



ACI Advanced Monitoring and Troubleshooting

Sadiq Memon, CCIE® No. 47508

Joseph Ristaino, CCIE® No. 41799

Carlo Schmidt, CCIE® No. 41842

Forewords written by **Yusuf Bhaiji**, Director of Certifications, Cisco Systems;
and **Ronak Desai**, VP of Engineering for the Data Center
Networking Business Unit, Cisco Systems

ACI Advanced Monitoring and Troubleshooting

Sadiq Memon (CCIE No. 47508)

Joseph Ristaino (CCIE No. 41799)

Carlo Schmidt (CCIE No. 41842)

Cisco Press

221 River St.

Hoboken, NJ 07030 USA

ACI Advanced Monitoring and Troubleshooting

Table of Contents

Cover

Title Page

Copyright Page

About the Authors

Dedications

Acknowledgments

Contents at a Glance

Contents

Foreword

Foreword

Introduction

Part I: Introduction to ACI

Chapter 1 Fundamental Functions and Components of Cisco ACI

ACI Building Blocks

Hardware Specifications

Nexus 9000 Platform

APIC Controller

ACI Key Concepts

Control Plane

Data Plane

VXLAN

Tenant

Table of Contents

VRF

Application Profile

Endpoint Group

Contracts

Bridge Domain

External Routed or Bridged Network

Summary

Review Key Topics

Review Questions

Chapter 2 Introduction to the ACI Policy Model

Key Characteristics of the Policy Model

Management Information Tree (MIT)

Benefits of a Policy Model

Logical Constructs

Tenant Objects

VRF Objects

Application Profile Objects

Endpoint Group Objects

Bridge Domain and Subnet Objects

Bridge Domain Options

Contract Objects

Labels, Filters, and Aliases

Contract Inheritance

Contract Preferred Groups

vzAny

Outside Network Objects

Physical Construct

Access Policies

Switch Policies

Interface Policies

Global Policies

VLAN Pools

Table of Contents

Domains

Attachable Access Entity Profile

Managed Object Relationships and Policy Resolution

Tags

Default Policies

How a Policy Model Helps in Diagnosis

Summary

Review Key Topics

Review Questions

Chapter 3 ACI Command-Line Interfaces

APIC CLIs

NX-OSStyle CLI

Bash CLI

ACI Fabric Switch CLIs

iBash CLI

VSH CLI

VSH_LC CLI

Summary

Reference

Chapter 4 ACI Fabric Design Options

Physical Design

Single- Versus Multiple-Fabric Design

Dark Fiber

Dense Wavelength-Division Multiplexing (DWDM)

Ethernet over MPLS (EoMPLS) Pseudowire

Multi-Pod

ACI Multi-Pod Use Cases

ACI Multi-Pod Scalability

Inter-Pod Connectivity Deployment Considerations

APIC Cluster Deployment Considerations

Multi-Site

Table of Contents

- Cisco ACI Multi-Site Orchestrator
- Cisco ACI Multi-Site Deployment Considerations
- Migration Scenarios
- Deployment Best Practices
- General Best Practices for Cisco ACI Multi-Site Design
- Remote Leaf
- Hardware and Software Support
- Recommended QOS Configuration for a Remote Leaf Solution
- Discovery of a Remote Leaf
- Remote Leaf Control Plane and Data Plane
- Remote Leaf Design Considerations
- ACI Multi-Pod and Remote Leaf Integration

Logical Design

- Design 1: Container-as-a-Service Using the OpenShift Platform and Calico CNI
 - Business Case
 - Design Solution
- Design 2: Vendor-Based ERP/SAP Hana Design with ACI
 - Business Case
 - Design Solution
- Design 3: vBrick Digital Media Engine Design with ACI
 - Business Case
 - Design Solution

Summary

Review Key Topics

Review Questions

Chapter 5 End Host and Network Connectivity

End Host Connectivity

- VLAN Pool
- Domain
- Attachable Access Entity Profiles (AAEPs)
- Switch Policies
- Switch Policy Groups
- Switch Profiles

Table of Contents

Interface Policies
Interface Policy Groups
Interface Profiles
Virtual Port Channel (VPC)
Configuring VPC
Defining the VPC Domain
Creating an Interface Policy
Creating a Switch Profile
Port Channel
Configuring a Port Channel
Access Port
Configuring an Access Port
Best Practices in Configuring Access Policies
Policy Best Practices
Domain Best Practices
AAEP Best Practices
Compute and Storage Connectivity
FEX Connectivity
Cisco Blade Chassis Servers UCS B-Series
Standalone Rack-Mount Servers
Connecting Storage in ACI
L4/L7 Service Device Connectivity
Connecting Firewalls
Connecting Load Balancers

Network Connectivity

Connecting an External Bridge Network
Extending EPGs Outside the ACI Fabric
Extending an ACI Bridge Domain Outside the Fabric
Connecting an External Routed Network
External Layer 3 Supported Routing Protocols
Configuring MP-BGP Spine Route Reflectors
Configuring External Routed Networks
GOLF

Table of Contents

Network Connectivity Between Pods and Sites

IPN Connectivity Considerations for Remote Leafs

Diagnosing Connectivity Problems

Summary

Review Questions

Chapter 6 VMM Integration

Virtual Machine Manager (VMM)

VMM Domain Policy Model

VMM Domain Components

VMM Domains

VMM Domain VLAN Pool Association

Attachable Access Entity Profile Association

VMM Domain EPG Association

EPG Policy Resolution and Deployment Immediacy

VMware Integration

Prerequisites for VMM Integration with AVS or VDS

Guidelines and Limitations for VMM Integration with AVS or VDS

ACI VMM Integration Workflow

Publishing EPGs to a VMM Domain

Connecting Virtual Machines to the Endpoint Group Port Groups on vCenter

Verifying VMM Integration with the AVS or VDS

Verifying the Virtual Switch Status

Verifying the vNIC Status

Microsoft SCVMM Integration

Mapping ACI and SCVMM Constructs

Mapping Multiple SCVMMs to an APIC

Verifying That the OpFlex Certificate Is Deployed for a Connection from the SCVMM to the
APIC

Verifying VMM Deployment from the APIC to the SCVMM

OpenStack Integration

Extending OpFlex to the Compute Node

ACI with OpenStack Physical Architecture

Table of Contents

- OpFlex Software Architecture
- OpenStack Logical Topology
- Mapping OpenStack and ACI Constructs
- Prerequisites for OpenStack and Cisco ACI
- Guidelines and Limitations for OpenStack and Cisco ACI
- Verifying the OpenStack Configuration
- Configuration Examples for OpenStack and Cisco ACI

Kubernetes Integration

- Planning for Kubernetes Integration
- Prerequisites for Integrating Kubernetes with Cisco ACI
- Provisioning Cisco ACI to Work with Kubernetes
- Preparing the Kubernetes Nodes
- Installing Kubernetes and Cisco ACI Containers
- Verifying the Kubernetes Integration

OpenShift Integration

- Planning for OpenShift Integration
- Prerequisites for Integrating OpenShift with Cisco ACI
- Provisioning Cisco ACI to Work with OpenShift
- Preparing the OpenShift Nodes
- Installing OpenShift and Cisco ACI Containers
- Updating the OpenShift Router to Use the ACI Fabric
- Verifying the OpenShift Integration

VMM Integration with ACI at Multiple Locations

- Multi-Site
- Multiple Virtual Machine Managers Across Sites
- Single Virtual Machine Manager Across Sites
- Remote Leaf

Summary

Chapter 7 L4/L7 Service Integration

Service Insertion

The Service Graph

- Managed Mode Versus Un-Managed Mode

Table of Contents

L4L7 Integration Use Cases

How Contracts Work in ACI

The Shadow EPG

Configuring the Service Graph

Step 1: Create an L4L7 Device

Step 2: Create a Service Graph Template

Step 3: Deploy the Service Graph from the Template

Step 4: Configure the L4L7 Parameters (Managed Mode Only)

Verifying the Service Graph Configuration

Service Graph Design and Deployment Options

Firewall as Default Gateway for Client and Server (Routed Mode)

Firewall Not the Default Gateway for Clients (Routed Mode)

Route Peering with a Firewall (Routed Mode)

Service Graph with Firewall (Transparent Mode)

Service Graph with ADC (One-Arm Mode with S-NAT)

Service Graph with ADC (Two-Arm Mode)

Service Graph with Two Service Nodes (Firewall with NAT and ADC in Two-Arm Mode)

Service Graph with Two Service Nodes (Firewall with No NAT and ADC in Two-Arm Mode)

Service Graph with Two Service Nodes (Firewall with No NAT and ADC in One-Arm Mode)

Service Graph with an Intrusion Prevention System (IPS)

Policy-Based Redirect (PBR)

PBR Design Considerations

PBR Design Scenarios

PBR Service Graph with an ADC (One-Arm Mode and No S-NAT)

PBR Service Graph with a Firewall (Two-Arm Mode and Routed)

Configuring the PBR Service Graph

Service Node Health Check

L4L7 PBR Tracking

L4L7 PBR Threshold

L4L7 PBR Health Groups

Common Issues in the PBR Service Graph

Unnecessary Layer 2 Traffic Redirection Toward the Service Node

Inability to Ping the Consumer Connector

Table of Contents

Routing on a Service Node

L4/L7 Service Integration in Multi-Pod and Multi-Site

Multi-Pod

Anycast Services in Multi-Pod

Multi-Site

Review Questions

Chapter 8 Automation and Orchestration

The Difference Between Automation and Orchestration

Benefits of Automation and Orchestration

Example 1

Example 2

REST API

Automating Tasks Using the Native REST API: JSON and XML

API Inspector

Object (Save As)

Visore (Object Store Browser)

MOQuery

Automation Use Cases

Automating Tasks Using Ansible

Ansible Support in ACI

Installing Ansible and Ensuring a Secure Connection

APIC Authentication in Ansible

Automation Use Cases

Use Case 1

Use Case 2

Orchestration Through UCS Director

Management Through Cisco UCS Director

Automation and Orchestration with Cisco UCS Director

Automation Use Cases

Summary

Review Questions

Part II: Monitoring and Management Best Practices

Table of Contents

Chapter 9 Monitoring ACI Fabric

Importance of Monitoring

Faults and Health Scores

Faults

Health Scores

Health Score Used in Proactive Monitoring

Health Score Used in Reactive Monitoring

Health Score with Interface Errors

ACI Internal Monitoring Tools

SNMP

Interface Failures Example

Syslog

Example: Leaf Membership Failure

Example: Spine/IPN Failure

NetFlow

Example: Network Visibility on a Border Leaf

ACI External Monitoring Tools

Network Insights

Network Insights for Resources (NIR)

Network Insights Advisor (NIA)

Example: Application Intermittent Disconnect Issue (Standalone Compute)

Example: Application Connectivity Issue (Virtual Compute)

Network Assurance Engine

NAE Installation

NAE Configuration and Initial Setup

Example: Subnet Reachability Issue

Tetration

Software Agents

Hardware Agents

Tetration Installation and Configuration

Tetration System Monitoring

Configuring Email Alerts

Enabling Syslog

Table of Contents

- Tetration Scopes
- Tetration Applications
- Tetration Code Upgrades
- Tetration Patch Upgrade
- Tetration Cluster Reboot
- Tetration Cluster Shutdown
- Example: Workload Security with Tetration

Monitoring Through the REST API

- Monitoring an APIC
- Monitoring CPU and Memory
- Monitoring Disk Utilization
- Monitoring Interfaces
- Monitoring the APIC Cluster State
- Monitoring Leafs and Spines
- Monitoring CPU Utilization
- Monitoring Memory Utilization
- Monitoring Power Supply Unit (PSU) Status
- Monitoring Fan Status
- Monitoring Module Status
- Monitoring Leaf/Spine Membership Status in a Fabric
- Monitoring Interface Status
- Monitoring Applications
- Monitoring Application Traffic Status
- Monitoring External Network Connectivity
- Monitoring the PBR Service Graph

Summary

Review Questions

Chapter 10 Network Management and Monitoring Configuration

Out-of-Band Management

- Creating Static Management Addresses
- Creating the Management Contract
- Choosing the Node Management EPG
- Creating an External Management Entity EPG

Table of Contents

Verifying the OOB Management Configuration

In-Band Management

Creating a Management Contract

Creating Leaf Interface Access Policies for APIC INB Management

Creating Access Policies for the Border Leaf(s) Connected to L3Out

Creating INB Management External Routed Networks (L3Out)

Creating External Management EPGs

Creating an INB BD with a Subnet

Configuring the Node Management EPG

Creating Static Management Addresses

Verifying the INB Management Configuration

AAA

Configuring Cisco Secure ACS

Configuring Cisco ISE

Configuring AAA in ACI

Recovering with the Local Fallback User

Verifying the AAA Configuration

Syslog

Verifying the Syslog Configuration and Functionality

SNMP

Verifying the SNMP Configuration and Functionality

SPAN

Access SPAN

Fabric SPAN

Tenant SPAN

Ensuring Visibility and Troubleshooting SPAN

Verifying the SPAN Configuration and Functionality

NetFlow

NetFlow with Access Policies

NetFlow with Tenant Policies

Verifying the NetFlow Configuration and Functionality

Summary

Table of Contents

Part III: Advanced Forwarding and Troubleshooting Techniques

Chapter 11 ACI Topology

Physical Topology

APIC Initial Setup

Fabric Access Policies

Switch Profiles, Switch Policies, and Interface Profiles

Interface Policies and Policy Groups

Pools, Domains, and AAEPs

VMM Domain Configuration

VMM Topology

Hardware and Software Specifications

Logical Layout of EPGs, BDs, VRF Instances, and Contracts

L3Out Logical Layout

Summary

Review Key Topics

References

Chapter 12 Bits and Bytes of ACI Forwarding

Limitations of Traditional Networks and the Evolution of Overlay Networks

High-Level VXLAN Overview

IS-IS, TEP Addressing, and the ACI Underlay

IS-IS and TEP Addressing

FTags and the MDT

Endpoint Learning in ACI

Endpoint Learning in a Layer 2Only Bridge Domain

Council of Oracle Protocol (COOP)

Updating the Managed Object (MO) Tree

Endpoint Learning in a Layer 3Enabled Bridge Domain

Fabric Glean

Remote Endpoint Learning

Endpoint Mobility

Table of Contents

Anycast Gateway

Virtual Port Channels in ACI

Routing in ACI

Static or Dynamic Routes

Learning External Routes in the ACI Fabric

Transit Routing

Policy Enforcement

Shared Services

L3Out Flags

Quality of Service (QoS) in ACI

Externally Set DSCP and CoS Markings

EPG QoS

Custom QoS Policy

Contract QoS

CoS Preservation in ACI

iTraceroute Class

QoS and Multi-Pod

DSCP Class-to-CoS Translation Policy

Multi-Pod

Multi-Site

Remote Leaf

Forwarding Scenarios

ARP Flooding

Layer 2 Known Unicast

ARP Optimization

Layer 2 Unknown Unicast Proxy

L3 Policy Enforcement When Going to L3Out

L3 Policy Enforcement for External Traffic Coming into the Fabric

Route Leaking/Shared Services

Consumer to Provider

Provider to Consumer

Multi-Pod Forwarding Examples

Table of Contents

ARP Flooding

Layer 3 Proxy Flow

Multi-Site Forwarding Examples

ARP Flooding

Layer 3 Proxy Flow

Remote Leaf

ARP Flooding

Layer 3 Proxy Flow

Summary

Review Key Topics

References

Review Questions

Chapter 13 Troubleshooting Techniques

General Troubleshooting

Faults, Events, and Audits

moquery

iCurl

Visore

Infrastructure Troubleshooting

APIC Cluster Troubleshooting

Fabric Node Troubleshooting

How to Verify Physical- and Platform-Related Issues

Counters

CPU Packet Captures

ASIC

ASIC Interface

Application

SPAN

Troubleshooting Endpoint Connectivity

Endpoint Tracker and Log Files

Enhanced Endpoint Tracker (EPT) App

Rogue Endpoint Detection

Table of Contents

Troubleshooting Contract-Related Issues

- Verifying Policy Deny Drops

Embedded Logic Analyzer Module (ELAM)

- Summary

- Review Key Topics

- Review Questions

Chapter 14 The ACI Visibility & Troubleshooting Tool

- Visibility & Troubleshooting Tool Overview

- Faults Tab

- Drop/Stats Tab

 - Ingress/Egress Buffer Drop Packets

 - Ingress Error Drop Packets Periodic

 - Storm Control

 - Ingress Forward Drop Packets

 - Ingress Load Balancer Drop Packets

- Contract Drops Tab

 - Contracts

 - Contract Considerations

- Events and Audits Tab

- Traceroute Tab

- Atomic Counter Tab

- Latency Tab

- SPAN Tab

- Network Insights Resources (NIR) Overview

- Summary

Chapter 15 Troubleshooting Use Cases

- Troubleshooting Fabric Discovery: Leaf Discovery

 - Solution

- Troubleshooting APIC Controllers and Clusters: Clustering

 - Solution

- Troubleshooting Management Access: Out-of-Band EPG

Table of Contents

Solution

Troubleshooting Contracts: Traffic Not Traversing a Firewall as Expected

Solution

Troubleshooting Contracts: Contract Directionality

Solution

Troubleshooting End Host Connectivity: Layer 2 Traffic Flow Through ACI

Solution

Troubleshooting External Layer 2 Connectivity: Broken Layer 2 Traffic Flow Through ACI

Solution 1

Solution 2

Troubleshooting External Layer 3 Connectivity: Broken Layer 3 Traffic Flow Through ACI

Solution

Troubleshooting External Layer 3 Connectivity: Unexpected Layer 3 Traffic Flow Through ACI

Solution

Troubleshooting Leaf and Spine Connectivity: Leaf Issue

Solution

Troubleshooting VMM Domains: VMM Controller Offline

Solution 1

Solution 2

Troubleshooting VMM Domains: VM Connectivity Issue After Deploying the VMM Domain

Solution 1

Solution 2

Solution 3

Troubleshooting L4L7: Deploying an L4L7 Device

Solution

Troubleshooting L4L7: Control Protocols Stop Working After Service Graph Deployment

Solution

Table of Contents

Troubleshooting Multi-Pod: BUM Traffic Not Reaching Remote Pods

Solution 1

Solution 2

Troubleshooting Multi-Pod: Remote L3Out Not Reachable

Solution

Troubleshooting Multi-Site: Using Consistency Checker to Verify State at Each Site

Solution

Troubleshooting Programmability Issues: JSON Script Generates Error

Solution

Troubleshooting Multicast Issues: PIM Sparse Mode Any-Source Multicast (ASM)

Solution

Summary

Appendix A Answers to Chapter Review Questions

Index