

CISCO *Live!*

Let's go



The bridge to possible

How to Setup an ACI Multi-Site with Single Pod and Multi-Pod

Max Ardica, Distinguished Engineer
Ramses Smeyers, Distinguished Engineer



CISCO *Live!*

BRKDCN-2919

Who Are the Presenters?



Max Ardica
Distinguished Engineer – DC Team
@maxardica



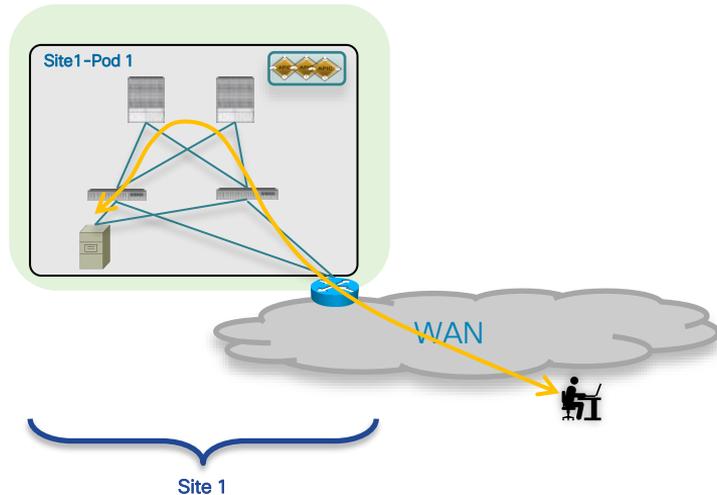
Ramses Smeyers
Distinguished Engineer – CX TAC
@rsmeyers

Agenda

CISCO *Live!*

- What We Want to Achieve Today
- Prerequisites
- Hardware Inspection and Installation
- Installing the First DC Site
- Expanding the Single Pod into a Multi-Pod Fabric
- Introducing Multi-Site and NDO
- Adding the DR Site on NDO
- NDO Additional Functionalities

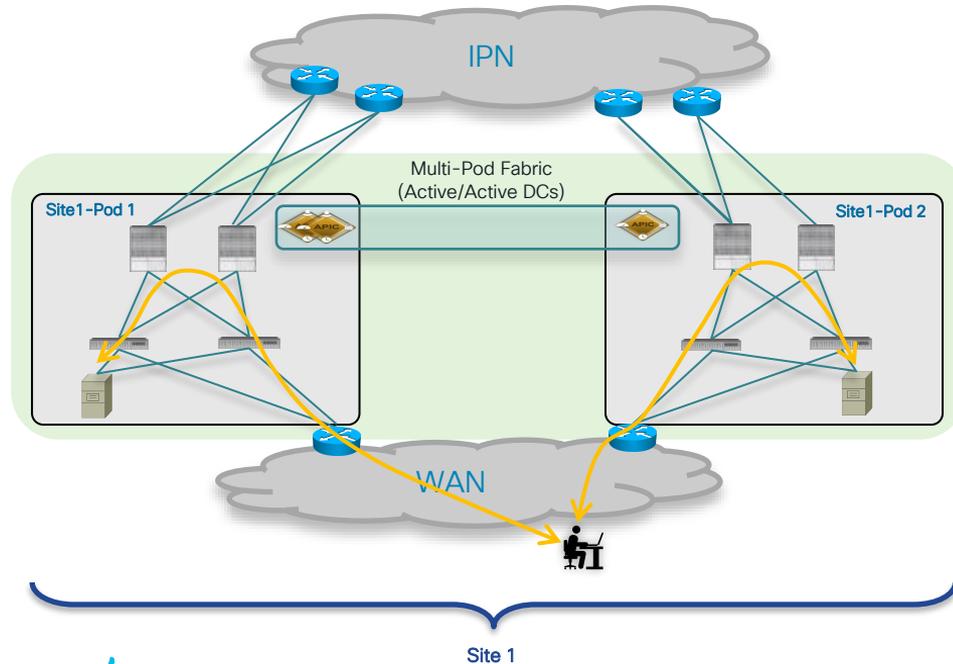
What We Want to Achieve Today



Starting Point

- RamMax Inc. starts its business operations with a single DC location
- Clients can successfully access the application inside the DC

What We Want to Achieve Today



- RamMax's business is very successful and there is a need to grow the size of the DC fabric
- Multi-Pod is the architecture of choice to extend the ACI fabric presence across two DC locations in the same metro area (operated as active/active DCs)

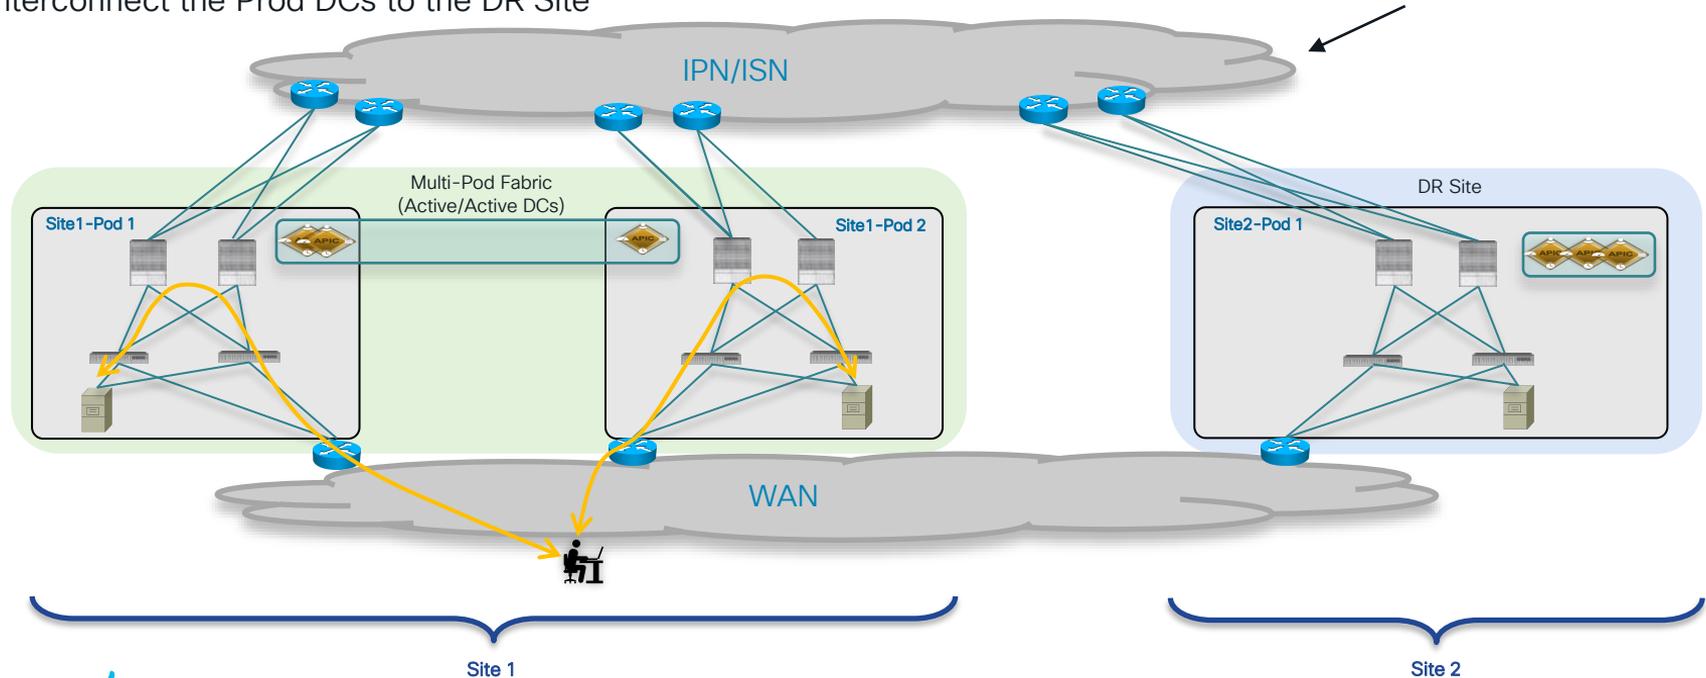
What We Want to Achieve Today

- RamMax decides to build a DR site to ensure business continuance and recovery
- Multi-Site is the architecture used to interconnect the Prod DCs to the DR Site



Nexus Dashboard
Orchestrator

Single external network used for
IPN and ISN



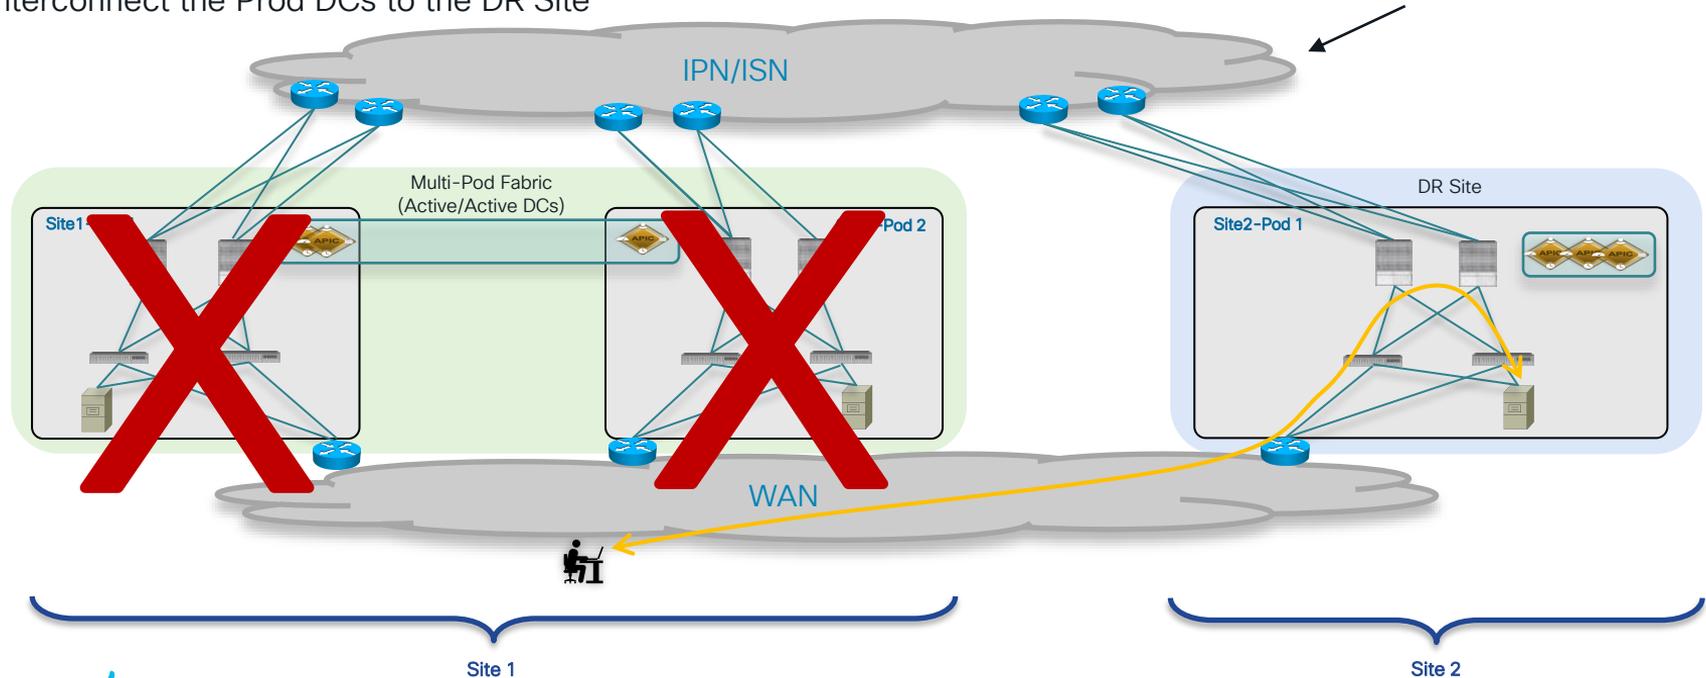
What We Want to Achieve Today

- RamMax decides to build a DR site to ensure business continuance and recovery
- Multi-Site is the architecture used to interconnect the Prod DCs to the DR Site



Nexus Dashboard Orchestrator

Single external network used for IPN and ISN



Prerequisites

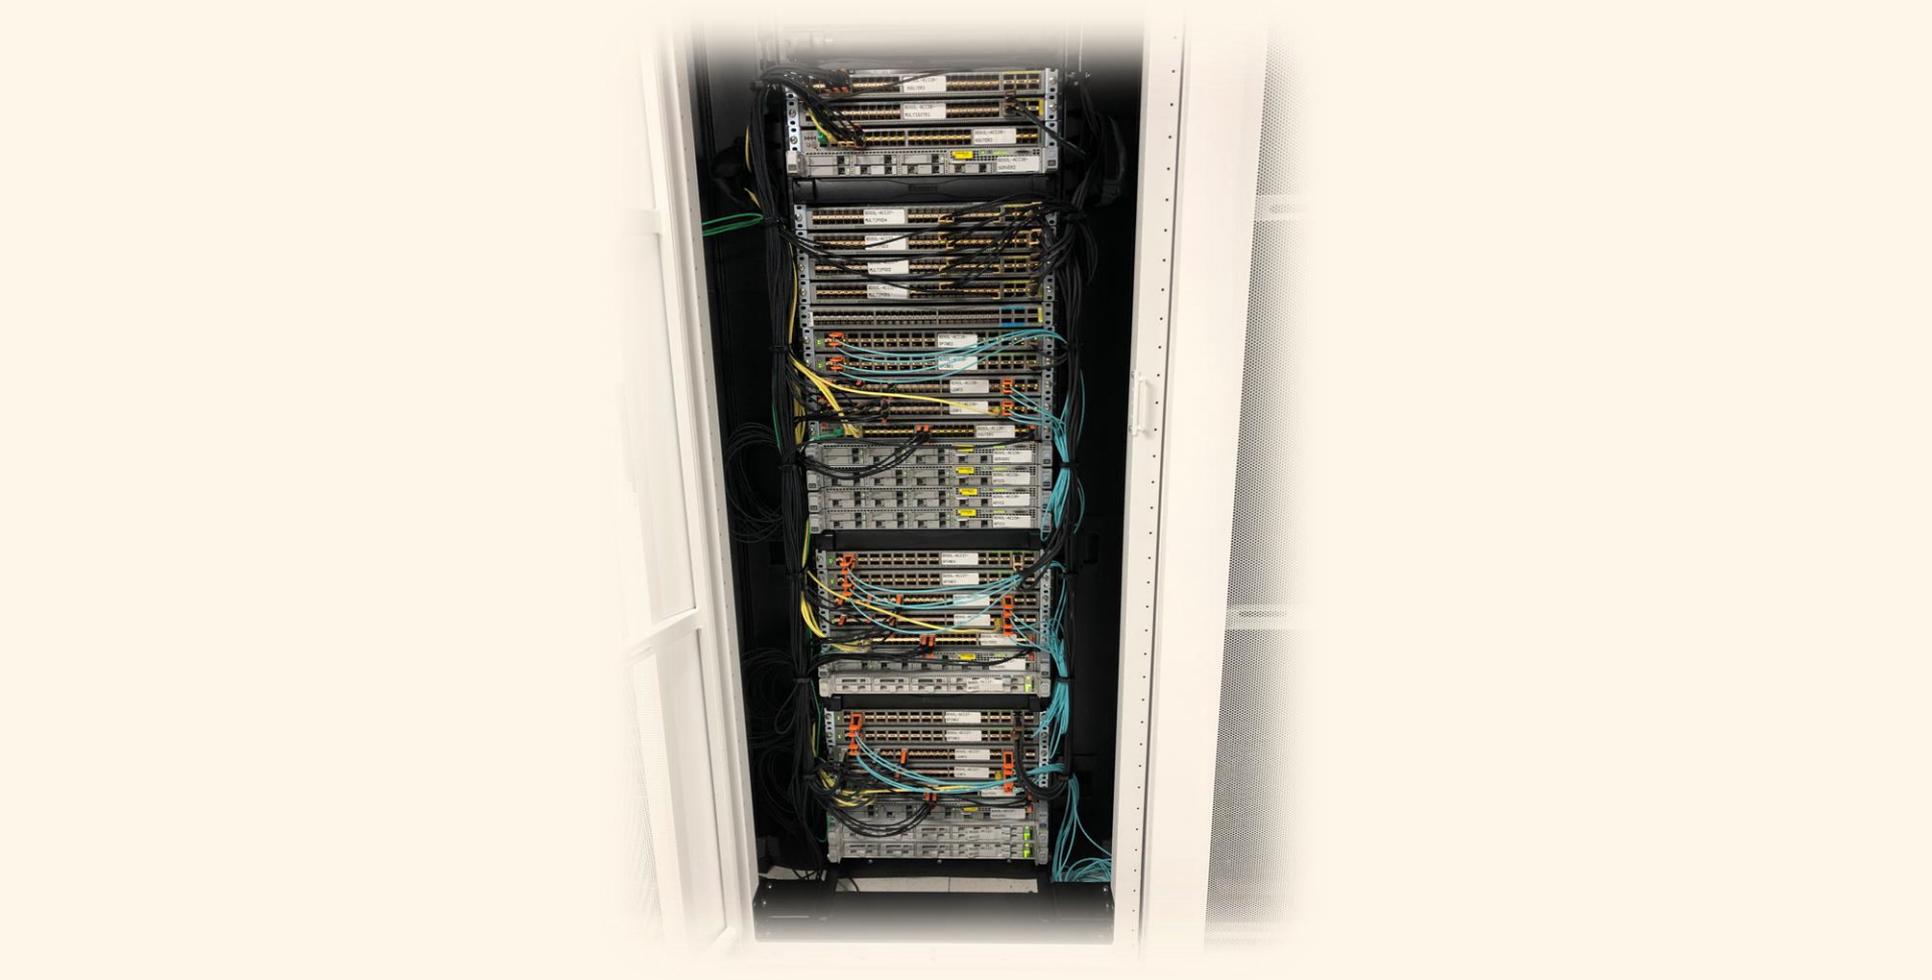


Prerequisites

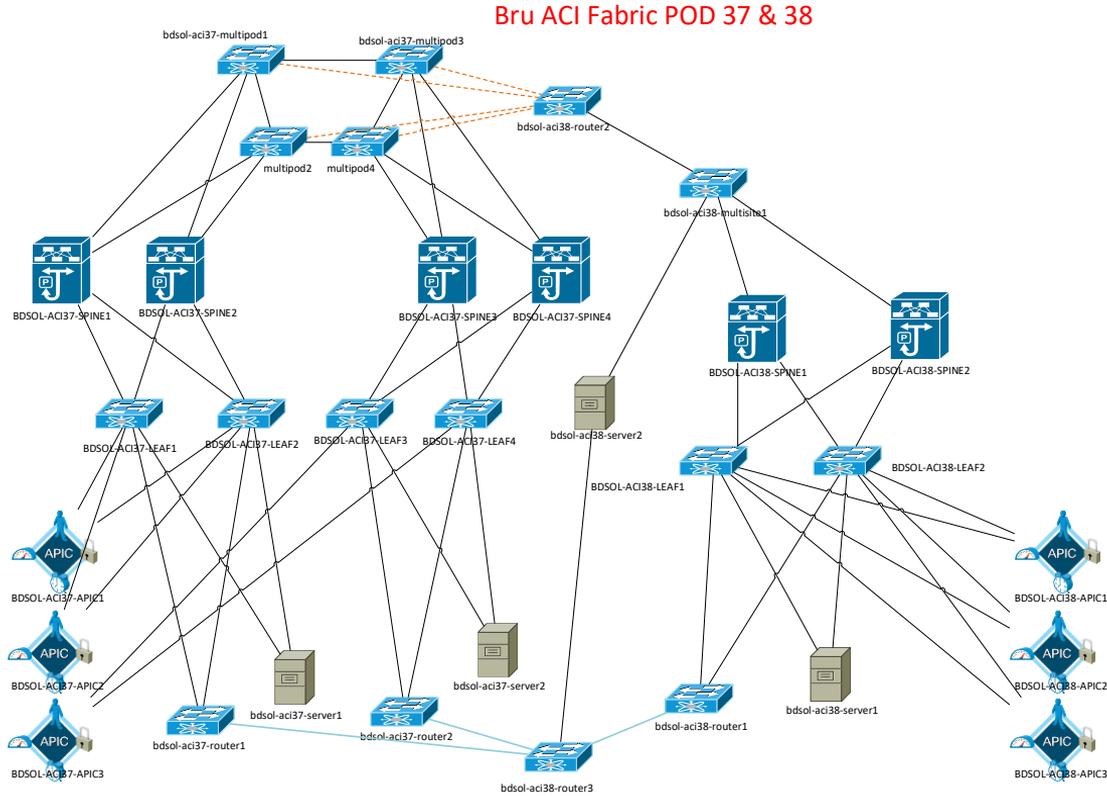
Before starting, you should have:

- For each APIC a routable IP addresses for OOB mgmt and CIMC
- Functional NTP server
- Serial number of all leaf and spine nodes
- Optionally but recommended:
 - 1 IP per leaf and spine for OOB
 - SCP / FTP / HTTP server (software)
 - Console / serial server
- Infrastructure VLAN / VTEP pool
- vCenter IP address and credentials

Hardware Inspection and Installation



Our Setup for Today (Detailed View)



Installing the First DC Site

Site1-Pod1 Configuration

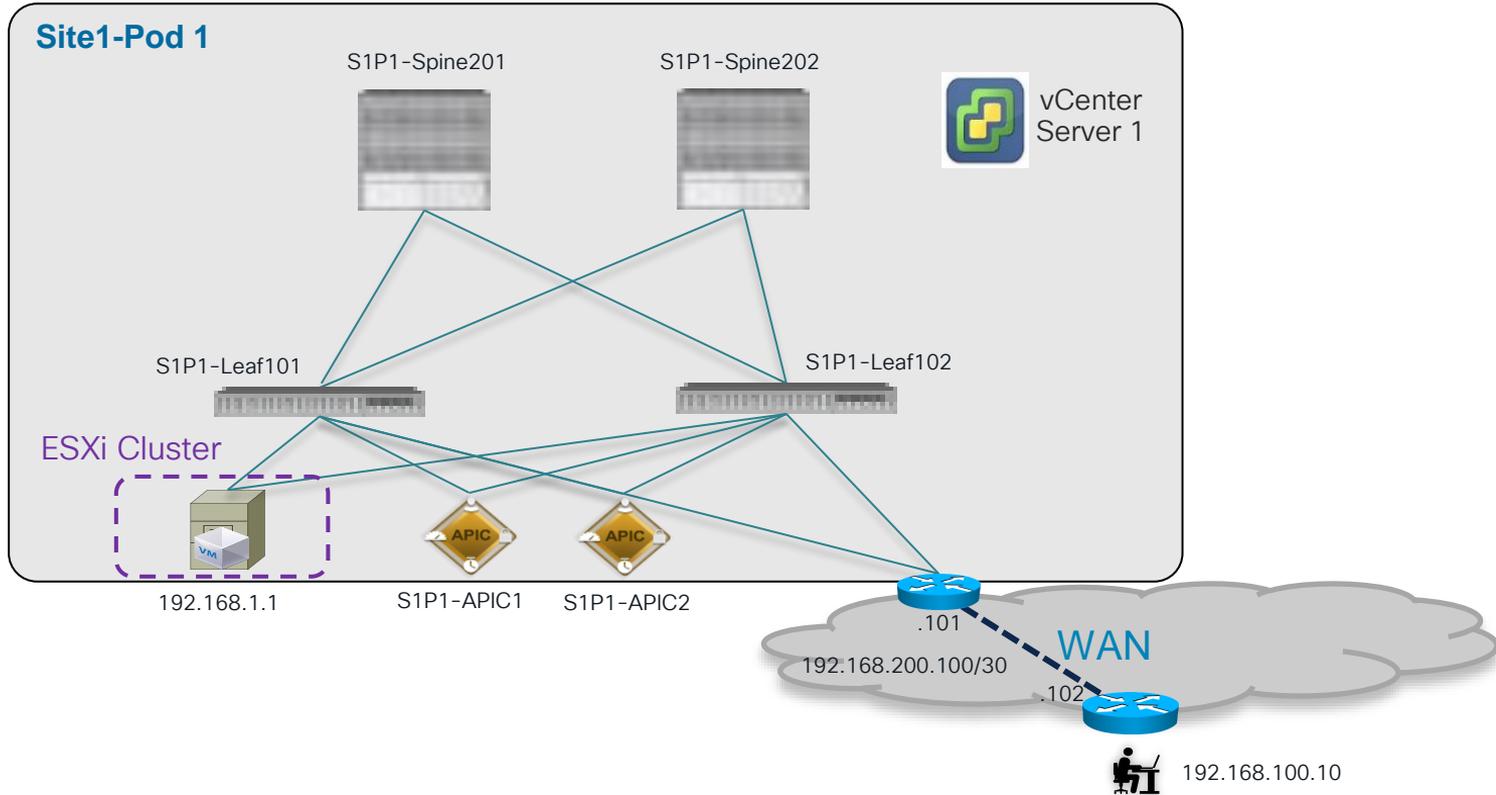
Installing the First DC Site

Site1-Pod1: Initial Fabric Setup (Already Done)

- APIC initial configuration (S1P1-APIC1) [only the 1st one for now]
- 1st leaf discovery
- Spines discovery
- 2nd leaf discovery
- S1P1-APIC2 configuration
- Verification
- OOB mgmt → IPs for leaf and spine nodes

Installing the First DC Site

Site1-Pod1 Fabric



Installing the First DC Site

Parameters for the APIC Initial Setup Script

	S1P1-APIC1	S1P1-APIC2
Fabric name	Fabric1	Fabric1
Fabric ID	1	1
Active controllers	3	3
Pod ID	1	1
Controller ID	1	2
TEP Pool	10.0.0.0/16	10.0.0.0/16
Infra VLAN	3937	3937

Installing the First DC Site

Site1-Pod1: Remaining Configuration Steps to Do

- NTP configuration
- Route Reflector for intra-BGP VPNv4 sessions
- VMM integration
- Tenant configuration with 'Ecommerce' running application
- 'Ecommerce' app connectivity verification
- L3Out creation and external connectivity verification

Demo 1

NTP, VMM, L3Out Configuration and Pod Verification

Expanding the Single Pod into a Multi-Pod Fabric



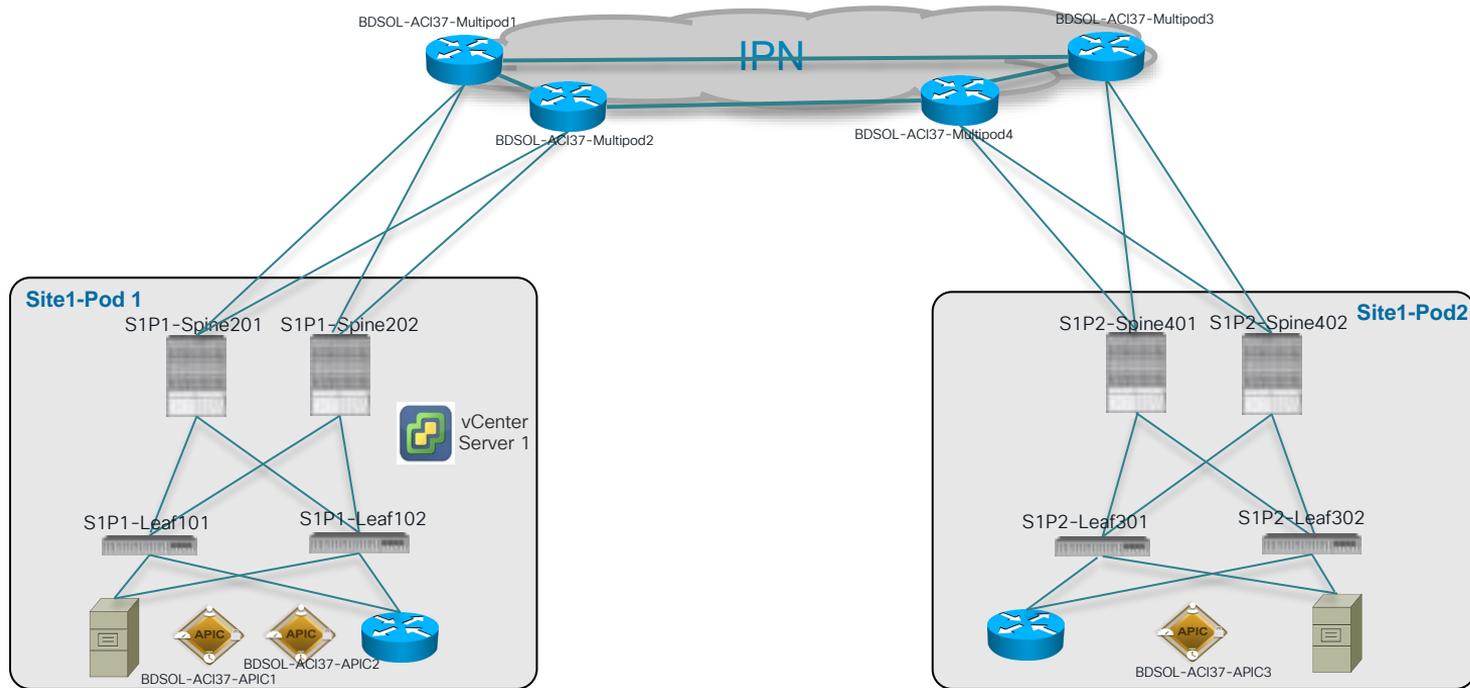
Expanding the Single Pod into a Multi-Pod Fabric

Adding the IPN and Site1-Pod2

- Step 1: verify the Inter-Pod Network (IPN) configuration
- Step 2: create the Multi-Pod fabric using the APIC Wizard
 - Add Site1-Pod1
 - Add Site1-Pod2
 - Discovery of Pod2's leaf and spines nodes
- Step 3: S1P2-APIC3 in Pod2 joins the APIC cluster
- Step 4: extend 'Ecommerce' Tenant to Pod2 (L3Out, ESXi host, access policies)
- Verification Steps:
 - Verify that the existing tenant configuration is extended into the Multi-Pod fabric
 - Verify East-West and North-South connectivity

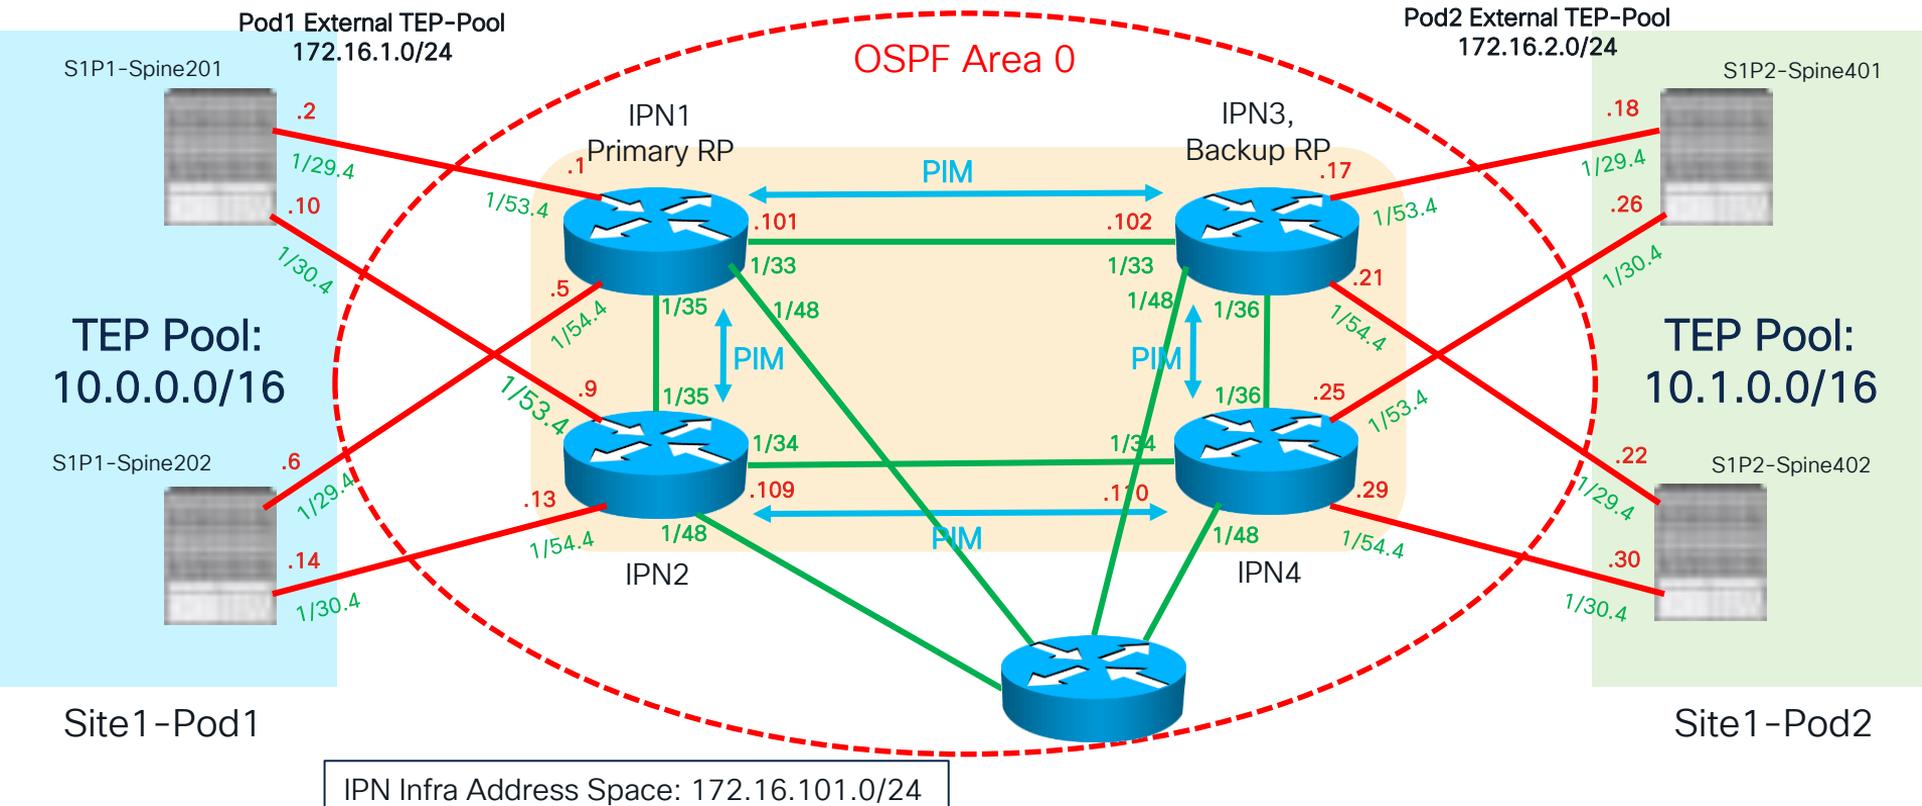
Expanding the Single Pod into a Multi-Pod Fabric

Adding the IPN and Site1-Pod2



Expanding the Single Pod into a Multi-Pod Fabric

Step1: Verify the Inter-Pod Network (IPN) Configuration



Demo 2

Verify the Inter-Pod Network (IPN) Configuration

Demo 3

Create the Multi-Pod Fabric Using the APIC Wizard

Expanding the Single Pod into a Multi-Pod Fabric

Step2: Create the Multi-Pod Fabric Using the APIC Wizard and import Pod2 Spine and Leaf Nodes

Nodes automatically discovered in Site1-Pod2 that need to be added to the APIC fabric membership table

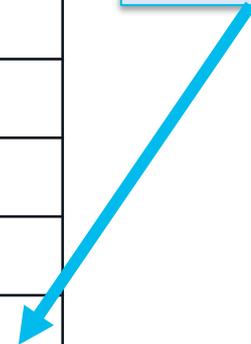
Node ID	Pod ID	Name	S/N
301	2	S1P2-Leaf301	FDO224702ET
302	2	S1P2-Leaf302	FDO223007J4
401	2	S1P2-Spine401	FDO22472FCV
402	2	S1P2-Spine402	FDO22391NP2

Expanding the Single Pod into a Multi-Pod Fabric

Step 3: S1P2-APIC3 in Pod2 Joins the APIC Cluster

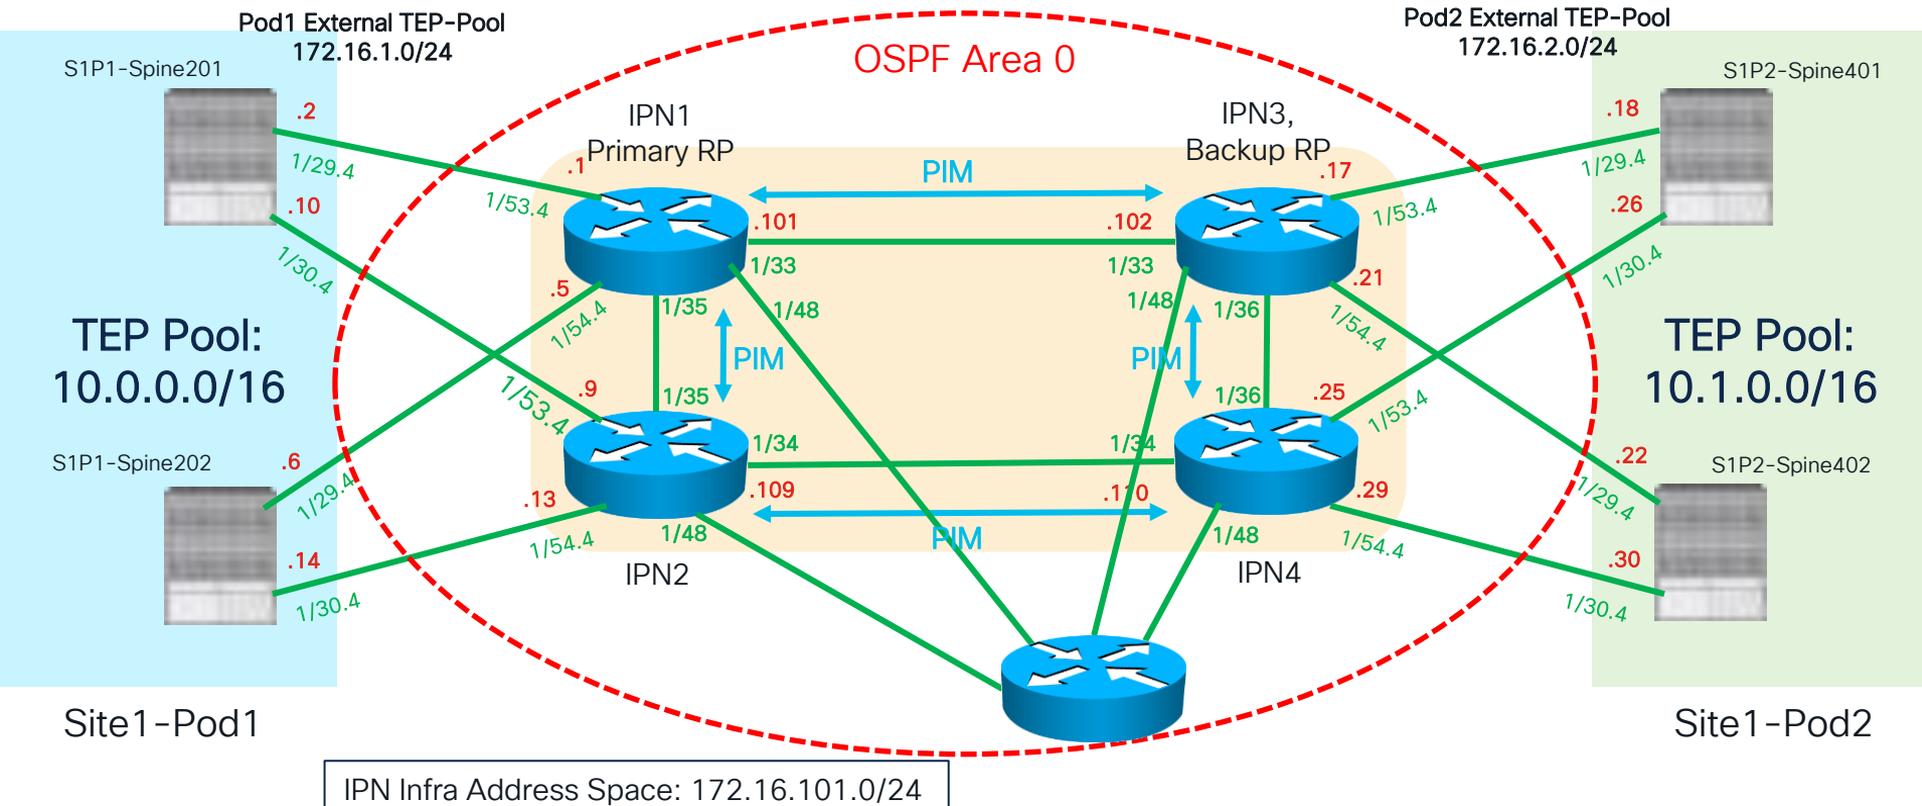
	S1P1-APIC1	S1P1-APIC2	S1P2-APIC3
Fabric name	Fabric1	Fabric1	Fabric1
Fabric ID	1	1	1
Active controllers	3	3	3
Pod ID	1	1	2
Controller ID	1	2	3
TEP Pool	10.0.0.0/16	10.0.0.0/16	10.0.0.0/16
Infra VLAN	3937	3937	3937

Pod2 uses TEP Pool 10.1.0.0/16 but S1P2-APIC3 resides in TEP Pool of Pod1



Expanding the Single Pod into a Multi-Pod Fabric

Full OSPF Adjacencies between Spines and IPN in both Pods

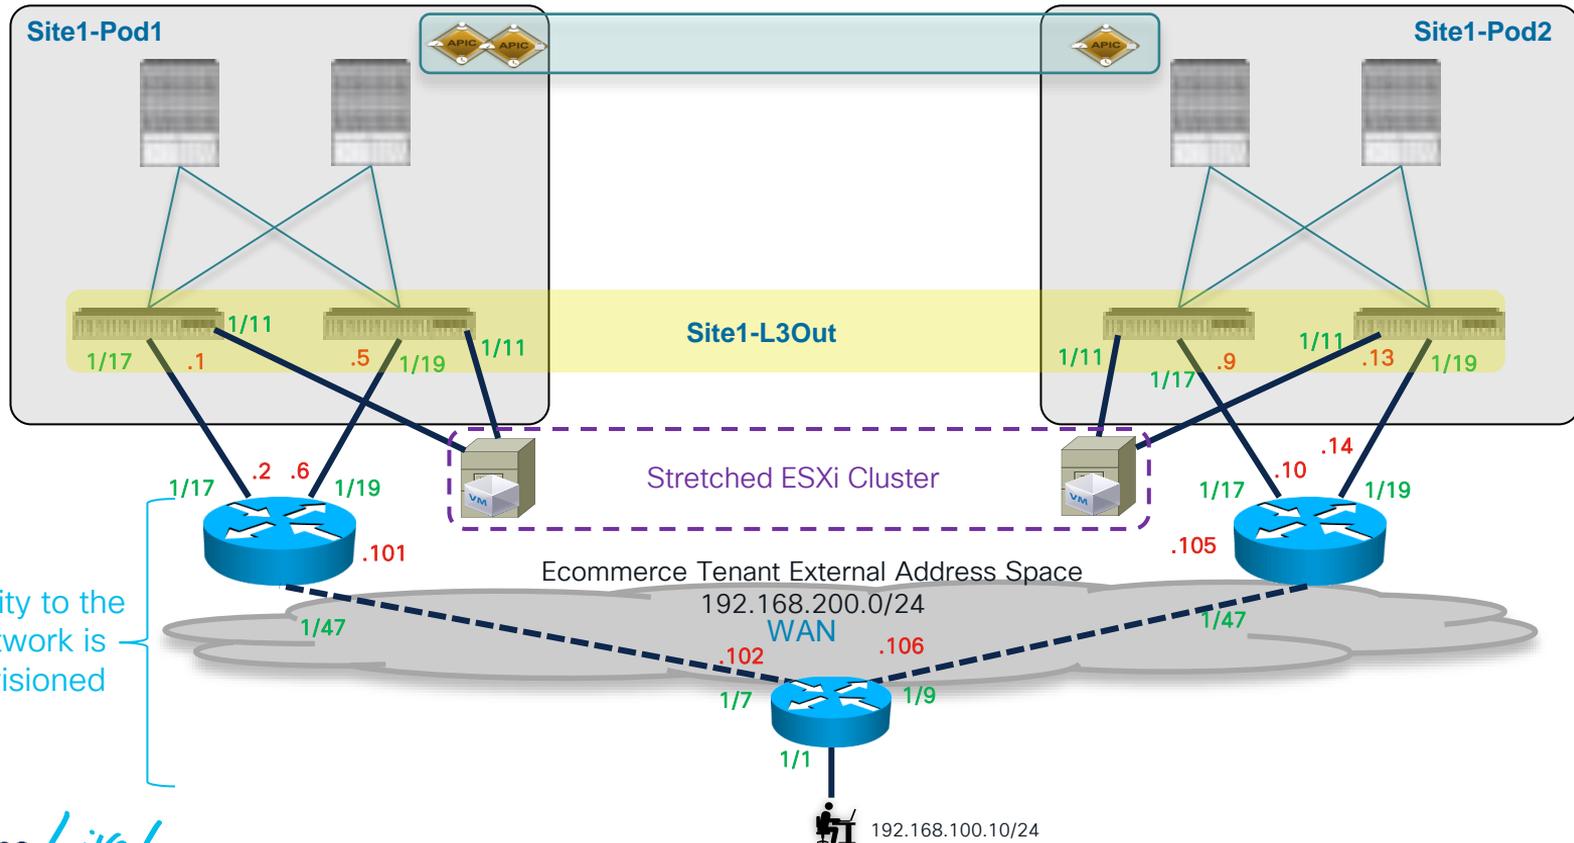


Demo 4

S1P2-APIC3 in Pod2 Joins the APIC Cluster

Expanding the Single Pod into a Multi-Pod Fabric

Step 4: Extend 'Ecommerce/Tenant' to Pod2 (L3Out, ESXi Host to VDS, etc.)



Connectivity to the WAN network is pre-provisioned

Demo 5

Extend 'Ecommerce' tenant
configuration to Pod2

Introducing NDO and Multi-Site



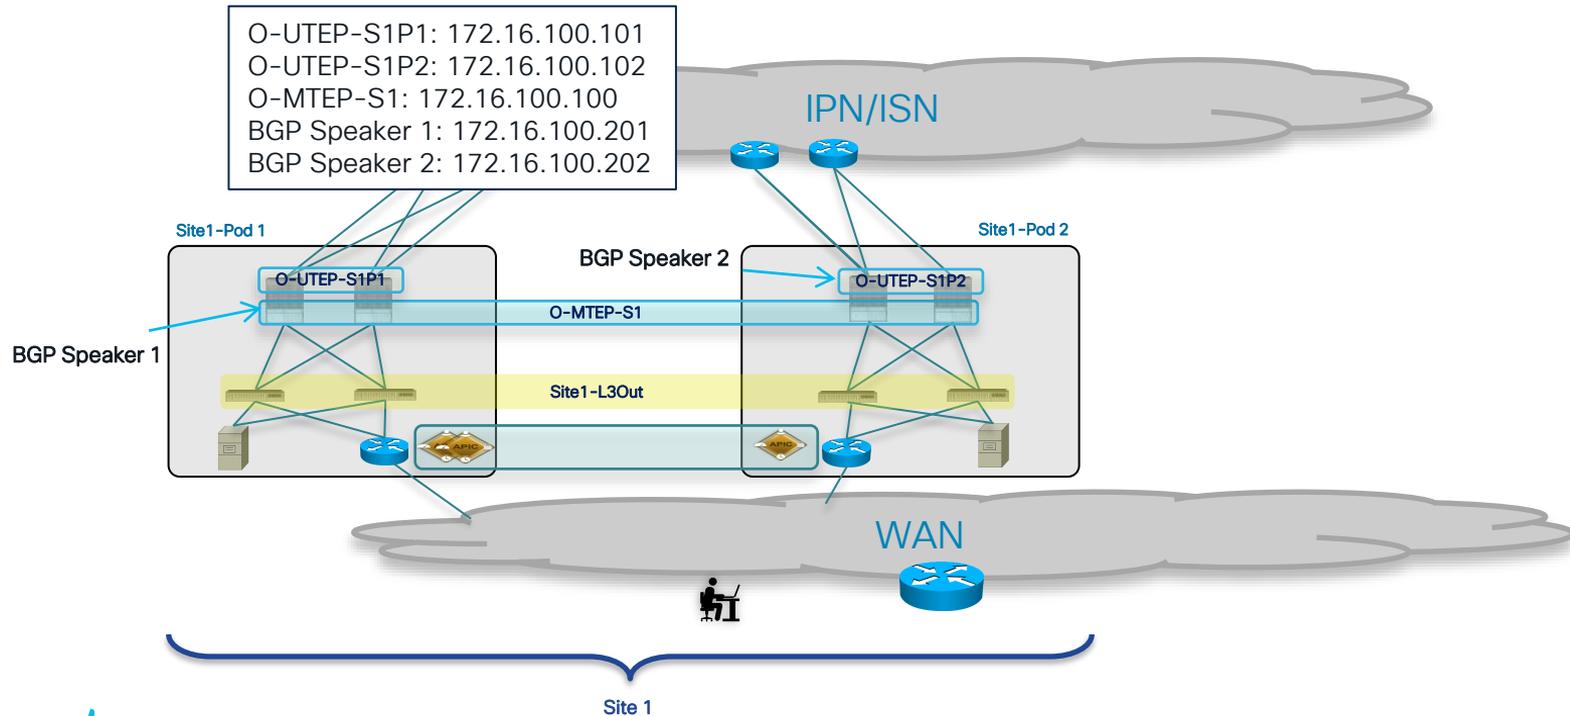
Introducing NDO and Multi-Site

Configuration Steps

- Initial setup of ND/NDO
- Adding the Multi-Pod fabric as first site on ND/NDO
 - Onboarding the Multi-Pod Fabric on ND
 - Setting the Multi-Pod Fabric as “Managed” on NDO
- Importing existing ‘Ecommerce’ tenant configuration on NDO

Adding the Multi-Pod Fabric on NDO

Assign Multi-Site TEP Addresses and BGP EVPN Router-IDs



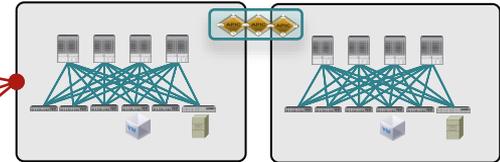
Schema Design (Initial State)

One Template for Site 1 local objects, plus 'Stretched' Templates

Schema Ecommerce



Site 1





Demo 6

Simplified Tenant Management through NDO

Adding the DR Site on NDO



Adding the DR Site on NDO

DR Site Preparation

- Through automation, create the necessary configuration for the DR site
 - Access policies
 - VMM domain integration

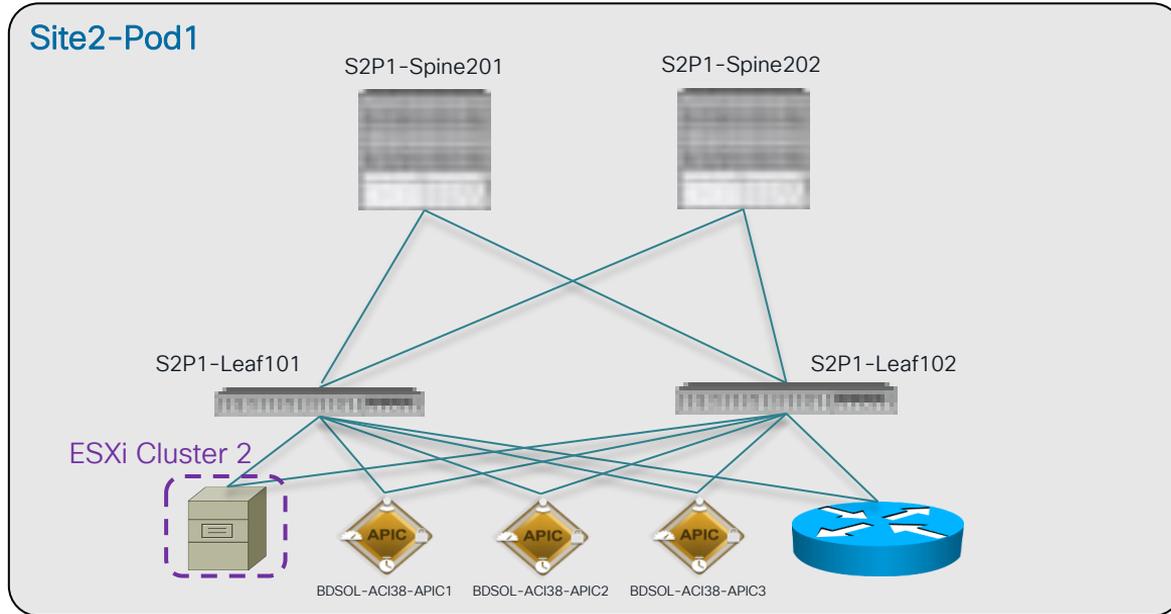
Adding the DR Site on NDO

Configuration Steps

- Adding the DR fabric as a second site on ND/NDO (onboarding it on ND, managing it on NDO, assigning Multi-Site TEP addresses, BGP EVPN Router-IDs and External TEP pool, establishing OPSF connectivity between spines and IPN routers)
- Verifying IPN connectivity
- Extending the tenant 'Ecommerce' to the DR site
- Create the Site2 local L3Out directly on NDO in the "L3OUt Template" (new NDO 4.1(1) feature)
- Extending the existing 'Ecommerce' **tenant configuration** to the DR site
- Verify East-West and North-South connectivity

Adding the DR Site on NDO

Site2-Pod1 Fabric

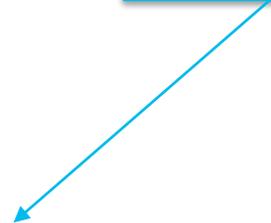


Adding the DR Site on NDO

Parameters for the APIC Initial Setup Script (Already Done)

	S2P1-APIC1	S2P1-APIC2	S2P1-APIC3
Fabric name	Fabric2	Fabric2	Fabric2
Fabric ID	1	1	1
Active controllers	3	3	3
Pod ID	1	1	1
Controller ID	1	2	3
TEP Pool	10.2.0.0/16	10.2.0.0/16	10.2.0.0/16
Infra VLAN	3937	3937	3937

Recommended
not to use
overlapping TEP
Pools with
existing sites

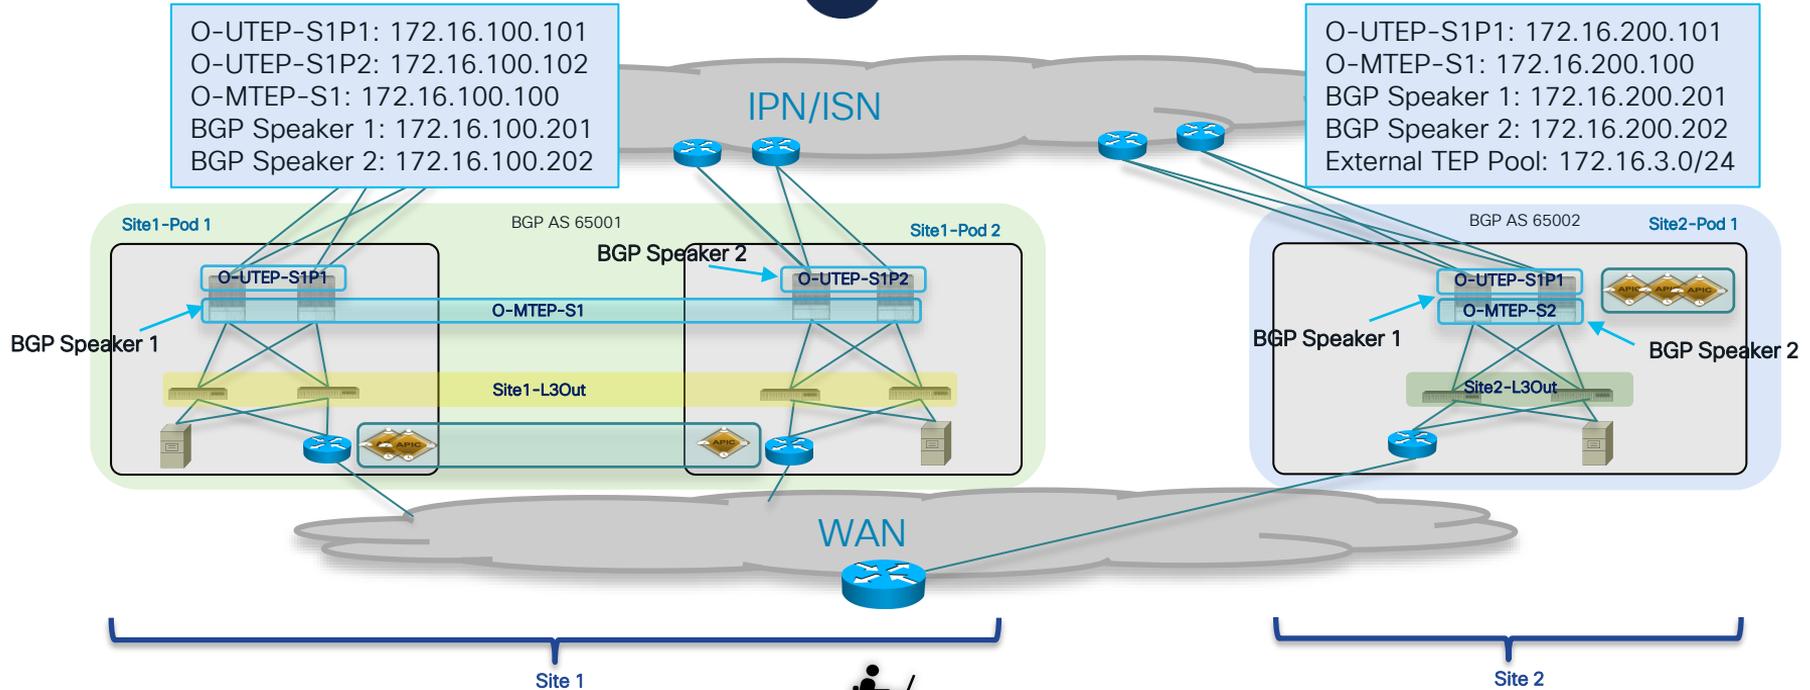


Adding the DR Site on NDO

Assign Routable TEP Addresses and BGP EVPN Router-IDs

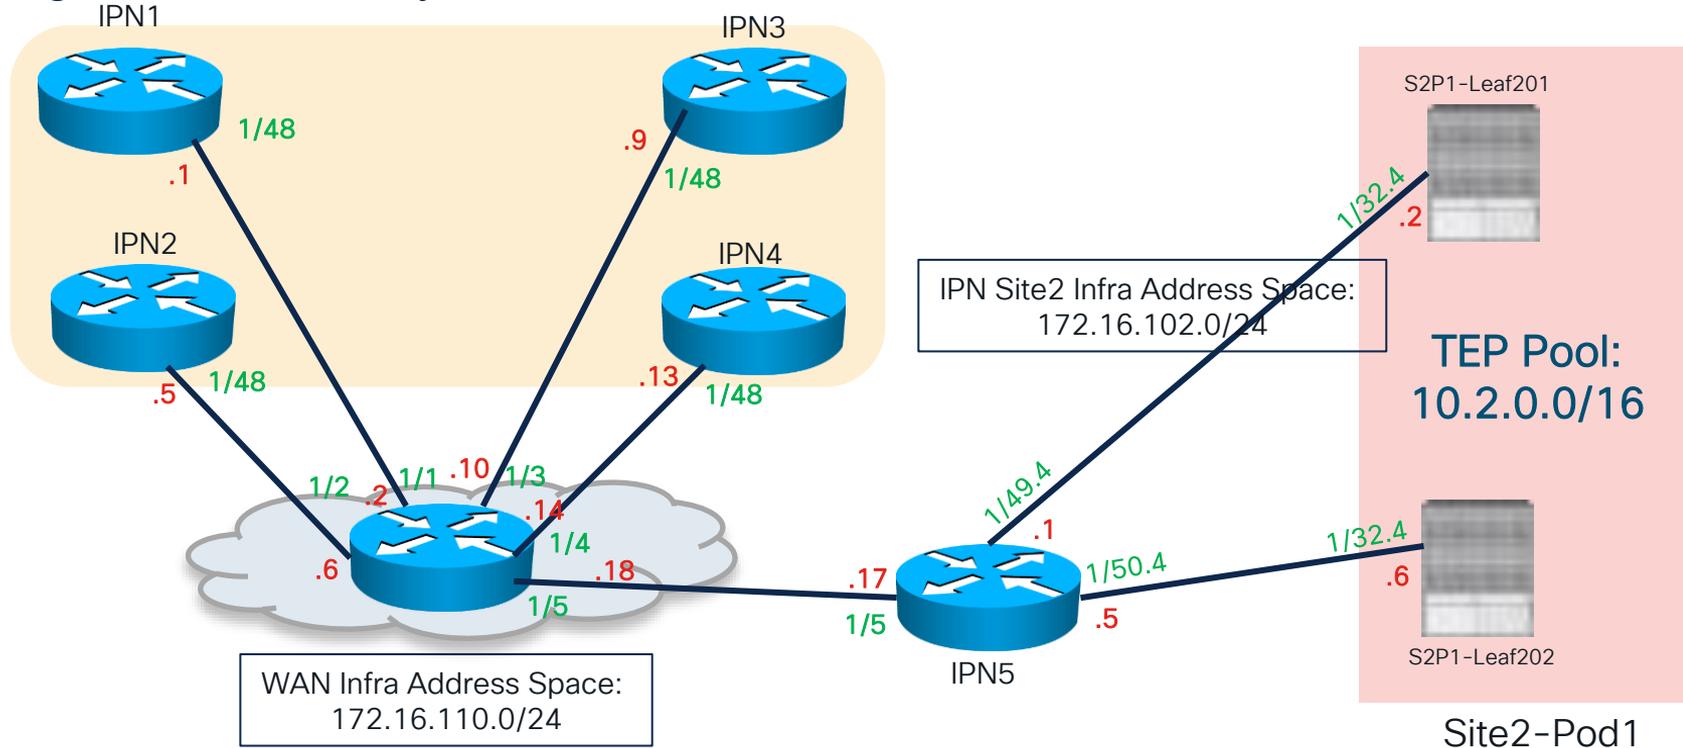


Nexus Dashboard
Orchestrator



Adding the DR Site on NDO

Verifying IPN connectivity

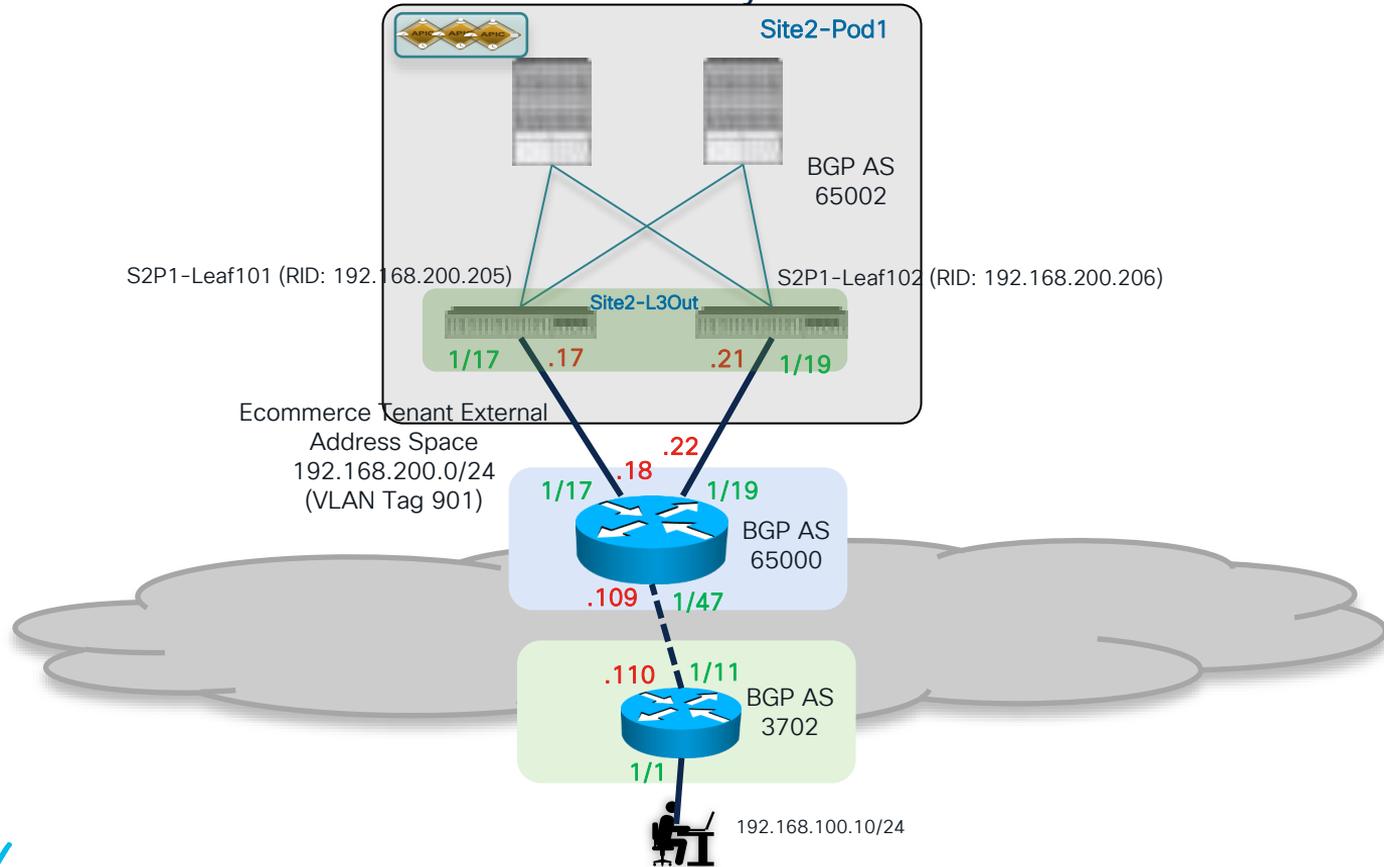


Demo 7

Adding the DR Site on NDO

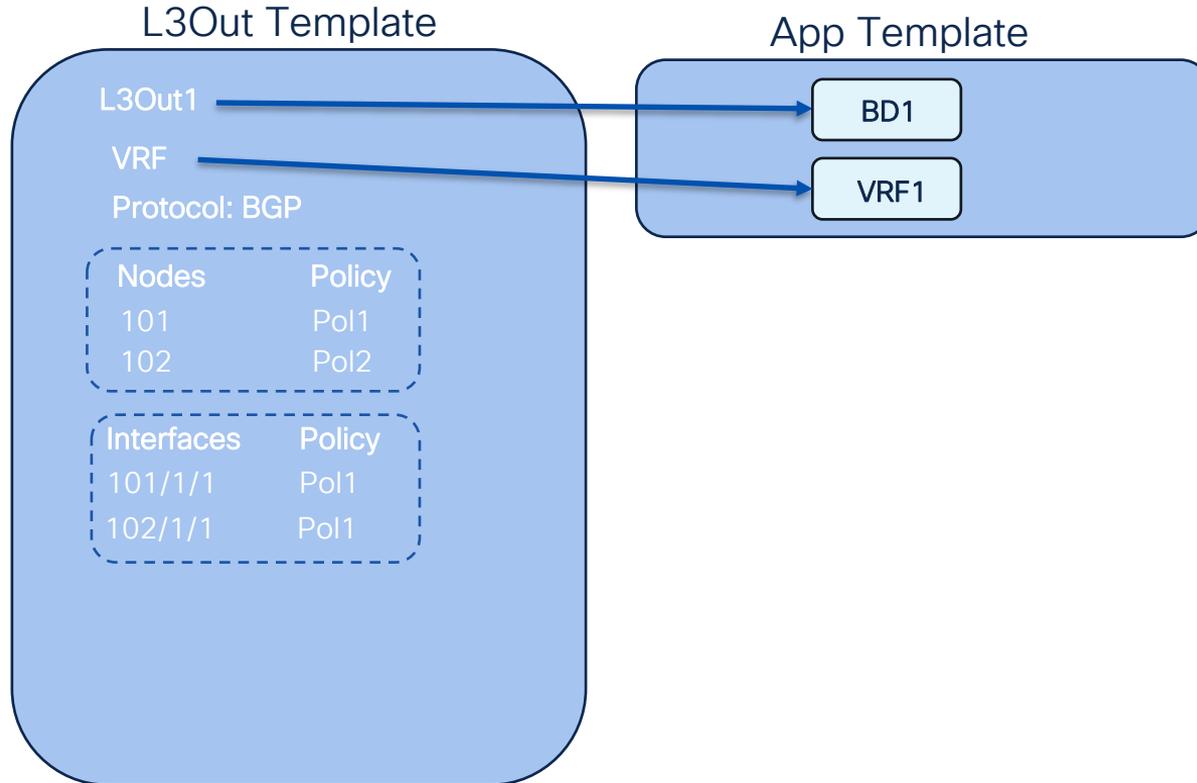
Adding the DR Site on NDO

Create a Local L3Out in the DR Site Directly on NDO



Create a Local L3Out Directly on NDO

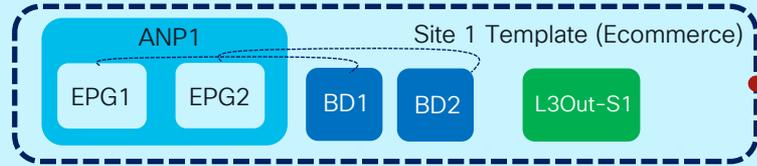
Cross-Referencing Different Types of Templates



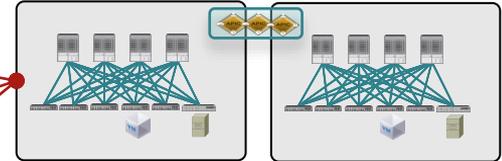
Schema Design (Today)

One Template for Site 1 local objects, plus 'Stretched' Templates

Schema Ecommerce



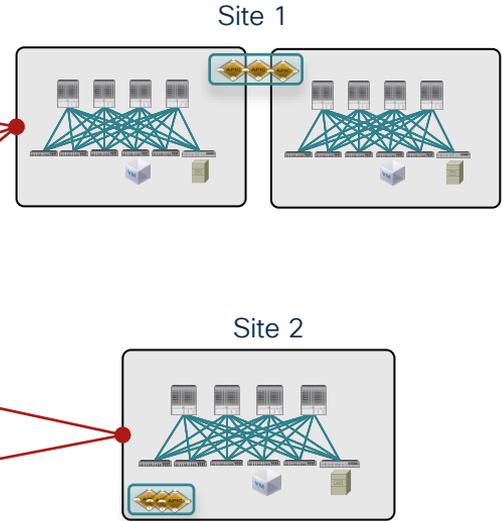
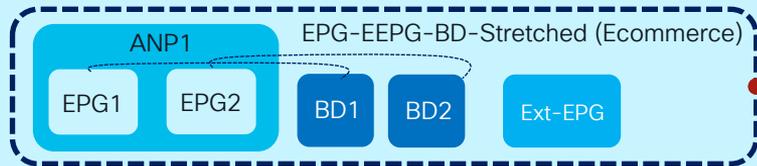
Site 1



Schema Design (Final State)

Migrating EPGs/BDs to the Stretched Template

Schema Ecommerce



NDO Additional Functionalities



NDO Additional Functionalities

- End host connectivity verification
- Host route advertisement (inbound traffic optimization)



Demo 8

NDO Additional Functionalities

Where to Go for More Information



- ✓ ACI Multi-Pod White Paper
<http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html?cachemode=refresh>
- ✓ ACI Multi-Pod Configuration Paper
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739714.html>
- ✓ ACI Multi-Pod and Service Node Integration White Paper
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739571.html>
- ✓ ACI Multi-Site White Paper
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html>
- ✓ Cisco Multi-Site Deployment Guide for ACI Fabrics
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html>
- ✓ ACI Multi-Site and Service Node Integration White Paper
<https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-743107.html>



The bridge to possible

Thank you

CISCO *Live!*



The Cisco Live! logo features the word "CISCO" in a bold, black, sans-serif font, followed by "Live!" in a black, cursive script font. The background of the entire image is a vibrant, multi-colored abstract pattern of overlapping, wavy bands in shades of red, orange, yellow, green, and blue, creating a sense of motion and energy.

CISCO *Live!*

Let's go