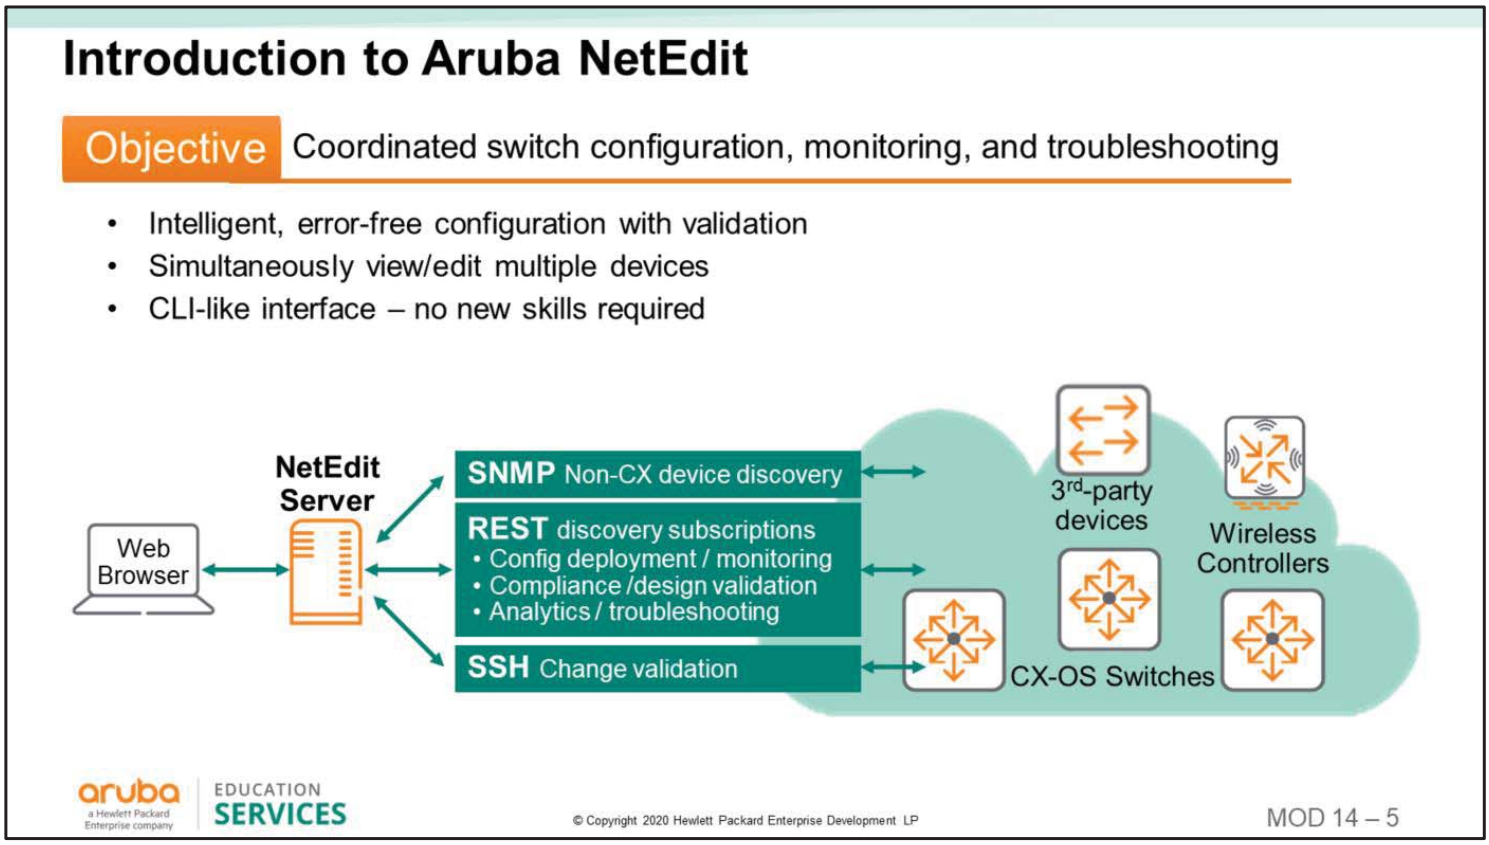

As networks grow, they become more challenging to maintain, especially when introducing a new protocol or feature. Changes to large networks must be prepared, designed, configured and validated on every single network device. CLI is a powerful configuration tool, but it is not scalable.

Aruba NetEdit provides for coordinated switch configuration, monitoring, and troubleshooting. It provides for intelligent, error-free configuration with validation for consistency and compliant, and the automation of multi-device change workflows without the overhead of programming. IT teams can coordinate end-to-end service roll outs, automate rapid network-wide changes, and ensure policy conformance after network updates.

Aruba NetEdit works closely with the embedded analytics in each ArubaOS-CX switch, provided by the Aruba Network Analytic Engine (NAE). You get intelligent assistance and continuous validation, with compliance and design validation, analytics, and troubleshooting.

As shown in the figure, NetEdit works with ArubaOS-CX switches via a Representational State Transfer Application Programming Interface (REST API). NetEdit uses SSH for change validation and supports SNMPv2c and SNMPv3 for non-CX device discovery.

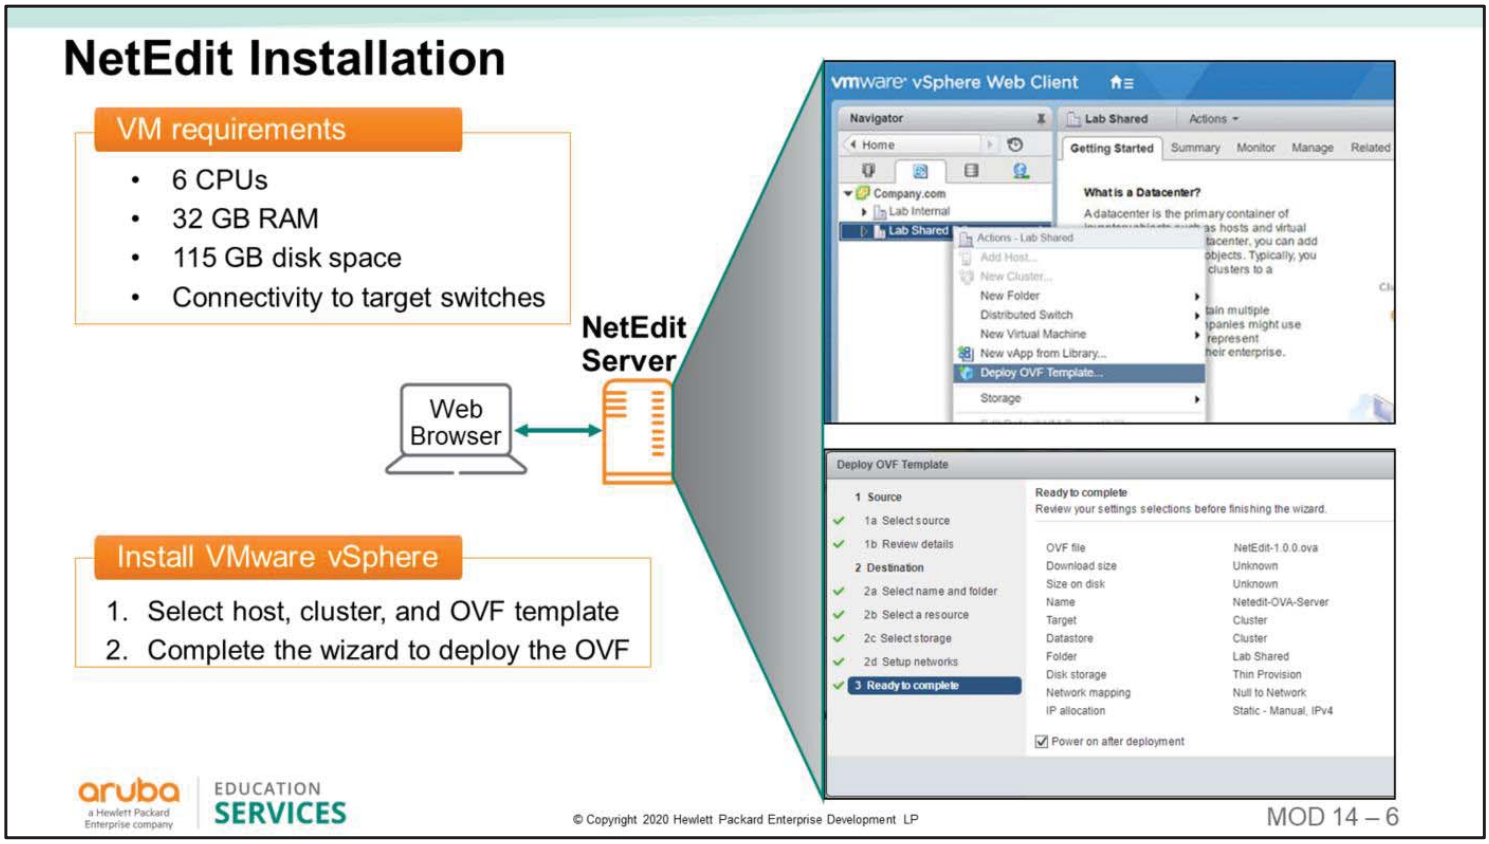

Aruba NetEdit is a web-based application and runs as an Open Virtualization Application (OVA) Virtual Machine (VM).

A NetEdit deployment in VMWare vSphere version 6.0 or higher has the following requirements:

- 6 CPUs
- 32 GB RAM
- 115 GB disk space
- Network connectivity to the target switches to be managed.

Install Procedure

1. Select the Host and Cluster and deploy the OVF Template.
2. Complete the wizard to deploy the OVF. After it the VM will be installed in the vSphere system.

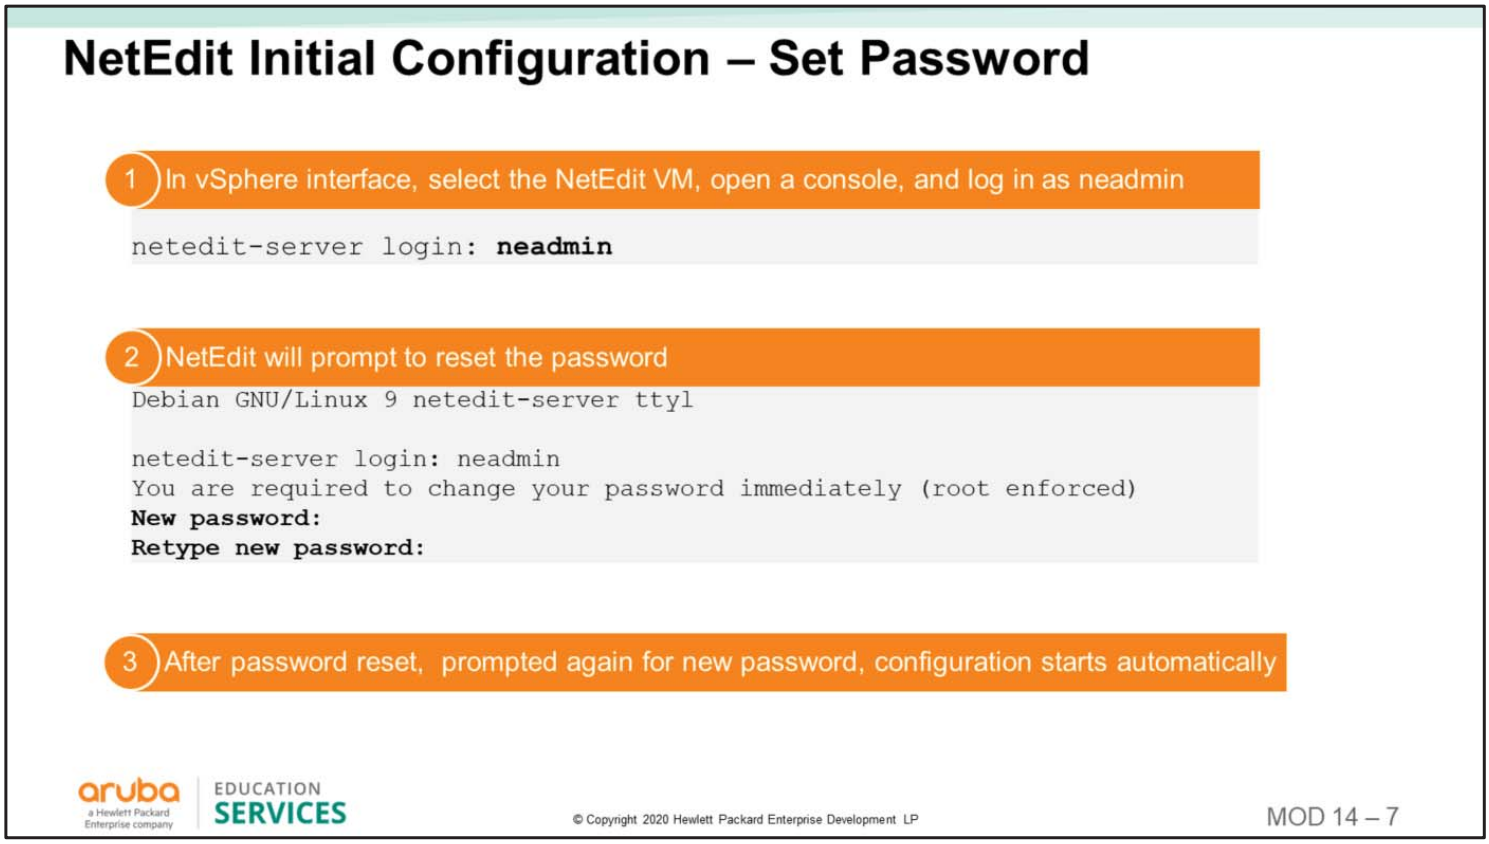

Initial Configuration

1. Use the vSphere interface to select the NetEdit VM, open a console and log in as neadmin.
2. NetEdit will prompt to reset the password.
3. After successfully resetting the password, the user will be prompted again for the new password, then the console-based configuration will automatically start.

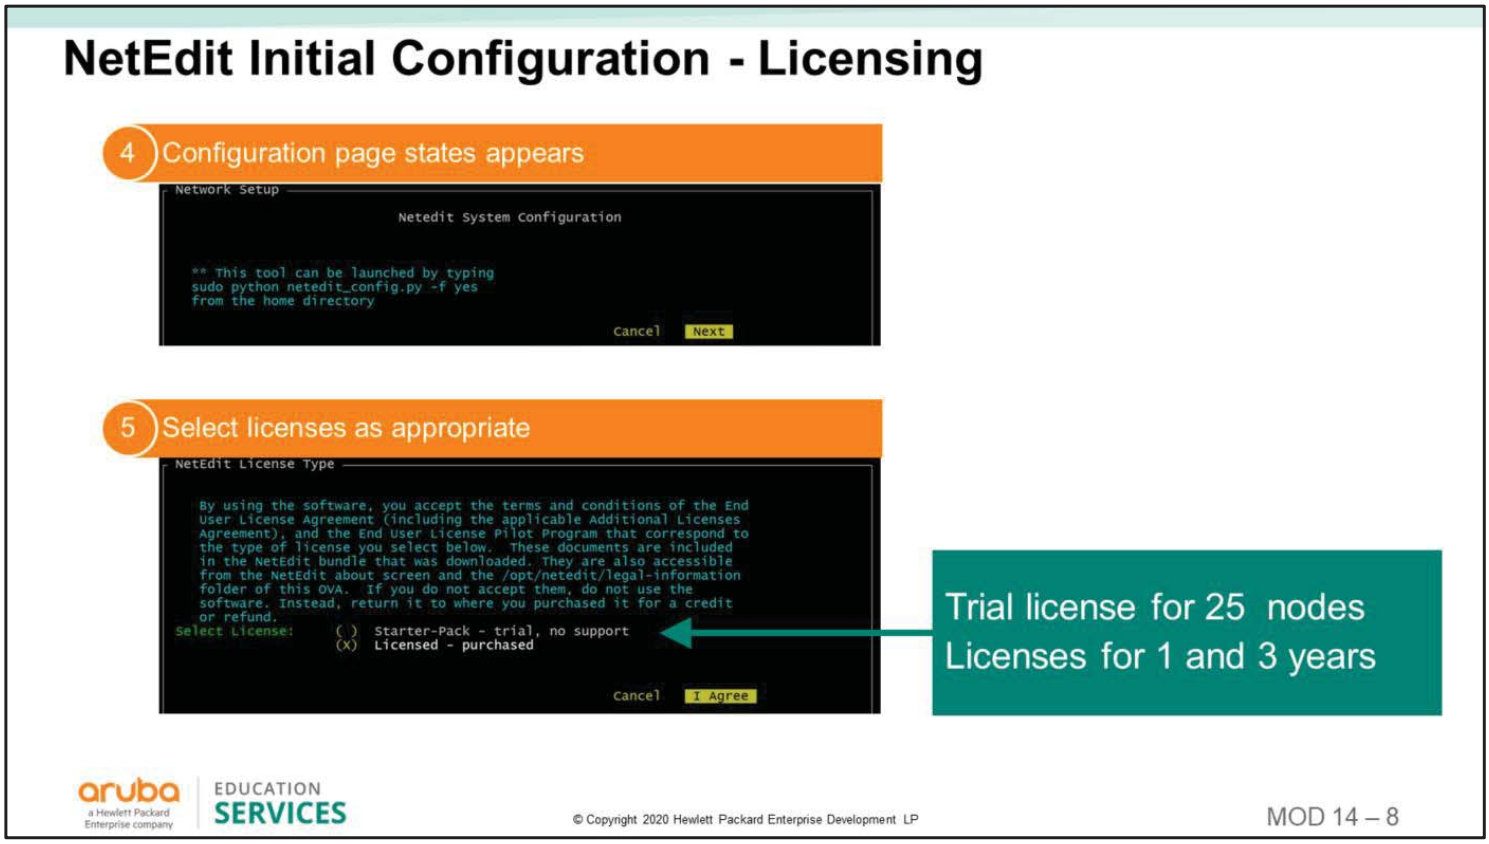

Initial Configuration

4. The first configuration page states that this tool can be rerun at any time.
5. Select the license that applies and then click I Agree.

NetEdit is currently available on a trial basis for up to 25 network switches. There are also licensing options for one- and three-year subscriptions for Aruba Support Services

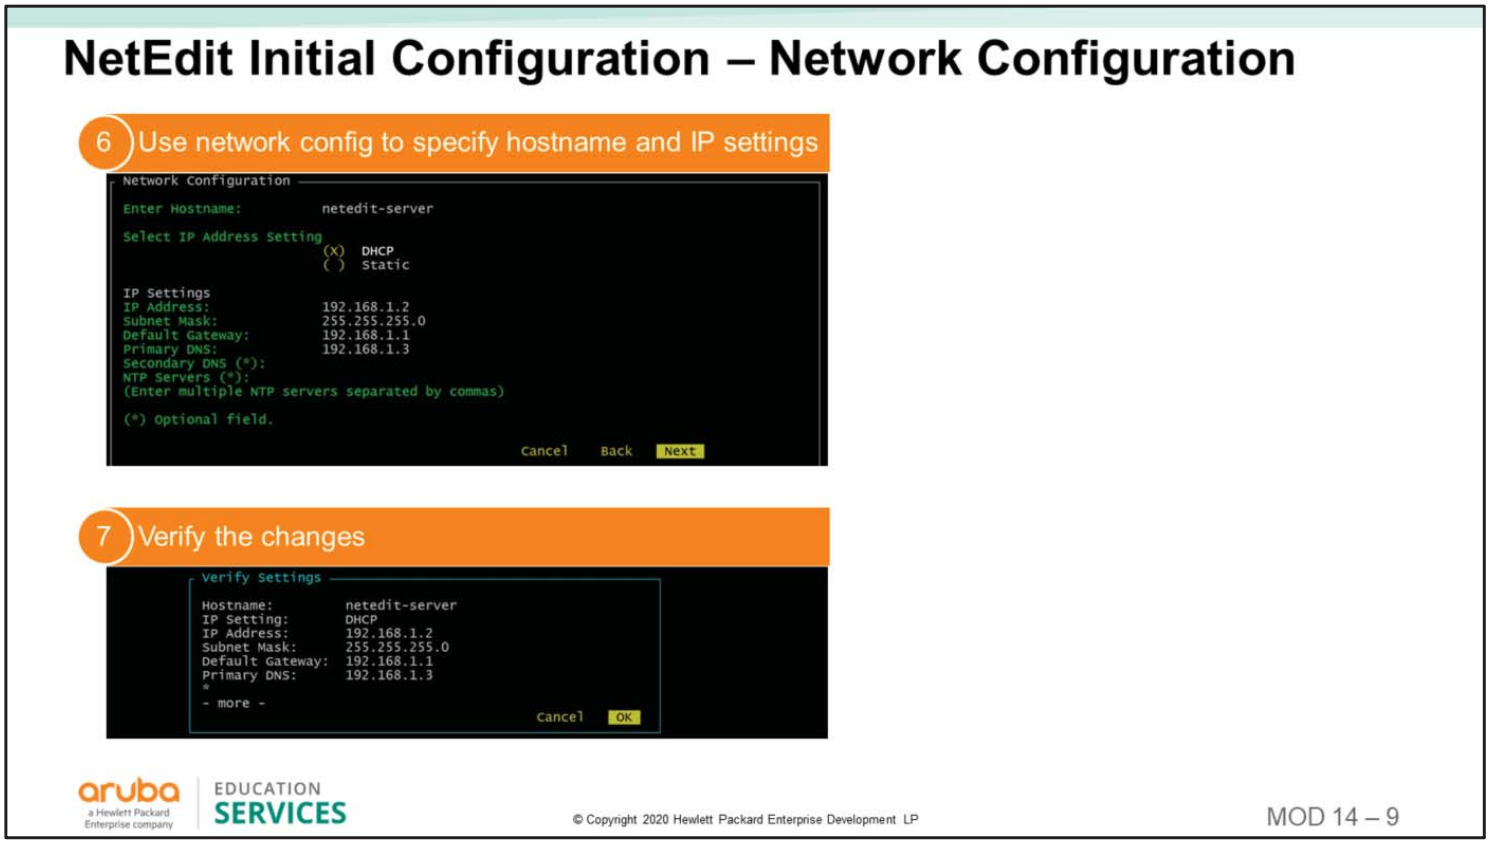

Initial Configuration

6. Use the network configuration screen to specify a hostname and IP settings and then click Next.
7. Verify the settings and click OK.

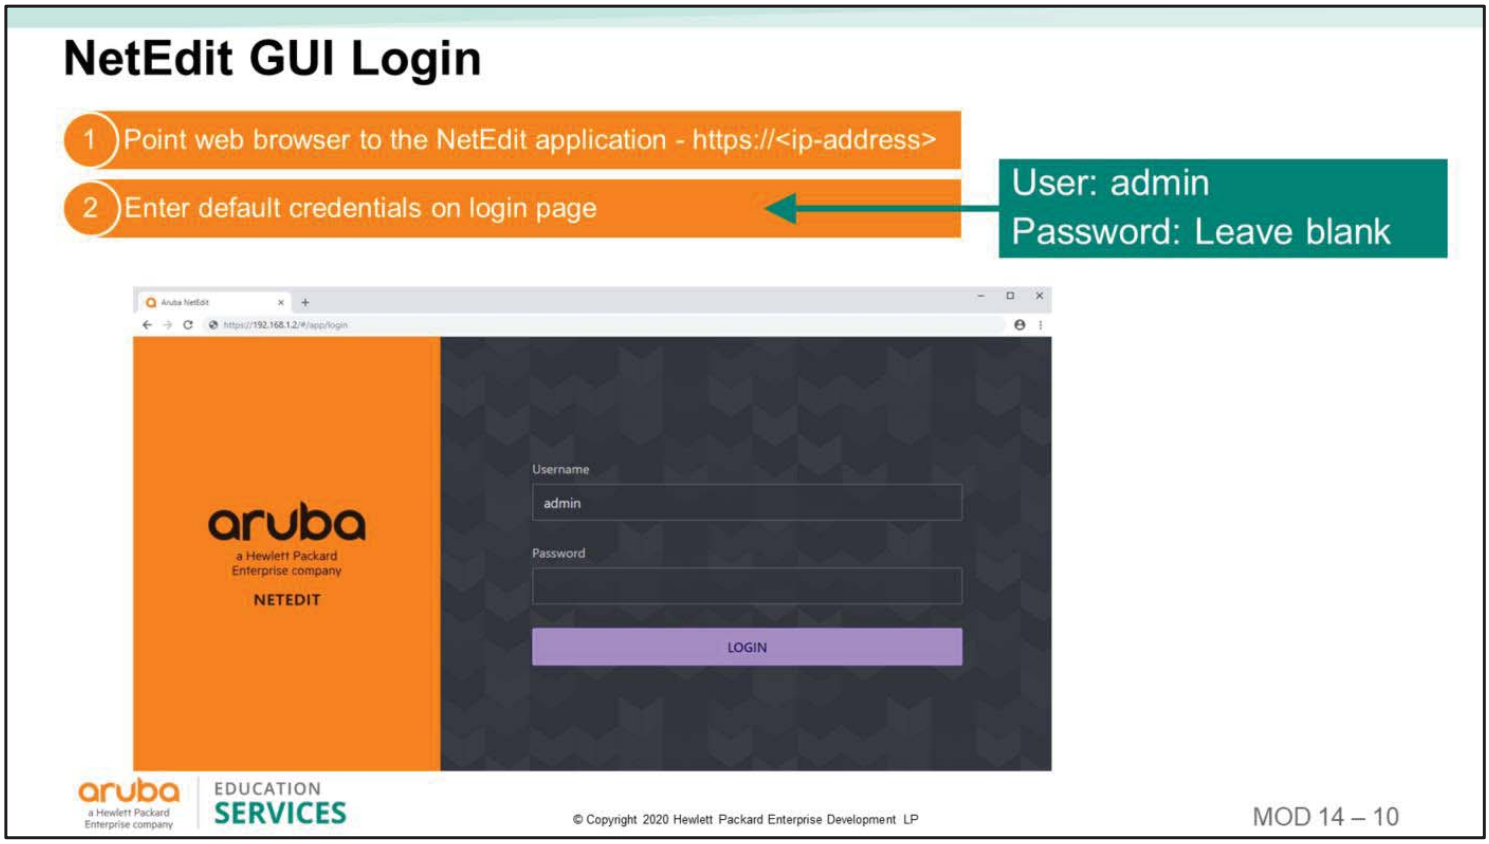

Login to NetEdit using the GUI

1. From a web browser, navigate to the NetEdit application using https://<ip-address>.
2. On the Login page enter the default credentials:
    - User: **admin**
    - Password: **Leave the field blank**

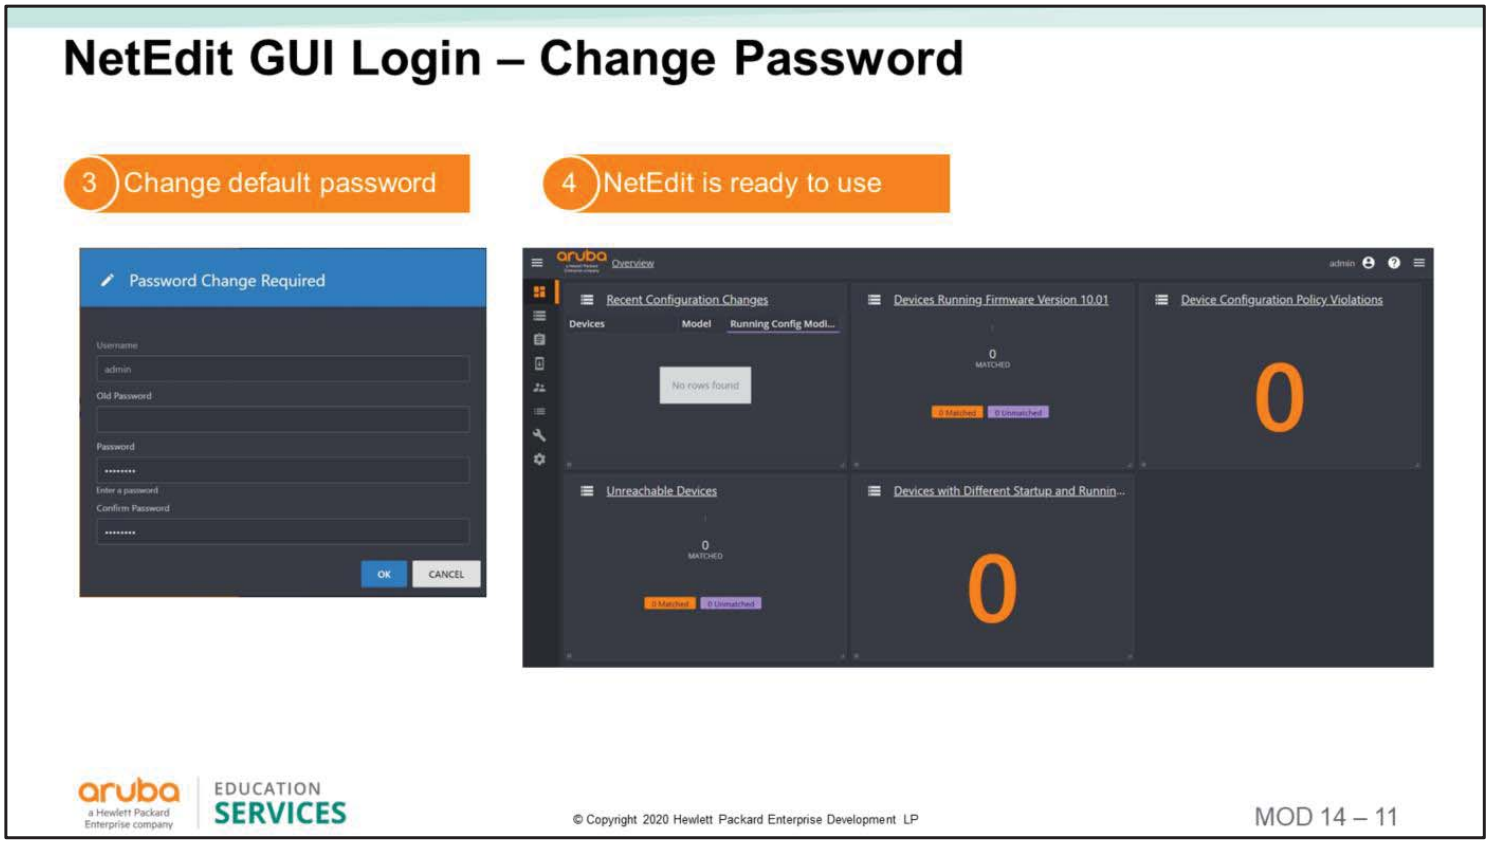

Login to NetEdit using the GUI

3. At the first login, a prompt to change the password displays.
    - Leave the Old Password field blank
    - Enter a new password
4. The overview page displays, this means that NetEdit is ready to use.

NetEdit Device Discovery

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

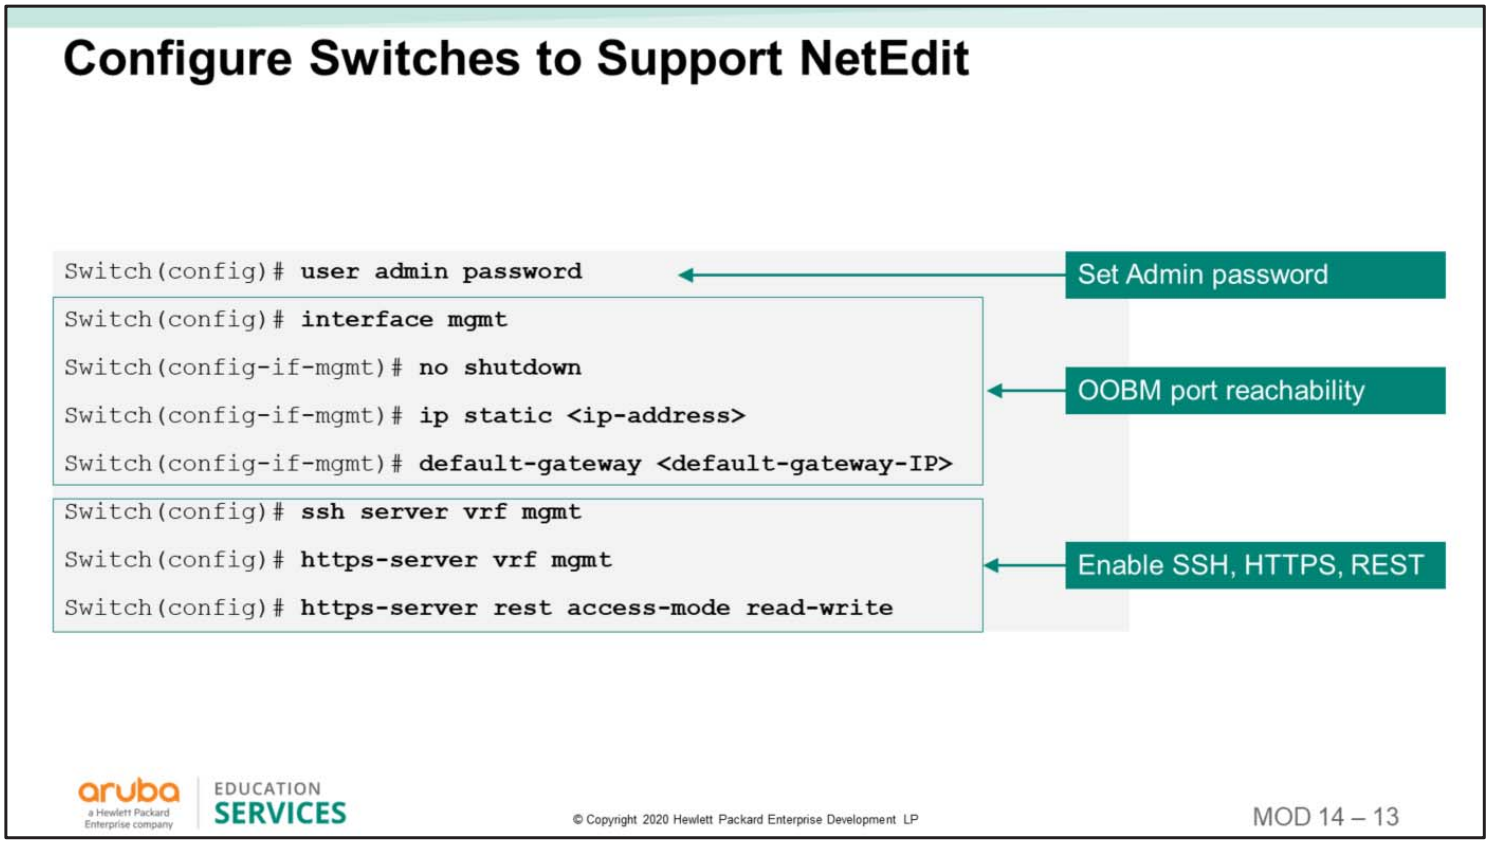

Aruba NetEdit require basic configuration on the switches.

Admin user must have a password.
Device must be reachable using the OOBM port.
SSH must be enabled using the **mgmt VRF.**
HTTPS must be enabled using the **mgmt VRF.**
REST interface must be enabled in read-write mode.

The following script can be used to meet these requirements.

Switch(config)# **user admin password**
Switch(config)# **interface mgmt**
Switch(config-if-mgmt)# **no shutdown**
Switch(config-if-mgmt)# **ip static <ip-address>**
Switch(config-if-mgmt)# **default-gateway <default-gateway-IP>**
Switch(config)# **ssh server vrf mgmt**
Switch(config)# **https-server vrf mgmt**
Switch(config)# **https-server rest access-mode read-write**

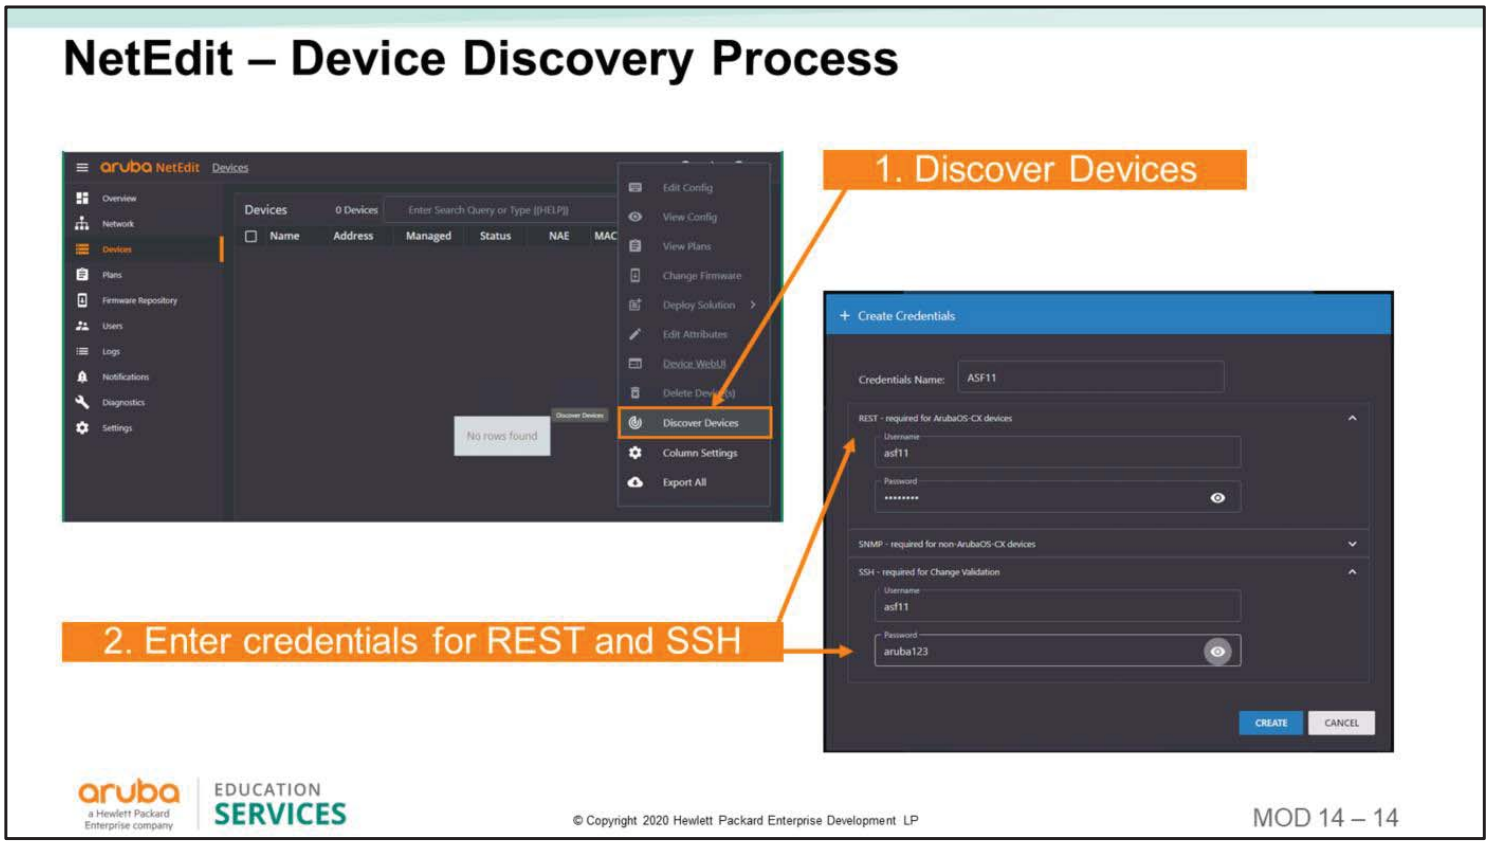

Aruba NetEdit supports an automated discovery process of ArubaOS-CX enabled switches. This process will permit NetEdit to monitor and manage the device.

Step1. Select **Discover Devices** under the **Actions** menu.
Step2. Enter the credentials for admin account in the ArubaOS-CX switch.

Note. NetEdit supports a manually adding devices.

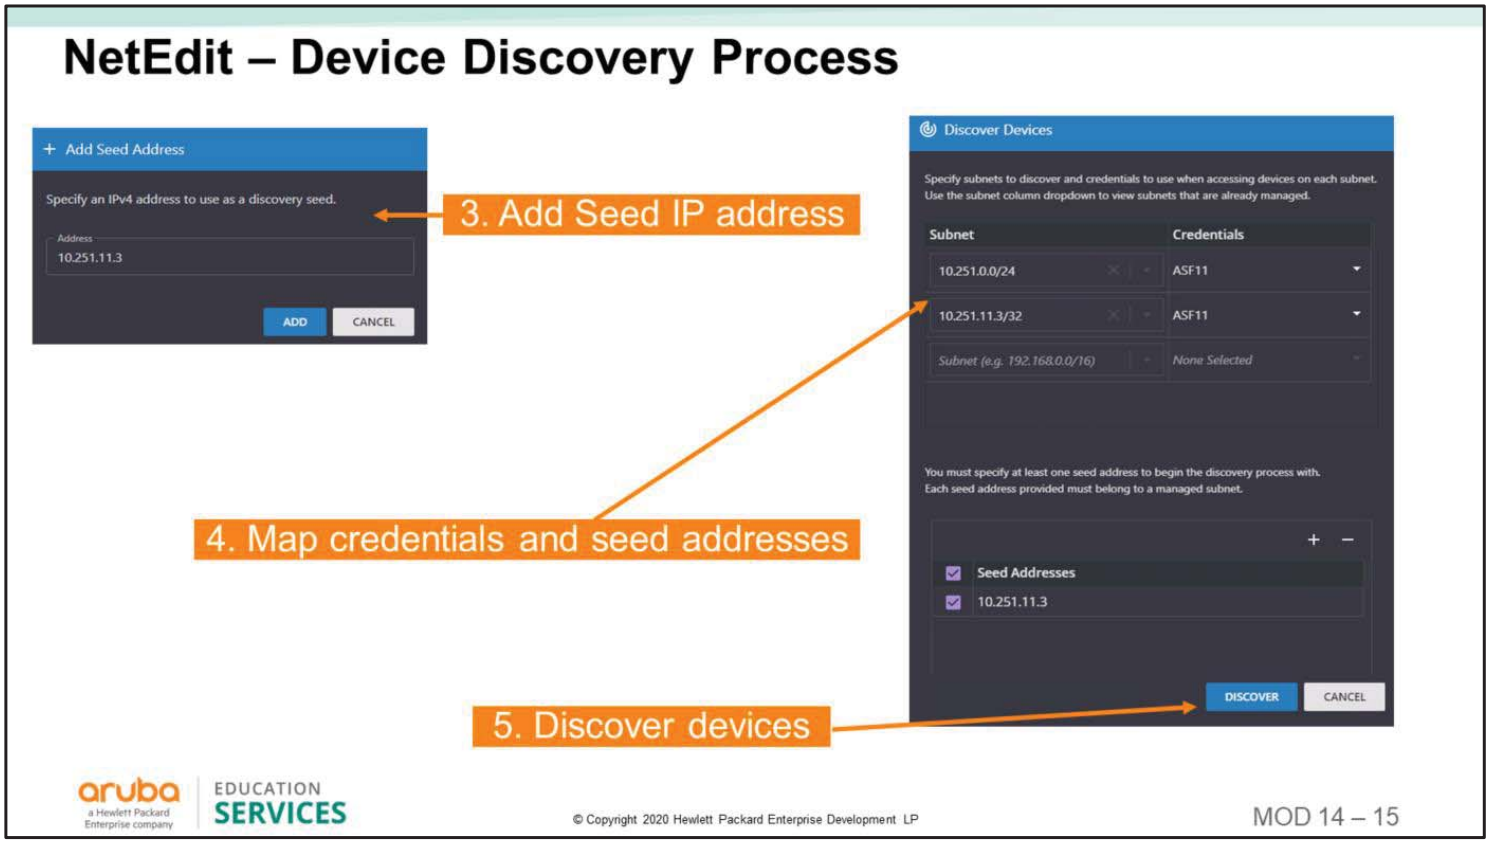

Step 3. Add seed IP address.
Step 4. Map the credentials and seed addresses.
Step 5. Discover the devices

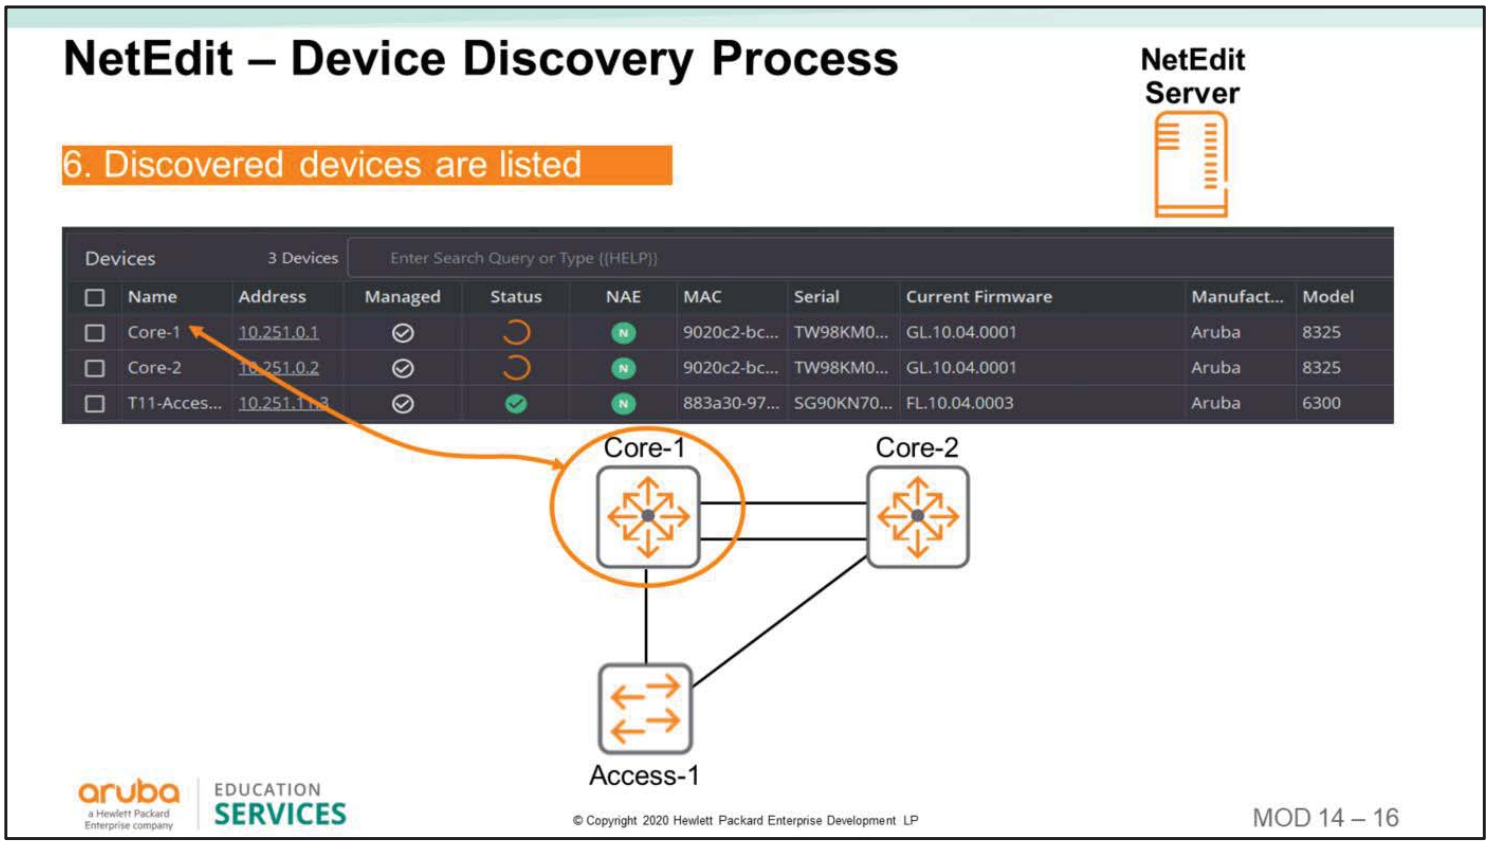

Step 6. Discover devices are listed.

NetEdit Device Monitoring

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

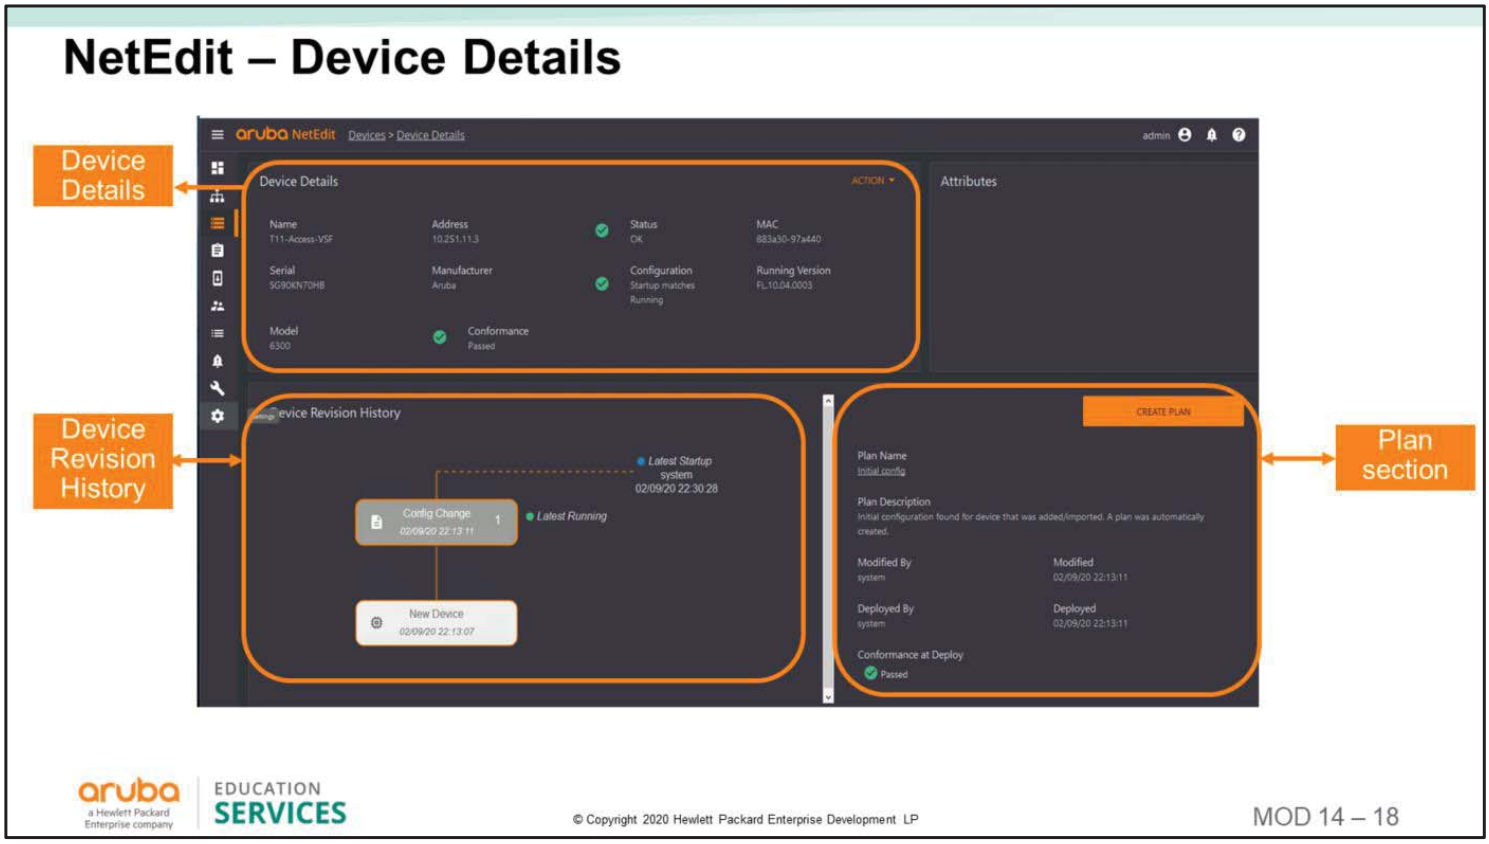

The Device Details section displays device general information including Hostname, Serial Number, Based MAC Address.
The Device Revision History section displays configuration changes that are applied to the device.
Create a plan section allows to configure a device.

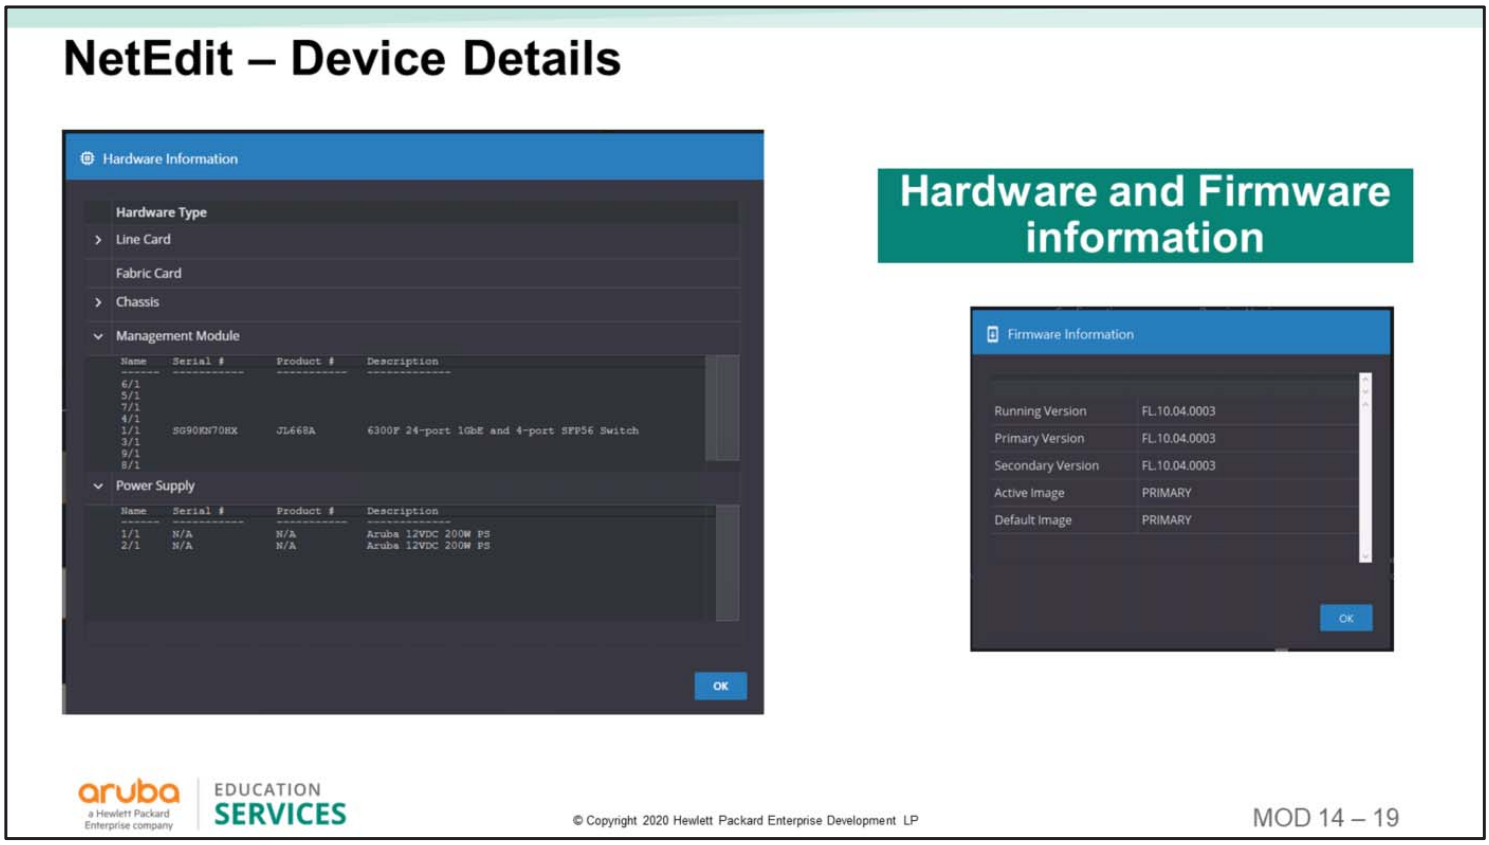

Monitoring using NetEdit includes verify the Hardware installed (such as transceivers, line cards and power supplies) on the ArubaOS-CX switch.
The firmware section will display the firmware information on both partitions (Primary and Secondary).

NetEdit – Device Details

Running configuration

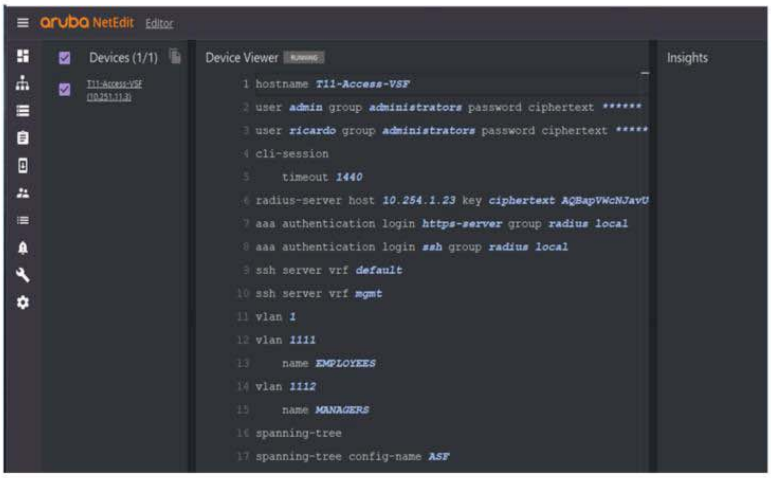

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 20

From NetEdit you can verify the running configuration of the selected switch.

NetEdit Device Management

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 21

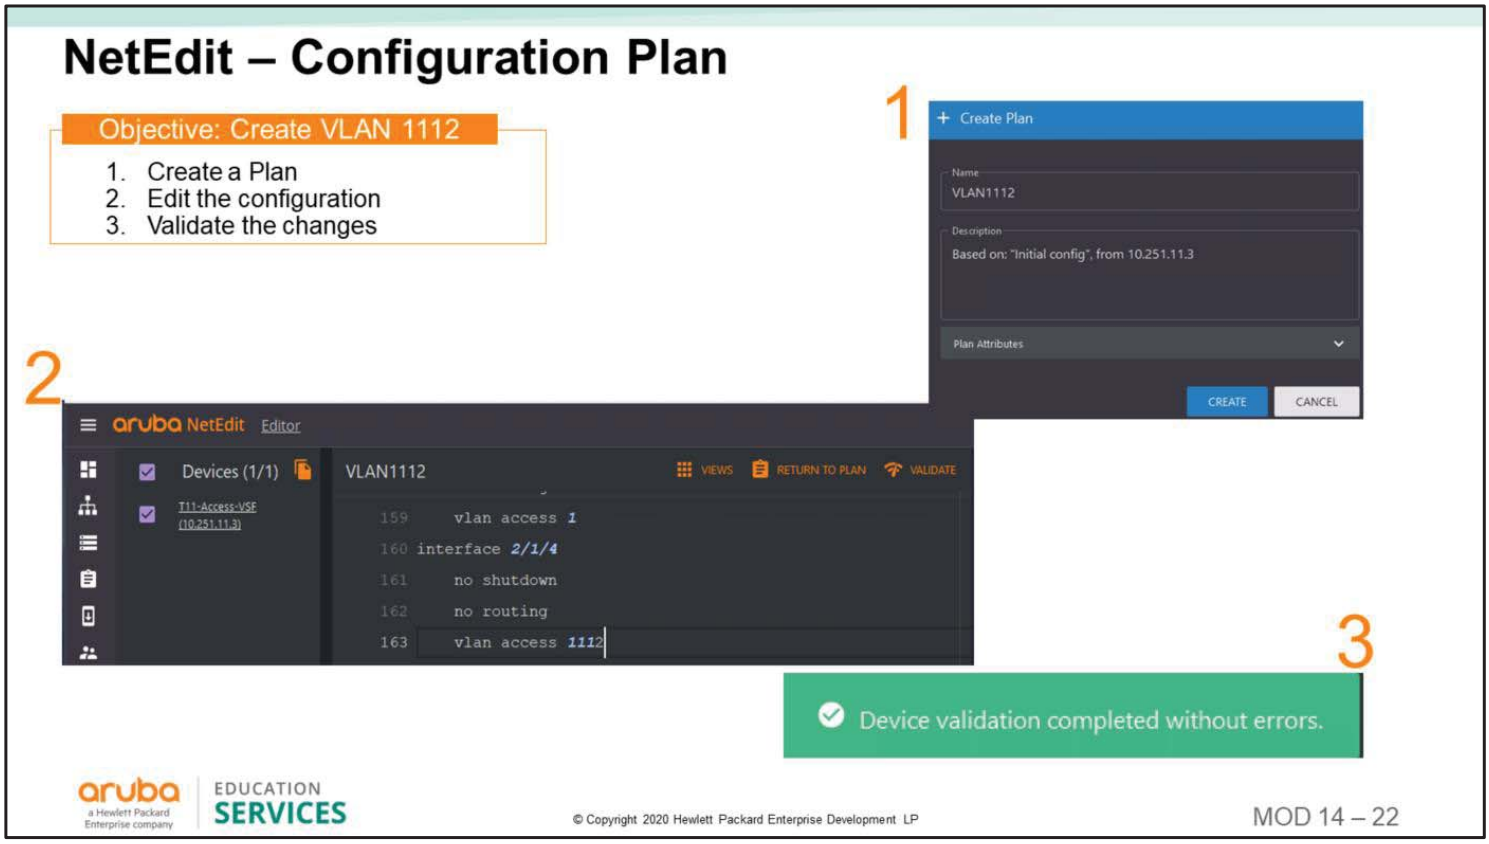

Every time you need to configure a device you will need to create a Configuration Plan.

As an example, consider that you will configure VLAN 1112 and assigned to the interface 2/1/4 using NetEdit.

Step 1. Create a Configuration Plan.
Step 2. Edit the configuration, this will feel and looks like configuring a port from the CLI.
Step 3. Validate the changes. No errors should appear in order to continue.

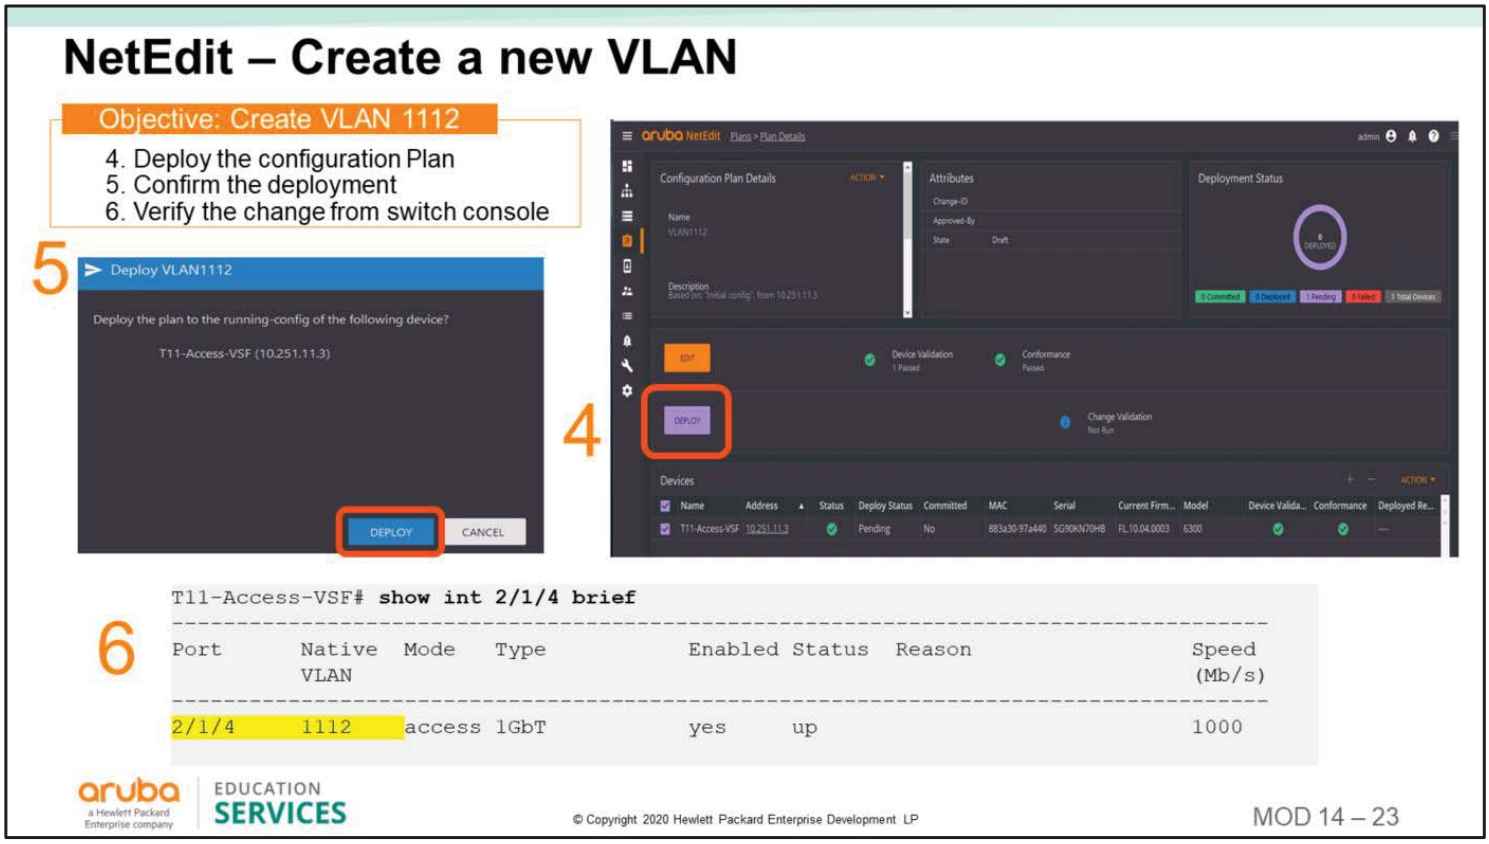

Step 4. Deploy the configuration Plan.
Step 5. Confirm the deployment of the Plan
Step 6. Verify the changes from the switch CLI.

Note. NetEdit supports configuring multiple devices using a single Configuration Plan.

Aruba CX Mobile App Overview

Objective Simplified switch install /configure

- Accelerates day zero configuration
- OS-CX connection via Bluetooth / Wi-Fi
- Monitor, configuration workflow
- Integrates with Aruba NetEdit

Requirements

- 8400, 8320 and 8325 AOS 10.02.0001
- 6300 AOS 10.4
- Apple iOS 12 or higher
- Android 5.0 or higher

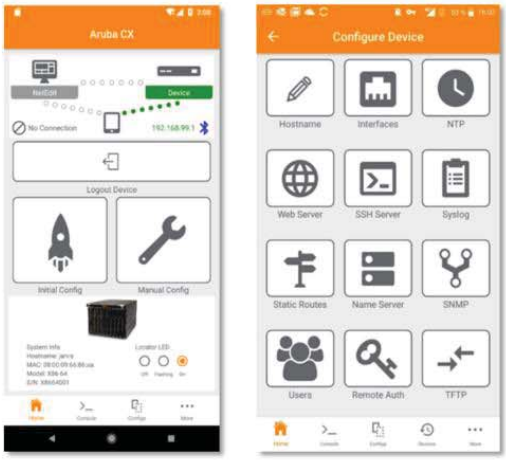

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 24

The **Aruba CX Mobile App** accelerates day zero configuration and deployment of ArubaOS-CX switches. No need for physical switch connections – connect your Apple or Android device via Bluetooth or Wi-Fi, for easy monitoring and configuration workflows and templates. Configure and monitor virtual switch stacks and PoE usage.

This application was created to address the complexity of switch configuration and installation for large deployments. You can validate that switches were configured and installed correctly, without relying on a network administration team for confirmation. The Aruba CX application integrates with Aruba NetEdit for intelligent configuration management, continuous validation, and overall network health.

ArubaOS-CX switches must run at least version 10.02.0001 for 8400, 8320 and 8325 models. The minimum code version for 6300 and 6400 switch series is 10.4.
Your mobile device must be an Apple device running iOS version 12 or higher, or an Android deice running version 5.0 or higher.

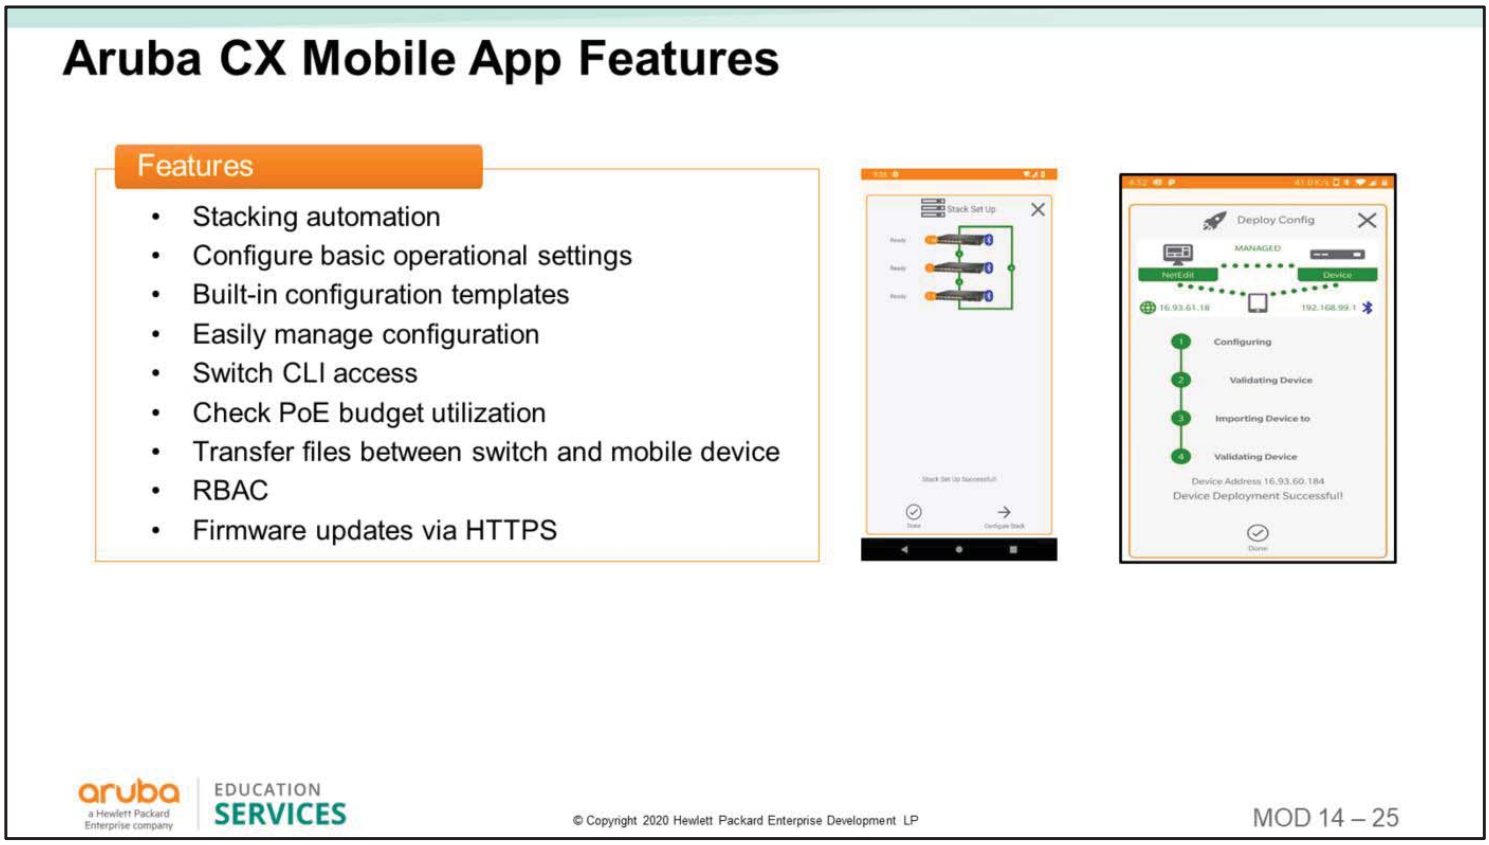

Key features

- The application uses Bluetooth and Wi-Fi to communicate with ArubaOS-CX switches.
- Is available for Apple iOS 12 or higher and Android 5.0 or higher platforms.
- Configuration of basic operational settings eliminates having to connect a console cable.
- Built-in configuration templates and ability to customize your own.
- Easily manage the running and startup configuration.
- Access the Switch CLI.
- Check PoE budget utilization.
- Transfer files between the switch and your mobile device.
- Role based Access Control (RBAC)
- Firmware updates via HTTPS
- Stacking automation
- Monitoring of PoE utilization

Question #1

Introduction to Aruba NetEdit

What is true about Aruba NetEdit connectivity options?

a. It uses SNMPv3 to discover and manage ArubaOS-CX switches.

b. It can use SNMPv2c to discover 3rd-party devices.

c. Many useful functions are available via a REST API.

d. You can simultaneously view and edit multiple devices.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 26

Correct answers:
It can use SNMPv2c to discover 3rd-party devices.
Many useful functions are available via a REST API.
You can simultaneously view and edit multiple devices.

Question #2

NetEdit Installation, NetEdit Initial Configuration – Licensing, NetEdit GUI Login, Configure Switches to Support NetEdit

Which of the options below describe valid NetEdit installation and configuration options?

a. NetEdit is a web-based application that runs on VMware.

b. The NetEdit trial license is limited to 25 weeks.

c. Initial NetEdit Web login requires username = admin, password = admin.

d. NetEdit connects to switches via any available switchport.

e. You can enable either SSH, HTTSP, or REST.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 27

Correct answers:
NetEdit is a web-based application that runs on VMware.

Question #3

Aruba CX Mobile App Overview, Aruba CX Mobile App Features

What is true about the Aruba CX Mobile App?

a. Android devices connect to OS-CX switches via Bluetooth or Wi-Fi.
b. Apple iOS devices connect to OS-CX switches via Bluetooth or Wi-Fi.
c. The app integrates with NetEdit.
d. The app can use SNMPv2c for integration with 3rd-part y devices.
e. You can transfer files between the switch and your mobile device.

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

Knowledge Check

© Copyright 2020 Hewlett Packard Enterprise Development LP

MOD 14 – 28

Correct answers:
Android devices connect to OS-CX switches via Bluetooth or Wi-Fi.
Apple iOS devices connect to OS-CX switches via Bluetooth or Wi-Fi.
The app integrates with NetEdit.
You can transfer files between the switch and your mobile device.

Lab Activity

aruba a Hewlett Packard Enterprise company | EDUCATION SERVICES

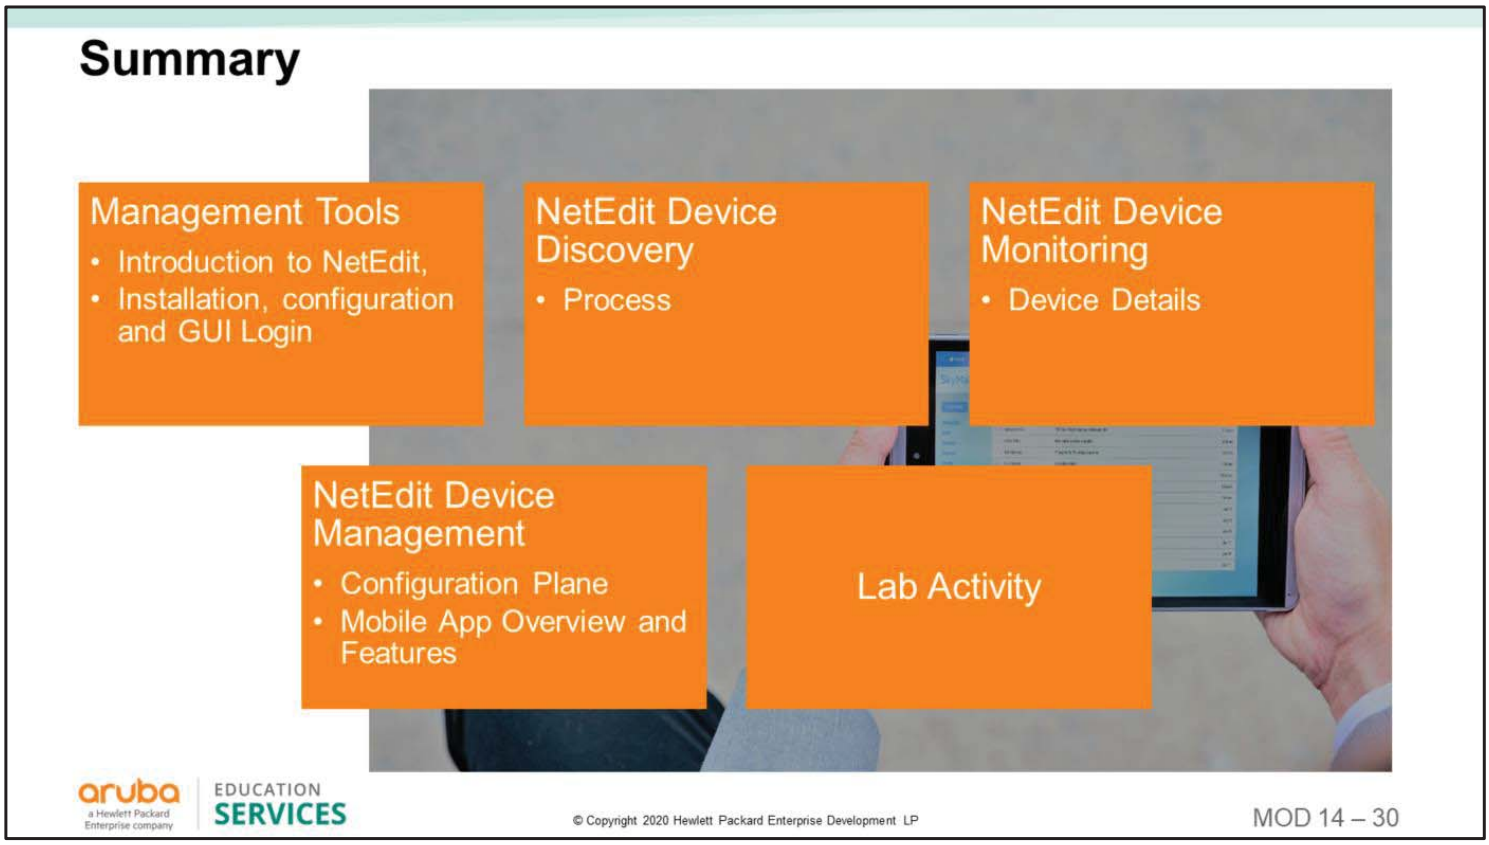

Congratulations! You have learned about some very useful ArubaOS-CX management tools – to improve your network administration efficiency, and the ease with which you perform many configuration and monitoring tasks.

You began by exploring NetEdit features, installation, configuration, and access, before learning how to configure ArubaOS-CX switches to support NetEdit.
Then you explored the Aruba CX Mobile App.

Glossary of Terms

Term	Definition
AOS-CX	Aruba's modern switch operating system built on a microservices architecture to deliver network-wide analytics and programmability.
VLAN	A virtual LAN (VLAN) standardized in IEEE 802.1Q, divides broadcast domains through the use of vlan tagging to provide segmentation, minimize broadcasts, and help to secure the layer two ethernet switching network.
Spanning-Tree	A suite of protocols that include standard IEEE 802.1d (STP), 802.1w RSTP, and 802.1s MSTP along with Cisco proprietary PVRSTP+ and PVSTP+. These were developed to enable a loop-free layer two topology while still allowing redundancy in the event of a network path failure.
VRRP	Virtual Route Redundancy Protocol is a standard (RFC 3768) that allows multiple layer three routers to provide a single virtual IP gateway address (VIP) within each VLAN for redundancy.
LAG	Link Aggregation or LAG is when multiple physical links are bundled into a single logical interface for improved throughput,resiliency, and loop-prevention. A LAG or multi-chassis LAG (MC-LAG) can be created dynamically using an open standard IEEE 802.3ad LACP or manually.
Subnetting	Taking a classful network address and dividing it into two or more smaller broadcast domains or subnets for the benefit of efficiently sizing network segments to the amount of host addresses needed.
OSPFv2 Single Area	The Open Shortest Path First <strong(OSPF) routing protocol version two allows for discovery and exchange of IPv4 routes between neighboring layer three devices. It is a standardized link-state protocol using link speed as a cost metric and Dijkstra's shortest path first algorithm to determine the best path when forwarding traffic across a single area within a single autonomous system (AS).
VSF	Virtual Switching Framework (VSF) is an Aruba stacking technology that virtualizes multiple physical devices into one virtual fabric to provide high availability, network simplification, and a single point of management.
NetEdit	An Aruba management appliance used to orchestrate the configuration of multiple switches in order to provide automation, client-side validation, compliance and analytics.
Computing Network	A group of computing resources that permit data exchange between devices.
LAN	A group of computer devices in the same geographic location.
WAN	A group of computer resources that communicate over long distances.
Internet	The global collection of large company networks (autonomous systems) and smaller compute devices connected over a public address space.
MPLS	Multi-Protocol Label Switching (MPLS) is a labeling technology often used by service providers (ISPs) to create scalable private pseudo-wire VPNs for customers across the public internet.

ATM	Asynchronous Transfer Mode (ATM) networks are provider based layer two connections that better enabled integrated services (voice, video, and data) by breaking customer traffic into smaller 53 byte packets to improve performance on legacy networks.
Frame Relay	A legacy way for provider networks (ISPs) to emulate layer two overlay connections or circuits between customer sites. Generally considered legacy today, Frame Relay was a replacement in the 1990s for X.25 analogue networks until ATM displaced Frame Relay shortly thereafter.
Dark Fiber	Describing Fiber Connections between sites that are unused or available for growth. Dark Fiber may also refer to fibre being leased by a privately-operated network rather than publicly available.
Protocol	A set of rules that computer devices follow to establish and maintain communication.
OSI	Open System Interconnect (OSI) is the standard communication model for computing devices created by the ISO to describe network communication into seven separate layers. This separation of network operating functions allow for industry specialization, standardization, and better understanding of networking processes.
Physical Layer	The first layer of the OSI model applying to the media that carries data broken down into bits of ones and zeros.
Data Link Layer	The second layer of the OSI model referred to as frames. Frames allow for forwarding encapsulated packets of data in a switched network. Switches typically are found at this layer using protocols such as Ethernet and Frame Relay.
Network Layer	The third layer of the OSI model referred to as packets. This layer is used to establish end-to-end communication across multiple LANs or WANs. Routers and PCs are found at this layer using protocols such as IPv4 and IPv6.
Transport Layer	The fourth layer of the OSI model referred to as segments. This layer controls link reliability using sequencing to segment packets and allowing error control. Devices such as firewalls are typically found at this layer leveraging protocols such as UDP and TCP.
Session Layer	The fifth layer of the OSI model enabling the setup, maintenance, and teardown of sessions between computing devices.
Presentation Layer	The sixth layer of the OSI model used to transform data into the format that applications can accept over a network connection. Often used when applications require compression, encryption, or code translation.
Application Layer	The seventh layer of the OSI model allowing the end user a front end way to leverage network traffic. For example a web browser would utilize HTTP as the application at this layer to present traffic to the end user.
PDU	Protocol Data Unit (PDU) is the header and payload of data at each encapsulation or de-encapsulation layer in the OSI model.
Media	The physical carrier of ones or zeros found at layer one of the of the OSI model. Commonly network ports will serialize the data as bits over the media of copper, fibre, or radio.
Optical Fiber	Media using light pulses to send bits along either reflective plastic coating (MMF) or glass silica (SMF) to obtain the higher performance and avoid EMI issues that may affect copper cabling.

MMF	Multi Mode Fiber (MMF) uses reflective plastic coated fiber to transport data causing multi wavelengths or modes of light. MMF often is used for high speed uplink between floors of buildings and may be referred to a "vertical fiber" due to having a lower cost and shorter distance than single mode fiber (SMF).
SMF	Single Mode Fiber (SMF) is glass silica with a much smaller core carrying only a single wavelength of light. SMF can achieve greater distances than MMF and therefore is often used between buildings or site to site.
Radio Frequency	The electromagnetic signals radios use as the media to transmit data modulating frequency phase and changes in amplitude.
Duplex	The communication system composed by two or more parties. Half-duplex describes sending or receiving but not both. Full-duplex is found more commonly in switched networks and allows sending and receiving communication without collisions.
Multicast	A one-to-many approach to sending traffic where a receiver registers a request to listen to traffic being streamed from a sender. Multicast is a more efficient than broadcast traffic way to has a large number of devices receive traffic from a small number of senders.
Unicast	A single sender forwarding traffic to a single recipient. This type of traffic is referred to as "one-to-one".
Broadcast	Traffic that is flooded within a subnet to reach all receivers. This type of traffic is referred to as "one-to-everyone".
Binary	A base-2 numbering system utilized in dotted decimal addressing such as IPv4 addresses.
Hexadecimal	A base-16 numbering system with numbers 0-9 and letters "A-F". This system is used in protocols such as Ethernet and IPv6 addresses.
TCP/IP Stack	Foundational protocol suite that describes the end-to-end data communication with a four layer model compared to the OSI seven layer model. The TCP/IP refer to the two protocols commonly used at the transport and internet layers of the stack.
TCP	Transmission Control Protocol is a Transport layer protocol responsible for flow and error control with sequencing and a 3-way handshake to reliably forward packets.
UDP	User Datagram Protocol is a Transport layer protocol for connectionless forwarding of traffic that does not require the overhead and reliability of TCP. A common use case is real time voice and video traffic.
Ethernet	An IEEE 802.3 standard of framing switching data at layer two of the OSI model.
MAC	Ethernet addresses that are 48-bits represented as 12 hexadecimal characters. The first 22 bits are registered through IANA as the OUI or Organizational Unique Indicator from the NIC manufacture. The last 24 bits are serialized to avoid duplicate MAC or Media Access Control addresses from occurring on a LAN.
NIC	Network Interface Card found in compute devices to allow network access. A NIC can be wired or wireless. NICs are assigned a unique Burned In Address (BIA) for layer two access to the LAN.
Preamble	The initial bits sent between two devices to synchronize timing on the link before sending a SFD indicator after which the bits that follow are the Frame itself.
SFD	The final eight bits in the Preamble after the link is synchronized indicating that the sender is ready to forward the fame itself.

EtherType	The portion of the layer two Ethernet frame that maps to which layer three protocol should process the payload.
Payload	The portion of a Frame or Packet that carries the data a compute device is sending or receiving.
FCS	Frame Check Sequence is a checksum performed at layer two of the frame to detect data that may become corrupted in while in transmit between switched devices.
TTL	The Time To Live portion of an IPv4 or IPv6 header is a routing loop prevention mechanism that counts down from an 8-bit number each time the packet is routed through a layer 3 device. If the router see the next hop as as zero then the packet is dropped an ICMP message is forwarded to the original sender.
Protocol	In the IPv4 header, the Protocol field maps which Layer four protocol the packet should processed by.
IPv4 Address	A 32-bit address used to indicate original sender or destination for a packet. IPv4 addressed are shown as four octets in dotted decimal format from 0 - 255 to represent 8 bits each.
TCP Three-Way Handshake	A layer four mechanism used by TCP to establish communication between a source and destination compute device through the use of SYN and ACK packets that are sequenced. A successful outcome of the three-way handshake is that a session is establish ahead of forwarding actual TCP data payloads.
SYN flag	The Synchronization flag of a tcp three-way handshake starts the sequenced communication from a sender to a receiver with the expectation being the receiver responds with an acknowledgement using the next sequence.
ACK flag	The Acknowledgement flag of a tcp three-way handshake validates a sequenced SYN packet was received successfully and therefore acknowledged using the expected next sequence number.
SYN/ACK flag(s)	A TCP three-way handshake segment with both the SYN and ACK flagged as a response to the initial SYN received. The SYN in this response starts a separate sequence from the receiver back to the sender.
TCP Sequence	A numbered sequence between a TCP sender and receiver used as tool for resending lost packets and windowing to efficiently take advantage of available bandwidth between two compute devices.
Port Numbers	Referring to the 16 bit sender and 16 bit receiver addresses used by TCP and UDP to packet a segment PDU to an application.
Destination Port	Maps the Layer four TCP or UDP port number to the receivers application layer protocol.
Source Port	Indicates a random well-known port (above 1024) used identify which session the traffic is sourced from. The source port also assists in firewall and NAT functions.
Window Size	The field in a TCP header that is used by a receiver indicates back to the sender the amount of data that can be sent without being lost. This establishes flow control to more efficiently utilize available bandwidth.
Switch	A layer two device used to connect computing devices into one or more Local Area Networks. Switches forward traffic by learning the mac address from the senders and mapping that MAC to the source port that heard the traffic. Switches own switchport mac address are transparent and do not affect the forwarding of frames.

Router	A layer three network device that forwards packets based on the destination IP address. Dedicated routers are often used as gateway devices connecting a LAN to a WAN or ISP with few port connections.
Multi-Layer Switch	A switch with high density of ports that also can router traffic at layer 3. Most commonly these layer 3 switches are utilized in the Core and Aggregation layers of campus networks.
Wireless Access Points	APs enable wireless users using 802.11 radios to bridge traffic to ethernet 802.3 cables. APs may be standalone, controller managed, or a hybrid of the two.
Firewall	A layer four function that applies a permit or deny policy on how matched traffic passing through the device based on security rules. Firewalls can be implemented as dedicated appliances or as function in layer three devices.
Hexadecimal	Hexadecimal addresses represented using numbers from 0-9 and letters A-F to represent the sixteen possible 4-bit decimal values from 0000 to 1111. Hexadecimal is the format used in IPv6 addresses.
Dotted Decimal	A method of expressing IPv4 addresses as octets of 8 bits represented as a value from 0 to 255. Each octet is separated by a dot.
Binary	A bit value representing either a 0 or a 1 to express numerical notation as a base of 2 rather than our typical 10 base..
Two-Tier	A network design noted as having a Core and Access layer without a Distribution layer between the two.
Three-Tier	A network design utilizing a Distribution layer typically to provide layer three feature for traffic from the access layer and to add scalability over a two-tier design.
SFP	Small Form-factor Pluggable single port slots that are compact, hot-pluggable network interfaces to be equipped with any suitable type of transceiver needed.
ZTP	Zero-Touch-Provisioning simplifies installation of the switch infrastructure by allowing for adding new devices to the network with little or no configuration by the person deploying the device.
PoE	Switchport Power Over Ethernet standards that include PoE 802.3af, PoE+ 802.3at and PoE++ (or uPoE 802.3bt)standards to meet increased power needs of attached ethernet devices.
Smart Rate	802.3bz standard of switchports where a single port can provide 1Gb/s, 2.5Gb/s, or 5Gb/s throughput often required for APs
Data Center	A facility that centralizes an organization's IT operation and equipment, as well as where it stores, manages and disseminates its data.
REST	REpresentational State Transfer protocol designed to allow access to web servers typically providing access to operations such as Create, Read, Update, or Deleting information on a device.
REST API	A programmable interface devices will expose to allow for REST command access.
Aruba NAE	Network Analytics Engine running on Aruba AOS-CX switches to provide monitoring analytics, assurance, automation, and remediation.
MM	Management Modules found on Aruba chassis based switches such as the 5400R, 6400, and 8400. The modules house control plane function in addition to providing management of the switch itself.
VSX	An Aruba redundancy stacking technology to allow switches or chasses with additional resiliency by keeping the control plane separated per switch while still providing MC-LAGs.

MC-LAG	Multi-Chassis Lag creates one virtualized switch composed by two individual physical switches. (see VSX)
AOS-CX	Aruba Operating System that is a modern database driven switch OS providing root analytics with NAE, REST APIs and Python scripting support, microservices architecture and state synchronization.
Jumbo Frames (Packets)	Frames where the maximum transmission unit (MTU) size exceeds the standard 1500 bytes.
VXLAN	Virtual Extensible LAN tunneling standard (RFC 7348) that forwards ethernet frames over layer 4 using UDP port 4789 to allow for scalable layer two access over a routed network core.
RSTP	Rapid Spanning Tree Protocol 802.1w standard that ensures a loop-free switching topology by blocking all redundant links without consideration for load-sharing.
MSTP	Multiple Spanning Tree Protocol 802.1s standard to ensure a loop-free topology switching topology. MSTP blocks redundant links within each instance created (typically one per gateway). Groups of vlans are mapped to an instance which allows for load-sharing traffic more efficiently than traditional PVSTP+.
IGMP	Internet Group Management Protocol is used to build a multicast SPT forwarding table for layer three devices.
Port Mirroring	Used to copy Tx or Rx (or both) frames that pass through a port and mirror them out a monitoring port for analysis.
BFD	Bidirectional Forwarding Detection is a protocol for determining is two devices can reach each other. Other slower protocols can utilize BFD to leverage the sub-second failure detection.
OSPF	Open Shortest Path First is a standardized routing protocol for providing link-state route synchronization across multiple segments of a network called Areas within a single autonomous system. OSPF determines best paths utilizing cost based on bandwidth and Dijkstra's SPF algorithm.
BGP	Border Gateway Protocol is the routing protocol used between internet autonomous systems and is therefore considered the routing protocol of the internet. BGP operates as a path-vector protocol at layer 4 (TCP 179) adding neighbors manually to synchronize extremely large routing tables as the only protocol that can scale to millions of routes.
ECMP	Equal Cost Multipath is a trait shared by routing protocols and static routes where traffic sessions are shared among multiple paths as long as the cost is equal.
ACL	Access Control Lists
RADIUS	Remote Access Dial-In User Service is a networking protocol and service to provide Authentication, Authorization, and Accounting (AAA) leveraging UDP ports 1812 and 1813.
TACACS+	Terminal Access Controller Access Control Server refers to a security protocol to provide centralized AAA access with additional security features specific to management.
802.1x	The IEEE standard providing authentication for a user or device attempting to access switchports or wireless. 802.1x represents the suite of EAP options available to supplicant to connect over a local area network to authenticate centrally with a RADIUS or TACACS+ server.
VRRP	A first hop redundancy protocol standardized in RFC 5798 allowing multiple layer three devices to share a virtual address acting as a redundant gateway within a vlan subnet.

UDLD	A layer two protocol used in conjunction with the layer one auto-negotiation of switchports by detecting neighbor identities an shutting down ports that are misconnected. Typically utilizes to safeguard fiber that may experience a unidirectional failure or be misconnected.
LACP	Link Aggregation Control Protocol that dynamical allows bundles of up to eight active ports into a single logical channel for link redundancy without relying on STP.
VSF	Aruba Virtual Switching Framework is a front side stacking protocol that results in a single logical switch with up to 10 members depending on switch model.
802.1p	An IEEE standard providing QoS tagging within the VLAN 3-bit class of service (CoS) field. This allows 0-7 priority queues per port that can differentiate traffic at layer two.
CoS	Class of Service is the three bit field in the VLAN header labeled PRI. See 802.1p.
Rate Limiting	Limitations on traffic bandwidth typically measured using a token bucket system. Exceeding the allocated bandwidth or tokens results in a chosen policy to be applied such as remarking traffic to a lower priority or dropping the offending traffic all together.
PIM	Protocol Independent Multicast is used to route multicast traffic efficiently between receivers or hosts and discovered sources of multicast traffic or servers. PIM is often deployed using either Dense Mode (DM) or Sparse Mode (SM).
SSH	Secure Shell used to securely access remote the terminal of remote devices instead of Telnet.
SSL/TLS	Secure Socket Layer or the improved Transport Layer security are transport protocols for establishing authenticated and encrypted links between compute devices.
SNMPv3	The latest version of Simple Network Management Protocol using polling the defined Management Information Bases (MIBs) of network devices to determine their status. Version 3 is noted for included cryptographically strong authentication and encryption as options.
Strict Priority	A QoS policy that empties the higher priority queues entirely before forwarding a lower queue. This often results in starving lower traffic queues; a solution to this problem being to rate limit an strict-priority assigned queues.
DWRR	Deficit Weighted Round Robin allows WRR (weighted round robin) queues to be aware of bandwidth that may normal be unused to then be shared with lower priority queues.
Console Port	Dedicated out of band management port using either a separate RJ-45 port of a new Type-C connector.
Operator	Aruba CLI prompt indicated by a ">" prompt. Operator context restricts commands to read-only access.
Manager	Aruba CLI prompt indicated by a "#" prompt. Manager contact allows read-write access.
Global Configuration	CLI mode with privilege to make configuration changes.
LLDP	Link Layer Discovery Protocol is a standard protocol used for discovering connected device information and capabilities if that device is local to the other end of a link.
ICMP	Internet Control Messaging Protocol standard (RFC 792) for layer three devices to exchange error messages and operational information between each other.

Ping	An ICMP application used for determining if a destination is reachable through the use of Echo request messages.
Traceroute	An ICMP application that leverages pings to test hop by hop reachability and path determination towards a destination.
Collision Domain	The segment of a network connected by a shared medium such as a hub that is affected when a collision occurs. Routers, switches, bridges, and VLANs are all separate collision domains.
Broadcast Domain	Defines the group of devices that are on the same network segment that are capable of receiving and responding broadcast frames.
Forwarding Table	The FIB or Forwarding Information Base of switches that build a table of known mac-address to the port they can be reached on.
CSMA/CD	Ethernet Carrier Sense Multiple Access Collision Detection mechanism used by Compute NICs and Switchports operating in half-duplex to detect and recover from frame collisions.
Hub	Multi-port repeater at layer 1 device that receives bits and repeats them out all ports including the one receiving the signal in the first place.
802.1Q	The "Dot1Q" IEEE standard covering tagging of frames using a 12 bit field to indicate the VLAN ID or VID.
ARP Cache	Address Resolution Protocol builds a cache table of layer three destinations to the mac-address they can be reached on.
Native Vlan	The vlan ID a switch will assign to any frame sent or received without an indicated vlan ID. This is typically the traffic on a tagged port that is left untagged.
Flood	A switch behavior to handle BUM traffic when the destination mac address is not found in the FIB table.
BUM traffic	Broadcast, Unknown unicast, and Multicast traffic. These traffic types typical result in flooding by the switch to discover the destination.
Loop	A condition at layer 2 or 3 where redundant paths can result in traffic propagating indefinitely and overwhelming the network. Loop protection mechanisms such as spanning-tree and TTL fields can prevent this flooding from impacting the network.
Broadcast Storm	When broadcast traffic is flooded out a loop in the network and no protection mechanisms are in place resulting in high CPU cycles and overwhelmed bandwidth.
Bridge	A term synonymous with a switch in modern networks.
Root Bridge or Switch	This is the switch in spanning-tree that is elected as the start of shortest path "trees" calculation to create loop-free pathways for all traffic to use the root elected bridge.
BPDU	Bridge Protocol Data Unit is a messaging protocol used by all versions of spanning tree to exchange control messages every 2 seconds for path calculation and root tree election. With RSTP and newer protocols, BPDUs are leveraged as keepalives.
BID	Bridge Identifier is a unique identifier used in spanning-tree composed of 2-bytes for priority and a 6-byte bridge mac-address as a tie breaker. Default BID is 32768 where lower is preferred.
STP	802.1d standard as the original version of spanning tree protocol developed to build and maintain redundant loop-free networks.
Designated Ports	A spanning-tree port assignment set to a forwarding state. Each segment will have one designated port. All ports on the root bridge are designated.
Root Ports	A spanning-tree port assignment where this port is the best path based on bandwidth to reach the root bridge and set to forwarding.

Alternate Ports	A spanning-tree port is not be best path to the root however can become the best path should the root port fail. Alternate ports are set to a discarding state.
Backup Ports	A spanning-tree port like an Alternate port that is set to discarding. Backup ports unlike alternate ports, to not offer the switch its best path to the root and only become designated if the other designated port on the segment fails.
Edge Ports	A spanning-tree enabled port that connects to endpoints and should not receive switch to switch BPDUs.
TC-BPDU	In spanning-tree if a topology change (TC) is detected when ports move to different state or if BPDUs are not longer received then a
CST	Common Spanning Tree defines one instance of spanning tree regardless of the number of VLANs.
LAG	Link Aggregation Group is a bundled group of multiple physical interfaces into a single logical interface. LAGs allow all links to be active and loop-free by using a hash algorithm to assign unique sessions to a unique physical link in the LAG therefore avoiding reliance on spanning-tree protocols.
Dynamic LAG	Used to establish and maintain a LAG between two switches that support LACP.
Load Sharing	A method of sharing multiple links such as with MSTP, ECMP, or LAGs where the goal is redundancy.
Load Balancing	Appliances where traffic is measured to assign traffic across multiple pathways where the focus is on balancing or weighting traffic based on bandwidth to achieve a desired distribution of traffic.
LACP Hash	The algorithm chosen to grant a more balanced distribution of traffic based on source and/or destination of Layer 2, Layer 3, or Layer 4 traffic types.
VRF	Virtual Route Forwarder is a separate virtual routing instance inside a single physical router.
Inter-VLAN Routing	Inter-VLAN routing connects separate VLANs into a routed internetwork of communicating devices.
DHCP Helper-Address	The ip helper-address assigned to gateway interface of a subnet in order to relay DHCP requests between a host and a remote DHCP server.
Routing	The activity of layer 3 devices analyzing a packet to determine what next hop address to forward traffic and out of which interface.
IPv4	The layer three connectionless protocol used to determine traffic destinations. IPv4 still bares the majority of network traffic in the world utilizing 32-bit addressing.
IPv6	The next layer three version of IP that expands addressing to 128-bit addresses, has native support for encryption, removes checksums for efficiency, and utilizing an addressing scheme that is geographically significant.
Network ID	The portion of an IP address indicated by the subnet mask as to the bit identifying the network the address is a part of.
Host ID	The less significant part of a network IP address (the right side of where the subnet mask separates network bit from host bits) that identifies an individual host bits.
IP Route	A entry in the routing-table indicating what networks are known and how to reach them by a next hop address and/or the local interface to send packets.
Next Hop IP Address	The IP address of the connected router on the network that can route packets to the destination.

Default Gateway	The device that routes all network destinations for the endpoint devices in a broadcast domain when there does not exist a more specific route.
Multi-layer Interface	An interface that carries a layer two and layer three address. In multi-layer switches, a switches virtual interface (SVI) assigned to a vlan serves this purpose.
Routing Table	The Routing Information Base (RIB) table is table that informs Routers and Multilayer switches of the best path to any given destination by storing IP Routes.
FHRP	First Hop Redundancy Protocol is a set of protocols that create a single coordinated gateway from two or more physical routers sharing the same subnet for redundancy.
VRRP	Virtual Router Redundancy Protocol is a FHRP protocol that is standardized (RFC 5798) using a Master, Backup, and Standby router architecture to provide a single VIP address to hosts.
VRID	Virtual Router ID is a unique number assigned in VRRP to refer to a Group ID matching which VLAN(s) applies to which instance number to allow for load sharing between VRRP gateways.
VIP	Virtual IP Address used in VRRP is the gateway assigned IP address per subnet and shared across VRRP members using a well known virtual MAC address of 00:00:5e:00:00:XX (XX being the VRID or group ID).
VRRP Preemption	The VRRP process of allowing a high priority member to resume the Master role in the event of an election due to network failure and resumption.
Network Class	The IPv4 address space broken into Classes assigned by letters A-E. Classes are assigned by the first four most significant bits in an address. For example, the address 192.168.1.1 is 32-bits where the first four bits are 1100. The lead 0 in the third MSB placement indicates a Class C address.
Unicast Address	The range of addresses found in Classes A, B, and C typically used for unicast traffic either public, private or reserved liked loopback and APIPA.
Multicast Address	The range of addresses found in the Class D range.
Octet	An 8-bit abbreviation ranging from 0-255.
Bit	A one or a zero value.
Subnet Mask	Applies to the address to split the network portion and a host portion.
Reserved Addresses	Reverencing the two reserved addresses within a network or subnet that start with the Network ID at as the first address in a range and the Local Broadcast Address as the last address. These are normally unused for host assignments. An example for 10.1.1.0/24 would be 10.1.1.0 as the Network ID and 10.1.1.255 as the Local Broadcast.
Private IPv4 Address	Defined in RFC 1918 as the portion of each Class reserved for private organizations and should never be seen on the public internet. Class A = 10.0.0.0/8, Class B= 172.16.0.0/12, and Class C= 192.168.0.0/16.
Public IPv4 Address	Addresses reserved with the Internet Assigned Number Authority (IANA) to provide a unique identifier for every host on the internet. This scope applies to Classes A,B, and C unicast traffic excluding private or reserved addresses.
Classful	A method of allocating a fixed number of hosts for each network by using only the default subnet mask by class. Class A uses a /8, Class B uses a /16, and Class C uses a /24. Considered inefficient today in favor of classes networks that allow for variable length subnet masks (VLSM) such as /27 or /12 as an example.

Fixed Subnet Mask	Using a subnet mask such as /24 for all subnets within an organization for simplicity at the cost of efficiency.
VLSM	Variable Length Subnet Masks utilize a large enough subnet mask to allow for the necessary hosts required now and in the anticipated future; for example /30s for point to point networks that will only ever require two addresses.
Static Routes	Routes added to the routing table by an administrator rather than a dynamic routing protocol.
Administrative Distance	A trust rating to determine which route is used in the routing table when multiple sources claim to provide the best way to the same route. The range is 0-255 with lower being preferred.
Floating Routes	A administrator providing a standby redundant route by adding another route with a less desirable (higher) priority over the existing best route.
IGP	Interior Gateway Protocol are routing protocols that run and are managed inside a single AS. Examples include OSPF, RIP, and IS-IS.
EGP	Exterior Gateway Protocols are enterprise-class routing protocols used between AS's. BGP is the only example commonly used today.
Distance Vector Routing Protocols	An early method of routing used by RIPv2 and Cisco's proprietary EIGRP where each router advertises its metric distance from a network and the vector or director packets must travel to get to the network.
Link-state routing protocols	Routing protocols that possess a shared database of all routes within a given network. Considered more robust and allow for greater features than Distance Vector Routing protocols. Examples include OSPF and IS-IS.
AS	Autonomous System is a collection of routers under a common administrative domain.
ISP	Internet Service Provider is a transitive network providing customer traffic access to the public internet. Used to provide routing between separate AS's.
WAN	Wide Area Network connecting remote customer routers across a service provider that creates a transit network using tunneling to appear as a "pseudo-wire" connection between two or more customer edge routers while hiding the providers equipment.
Stacking	Allows you to manage a group of switches as a single device – a Virtual Switch.
VSF	Aruba Virtual Switching Framework (VSF) defines a virtual switch, which is composed by individual physical switches interconnected using Ethernet links. All member switches share a single control plane.
VSX	Aruba Virtual Switching Extension (VSX) is a virtualization technology for ArubaOS-CX switches to provide a more resilient and controlled stacking solution.
Data Plane	Sends and receives frames using specialized hardware called ASICs to build the forwarding table. (See FIB).
Control Plane	Logic used to determine what to do with the date that a switch or router has received. Used to process routing, switching, security, and flow optimization that results in updating the Data Plane.
Management Plane	A separate access method from normal traffic to allow for configuration and monitoring of a device. Often Out-Of-Band management is achieved by a separate physical ports, vlan, or VRF.

ASIC	Application Specific Integrated Circuits are specialized processors utilized by switches to provide line-rate forwarding and processing of frames.
MAD	Multi-Active Detection used between stacked switches to determine if the other member of the broken stack is still forwarding in order to avoid duplication of packets.
Router ID	A 32-bit IP address in dotted-decimal notation used by OSPFv2 to identify the router in the LSDB
OSPF Hello	Messages exchanged between OSPF neighbors to establish adjacency and become neighbors. Must agree on Area ID and Type, Hello and Dead timers, and Authentication.
Passive Interface	An OSPF assignment where the interface will stop sending and receiving OSPF packets out. However, it continues to advertise about that interface to other routes
LSDB	The OSPF topology database – a list of every single link (subnet), every single router, and exactly how those routers and subnets are interconnected.
LSA	OSPF Link State Advertisements of known routes and costs associated with the route. LSAs are exchanged using multicast by default.
OSPF Neighbor Table	A list of directly attached neighbor adjacencies for exchanging OSPF routes. Neighbors must agree on Hellos and go through the states of: Down, INIT, 2-WAY, Exstart, Exchange, Loading, and Full.
SPF	Short Path First is the OSPF Dijkstra algorithm to determine cost to reach a destination based upon each link cost. Cost is measure by dividing the link bandwidth into a standard reference bandwidth.
OSPF Point to Point	OSPF network type between only two peers on a link.
OSPF Broadcast	Two or more OSPF peers on a shared link.
OSPF DR/BDR	OSPF Designated Router is the router that is elected to coordinate peer interactions in the broadcast domain. A backup is also elected to take over if the DR fails.
Area	An Area is a group of OSPF routers that share the same Link State database. All router must be part of an area.
Backbone Area	The OSPF area 0 that is mandatory in an AS.
LSA type 1	Used by OSPF routers to announce themselves to the DR. Used area wide.
LSA type 2	Used by OSPF when there is a Broadcast network type link, to announce an elected DR and BDR. Intra-area LSA.
OOBM	Out of Band Management uses an physical port such as the management port, a dedicated vlan, or a VRF to send and receive management traffic to provide better device security.
Management VRF	A dedicated VRF for the purpose of management. In AOS-CX, the management port must be a member of this VRF only.
AAA	The standard description of Authentication, Authorization, and Accounting to management security in a network.
Authentication	Part of AAA where identity is established through one or more factors of authentication.
Authorization	Part of AAA where access rights are assigned and is usually based upon the Authentication results (or lack thereof).
Accounting	Part of AAA to track what the user of device did on the network.
RBAC	Role Based Access Control defines user authorization group creation and assigning users to those groups.

SNMP	Simple Network Management Protocol provides centralized polling, trapping, and management of events from multiple network devices to a single server.
Checkpoint	A snapshot of the switch running configuration and its relevant metadata. Available on Aruba CX devices.
NetEdit	An Aruba server appliance that provides for coordinated switch configuration, monitoring, and troubleshooting. It provides for intelligent, error-free configuration with validation for consistency and compliant, and the automation of multi-device change workflows without the overhead of programming.
Seed Device	Used as a starting point for network discovery to branch out to connected devices based on hops.

3333 Scott Blvd, Santa Clara, CA 95054
TEL: 408.227.4500 | FAX: 408.227.4550
www.ARUBANETWORKS.com

EDU-CXF-ILT-v20.21