



Deploying ACI

The complete guide to planning,
configuring, and managing
Application Centric Infrastructure

Frank Dagenhardt, CCIE® No. 42081

Jose Moreno, CCIE® No. 16601

With contributions from **Bill Dufresne**, CCIE® No. 4375

Deploying ACI

The complete guide to planning, configuring, and managing Application Centric Infrastructure

Frank Dagenhardt, CCIE No. 42081,
Jose Moreno, CCIE No. 16601,

With contributions from
Bill Dufresne, CCIE No. 4375

Cisco Press

800 East 96th Street

Indianapolis, Indiana 46240 USA

Deploying ACI: The complete guide to planning, configuring, and managing Application Centric Infrastructure

Table of Contents

Cover

Title Page

Copyright Page

Dedication

About the Author

Contents

Introduction

Chapter 1 Youve Purchased ACI. Now What?

- Industry Trends and Transitions

- Next-Generation Data Center Concepts

 - New Application Types

 - Automation, Orchestration, and Cloud

 - End-to-End Security

- Spine-Leaf Architecture

 - Existing Infrastructure and ACI (Places in the Network)

- ACI Overview

- ACI Functional Components

 - Nexus 9500

 - Nexus 9300

 - Application Centric Infrastructure Controllers

- Protocols Enabling the ACI Fabric

Table of Contents

Data Plane Protocols

Control Plane Protocols

Interacting with ACI

GUI

NX-OS CLI

Open REST API

Introduction to the Policy Model

Application Network Profiles and Endpoint Groups

VRFs and Bridge Domains

Fabric Topologies

Single-Site Model

Multi-Pod Model

Multi-Site Model

Summary

Chapter 2 Building a Fabric

Building a Better Network

Fabric Considerations

Roles of a Leaf

Fixed vs. Modular Spine

Integration Planning and Considerations

Security Considerations

Phased ACI Migration

Network-Centric Mode: Single Tenant

Network-Centric Mode: Multiple Tenant

Evolution to Application-Centric Mode

Microsegmentation

Bare-Metal Workloads

Virtualized Workloads

Containers

Table of Contents

Virtual Machine Manager (VMM) Integration

- AVS

- VMware

- Microsoft

- OpenStack

Layer 4-7 Services

- Managed Mode

- Unmanaged Mode

Additional Multisite Configurations

- Cisco ACI Stretched Fabric

- Cisco ACI Multi-Pod

- Cisco ACI Multi-Site

- Cisco ACI Dual-Fabric Design

- Pervasive Gateway

- VMM Considerations

Summary

Chapter 3 Bringing Up a Fabric

Out of the Box

- Suggested Services

- Management Network

- Out-of-Band Network

- In-Band Network

- What to Expect when You Configure a Controller

- Fabric Infrastructure IP Range Recommendations

- Fabric Infrastructure VLAN Recommendations

- Cluster Size and APIC Controller ID

- About High Availability for APIC Cluster

Logging In to the GUI for the First Time

- Basic Mode vs. Advanced Mode

Table of Contents

System Tab

Tenants Tab

Fabric Tab

VM Networking Tab

L4-L7 Services Tab

Admin Tab

Operations Tab

Apps Tab

Discovering the Fabric

Fabric Extenders

Required Services

Basic Mode Initial Setup

Management Network

NTP

Route Reflectors

VLAN Domains

Advanced Mode Initial Setup

Access Policies

VLAN Pools and Domains

Attachable Access Entity Profiles

Interface Policies

Interface Policy Groups

Interface Profile

Switch Profile

Management Network

Fabric Policies

NTP

Route Reflectors

Managing Software Versions

Table of Contents

- Firmware Repository
- Controller Firmware and Maintenance Policy
- Firmware Groups and Policy
- Maintenance Group and Maintenance Policy
- Using the Scheduler

Configuration Management

- Configuration Snapshots
- Configuration Backup

Summary

Chapter 4 Integration of Virtualization Technologies with ACI

Why Integrate Cisco ACI with Virtualization Technologies?

Networking for Virtual Machines and Containers

- Benefits of Cisco ACI Integration with Virtual Switches
- Comparing ACI Integration to Software Network Overlays
- Virtual Machine Manager Domains
- EPG Segmentation and Micro-Segmentation
- Intra-EPG Isolation and Intra-EPG Contracts
- Cisco ACI Integration with Virtual Switches in Blade Systems
- OpFlex
- Deployments over Multiple Data Centers

VMware vSphere

- Cisco ACI Coexistence with the vSphere Standard Switch
- Cisco ACI Coexistence with the vSphere Distributed Switch
- Cisco ACI Integration with the vSphere Distributed Switch
- vCenter User Requirements
- Micro-Segmentation with the VDS
- Blade Servers and VDS Integration
- Cisco ACI Integration with Cisco Application Virtual Switch
- Cisco AVS Installation

Table of Contents

Blade Servers and AVS Integration

Distributed Firewall

Virtual Network Designs with VDS and AVS

Cisco ACI Plug-in for vSphere vCenter Server: Configuring ACI from vCenter

Cisco ACI Coexistence with VMware NSX

Microsoft

Introduction to Microsoft Hyper-V and SCVMM

Preparing for the Integration

Micro-Segmentation

Blade Servers and SCVMM Integration

OpenStack

ML2 and Group-Based Policy

Installing Cisco ACI Integration with OpenStack

Cisco ACI ML2 Plug-in for OpenStack Basic Operations

Cisco ACI ML2 Plug-in for OpenStack Security

Cisco ACI ML2 Plug-in for OpenStack and Network Address Translation

Cisco ACI GBP Plug-in for OpenStack

Docker: Project Contiv

Docker Networking

Kubernetes

Kubernetes Networking Model

Isolation Models

Creating a New EPG for Kubernetes Pods

Assigning a Deployment or a Namespace to an EPG with Annotations

Visibility in ACI for Kubernetes Objects

Public Cloud Integration

Summary

Chapter 5 Introduction to Networking with ACI

Table of Contents

Exploring Networking in ACI

- Groups and Contracts
- Contracts Are ACLs Without IP Addresses
- Filters and Subjects
- Concept of Direction in Contracts
- Understanding the Bidirectional and Reverse Filter Options
- Configuring a Single Contract Between EPGs
- Using vzAny
- Contract Scope
- Contracts and Filters in the Common Tenant
- VRFs and Bridge Domains
- VRF Design Considerations
- Bridge Domain Design Considerations
- VRFs and Bridge Domains in the Common Tenant
- VRFs in the Common Tenant and Bridge Domains in User Tenants
- Layer 3 External Connection in the Common Tenant with VRFs and Bridge Domains in User Tenants
- Ingress Versus Egress Filtering Design Recommendations
- Connecting External Networks to the Fabric
- L2 Connections
- Basic Mode GUI
- Advanced Mode Access Policies

Network-Centric VLAN=BD=EPG

- Applying Policy to Physical and Virtual Workloads
- Moving Devices to the Fabric, VLAN by VLAN
- Unenforced vs. Enforced VRF
- L3 Connections to the Core
- Layer 3 Out and External Routed Networks
- L3 Out Simplified Object Model
- Border Leafs

Table of Contents

Migrating the Default Gateway to the Fabric

Summary

Chapter 6 External Routing with ACI

Layer 3 Physical Connectivity Considerations

Routed Ports Versus Switched Virtual Interfaces

Outside Bridge Domains

Bidirectional Forwarding Detection

Access Port

Port Channel

Virtual Port Channel

Gateway Resiliency with L3 Out

Hot Standby Routing Protocol

Routing Protocols

Static Routing

Enhanced Interior Gateway Routing Protocol

Open Shortest Path First

OSPF Summarization

Border Gateway Protocol

BGP Route Profile

Outbound BGP Policy

BGP Protocol Statistics

External Endpoint Groups and Contracts

External Endpoint Groups

Contracts Between L3 Out EPGs and Internal EPGs

Multitenant Routing Consideration

Shared Layer 3 Outside Connection

Transit Routing

Supported Combinations for Transit Routing

Loop Prevention in Transit Routing Scenarios

Table of Contents

WAN Integration

Design Recommendations for Multitenant External Layer

Connectivity

Quality of Service

User-Defined Classes

Reserved Classes

Classification and Marking

Multicast

Multicast Best-Practice Recommendations

Scenario 1: Leaf Switches Not Based on Cisco Nexus EX Platform

Scenario 2: Leaf Switches Based on Cisco Nexus EX Platform

Scenario 3: Hybrid Fabric with Leaf Switches Both Based on and Not Based on
Cisco Nexus EX Platform

Multicast Configuration Overview

Minimum Multicast Configuration: PIM-ASM

Minimum Multicast Configuration: PIM-SSM

Summary

Chapter 7 How Life Is Different with ACI

Managing Fabrics versus Managing Devices

Centralized CLI

System Dashboard

Tenant Dashboards

Health Scores

Physical and Logical Objects

Network Policies

Fabric-wide Policies

Comparing the ACI Controller to Traditional Network Management Systems

Troubleshooting the Deployment of Global Policies

Configuring Multiple Ports at the Same Time

Table of Contents

Maintaining the Network

- Fault Management
- Faults Across the Network
- Fault Lifecycle
- Immediate Fault Reporting for Change Validation
- Configuration Management
- Evaluating Change Impact
- Configuration Zones: Running Changes Gradually
- Centralized Change Description
- Atomicity of Network Changes
- Configuration Snapshots
- Network Audit Trails
- Upgrading the Software

Breaking the Shackles of IP Design

- Access Control Lists Without IP Addresses
- QoS Rules Without IP Addresses
- QoS Rules Without TCP or UDP Ports

Physical Network Topology

- ACI as a Clos Fabric and Design Implications
- Connecting Endpoints to Leaf Switches
- Scaling an ACI Fabric Means Adding More Leaf Switches
- Fabric Topology and Links
- Individual Device View
- Port View

Changing the Network Consumption Model

Summary

Chapter 8 Moving to Application-Centric Networking

Network-Centric Deployments

- Removing Packet Filtering in Network-Centric Deployments

Table of Contents

Increasing Per-Leaf VLAN Scalability

Looking at the Configuration of a Network-Centric Design

Application-Centric Deployment: Security Use Case

Whitelist vs. Blacklist Models

Enforced vs. Unenforced: ACI Without Contracts

Endpoint Groups as a Zone-Based Firewall

Dynamic EPG Relationships: Micro-Segmentation EPGs

Multiple EPGs in the Same Subnet

Contract Security Model

Inter-EPG Communication

Contract Scope

Contract Subject Settings

Filter Settings

Contract Subject Labels

Contract Inheritance

Stateful Firewalling with Cisco Application Virtual Switch

Intra-EPG Communication

Any EPG

Contract Definition Best Practices to Efficiently Use Resources

Application-Centric Deployment: Operations Use Case

Application-Centric Monitoring

Quality of Service

Impact Analysis

Asset Allocation

Migrating to an Application-Centric Model

Disable Bridge Domain Legacy Mode

Disable VRF Unenforced Mode

Create New Application Profiles and EPGs

Move Endpoints to the New EPGs

Table of Contents

Fine-Tune Security Rules

How to Discover Application Dependencies

Focus on New Applications

Migrate Existing Applications

Legacy Application Dependency Mapping

Cisco Tetration Analytics

Summary

Chapter 9 Multi-Tenancy

The Need for Network Multi-Tenancy

Data-Plane Multi-Tenancy

Management Multi-Tenancy

Multi-Tenancy in Cisco ACI

Security Domains

Role-Based Access Control

Physical Domains

Logical Bandwidth Protection Through Quality of Service

What Is a Tenant? What Is an Application?

Logical Separation for Lines of Business

Logical Separation for Security or Compliance

Moving Resources to Tenants

Creating the Logical Tenant Structure

Implementing Management Multi-Tenancy

Moving EPGs and Contracts

Exporting and Importing Contracts for Inter-Tenant Communication

Implementing Data-Plane Multi-Tenancy

When to Use Dedicated or Shared VRFs

Multi-Tenant Scalability

External Connectivity

Shared External Network for Multiple Tenants

Table of Contents

Inter-Tenant Connectivity

- Inter-VRF External Connectivity

- Inter-VRF Internal Connectivity (Route Leaking)

L4-7 Services Integration

- Exporting L4-7 Devices

- Multi-Context L4-7 Devices

Use Cases for Multi-Tenancy Connectivity

- ACI as Legacy Network

- Granting Network Visibility to Other Departments

- Network Shared Across Organizations with Shared Services

- External Firewall Interconnecting Multiple Security Zones

- Service Provider

Summary

Chapter 10 Integrating L4-7 Services

Inserting Services

- How We Do It Today

- Managed vs. Unmanaged

- Ecosystem Partners

- Management Model

- Functional Profiles

Security for All Hosts

- Building an End-to-End Security Solution

- Integrating Firewalls

- Service Node Failover

- Deploying Clustering for Physical Appliances (Cisco ASA Cluster)

- Virtual versus Physical

- Integrating Security Monitoring

- Integrating Intrusion Prevention Systems

- Copy Service

Table of Contents

Integrating Server Load Balancing and ADC

Two-node Service Graph Designs

Summary

Chapter 11 Multi-Site Designs

Bringing Up a Second Site

Stretched Fabric Design

Site-to-Site Connectivity Options

Stretched ACI Fabric Preserves VM Mobility

Loss of a Single APIC

Split Fabric

Standby APIC

Multiple-Fabric Design

Cisco Data Center Interconnect

Transit Leaf and L3 Out Considerations

DCI or Inter-Pod Network Considerations

Multiple Fabric Connectivity Options

Multi-Pod Architecture

ACI Multi-Pod Use Cases and Supported Topologies

ACI Multi-Pod Scalability Considerations

Inter-Pod Connectivity Deployment Considerations

IPN Control Plane

IPN Multicast Support

Spines and IPN Connectivity Considerations

Pod Auto-Provisioning

APIC Cluster Deployment Considerations

Reducing the Impact of Configuration Errors with Configuration Zones

Migration Strategies

Multi-Site Architecture

APIC Versus Multi-Site Controller Functionalities

Table of Contents

Multi-Site Schema and Templates

Multi-Site Use Cases

Stretched Bridge Domain with Layer 2 Broadcast Extension (Option 3)

Stretched Bridge Domain with No Layer 2 Broadcast Extension (Option 2)

Stretched EPG Across Sites (Option 1.1)

Stretched VRF with Inter-Site Contracts (Option 1.2)

Shared Services with Stretched Provider EPG

Multi-Site and L3 Out Considerations

Layer 3 Multicast Deployment Options

Migration of Cisco ACI Fabric to Cisco ACI Multi-Site

Summary

Chapter 12 Troubleshooting and Monitoring

You Have a Poor Health Score. Now What?

NX-OS CLI

Connecting to the Leaf Switches

Linux Commands

Mapping Local Objects to Global Objects

VLAN IDs

Legacy Mode

Port Channels

Some Useful Leaf Commands

ping

Troubleshooting Physical Issues

Troubleshooting Cabling

Troubleshooting Switch Outages

Replacing a Fabric Switch

Troubleshooting Contracts

Troubleshooting Tools in ACI

Hardware Diagnostics

Table of Contents

Dropped Packets: Counter Synchronization

Atomic Counters

Traffic Mirroring: SPAN and Copy Services

SPAN Destination Groups

ERSPAN Types

SPAN Source Groups

Cisco ACI Scalability for SPAN Sessions

Nexus Data Broker

Troubleshooting Wizard

Defining the Troubleshooting Session

Faults in the Troubleshooting Wizard

Statistics in the Troubleshooting Wizard

Contract Information in the Troubleshooting Wizard

Events and Audits in the Troubleshooting Wizard

Traceroute in the Troubleshooting Wizard

Atomic Counters in the Troubleshooting Wizard

Configuring SPAN from the Troubleshooting Wizard

Endpoint Tracker

Effectively Using Your Fabric Resources

Using Traffic Map to Find Bandwidth Bottlenecks

Using Capacity Dashboard to Detect Resource Bottlenecks

Using ACI Optimizer to Plan for Changes

Monitoring Policies and Statistics

SNMP Policies

Syslog Policies

Statistics

Third-Party Monitoring Tools with ACI Support

IBM Tivoli Netcool

SevOne

Table of Contents

ScienceLogic

Splunk

Zenoss

Summary

Chapter 13 ACI Programmability

Why Network Programmability? Save Money, Make Money!

What Is Wrong with Previous Network Automation Concepts?

SNMP

Network Configuration Protocol and YANG

Programming Interfaces and SDKs

What Is REST?

What Is a Software Development Kit?

Cisco ACI Programming Interfaces

Cisco ACI REST API

REST API Authentication

API Inspector

REST API Clients

Using REST APIs in Programming Languages

Cisco ACI Object Model

Debug Information in the GUI

Visore

moquery

Cisco ACI Software Development Kits

Python SDK: Cobra

Simplified Python SDK: ACI Toolkit

Ruby SDK

PowerShell SDK

Where to Find Automation and Programmability Examples

Developing and Testing Your Code Without an ACI Fabric at Hand

Table of Contents

Cisco DevNet

dCloud

Cisco ACI Simulator

Increasing Operational Efficiency Through Network Automation

Offering Visibility to the Network

Externalizing Network Configuration

Externalizing Switch Port Configuration

Externalizing Security Configuration

Horizontal Automation Integrations

Horizontal Integration Examples Embedded in the Product

Horizontal Integration Example Through External Automation

Automating the Generation of Network Documentation

Enabling Additional Business Models Through Network Automation

Agile Application Deployment and DevOps

Continuous Deployment and Continuous Integration

Linux Containers and Microservices Architectures

Configuration Management Tools

Private Cloud and IaaS

Integration with Cisco Enterprise Cloud Suite

Integration with VMware vRealize Suite

Integration with Microsoft Azure Pack and Azure Stack

Integration with OpenStack

Hybrid Cloud

Platform as a Service

ACI Integration with Apprenda

Mantl and Shipped

Cisco ACI App Center

Summary

Index

Table of Contents