White Paper

# Cisco Nexus 9000 Series Switches In-Place Migration from NX-OS Mode to ACI Mode

## Executive Summary

Customers of all sizes and in all industries that operate networks in high-availability data centers are facing similar scaling challenges associated with the dynamic nature of hybrid bare-metal and virtualized operating system environments. Among these challenges are VLAN provisioning, Spanning Tree Protocol use, and data center–wide visibility for east-west and north-south traffic flows.

Cisco Nexus® 9000 Series Switches have two modes of operation: the standalone NX-OS mode using Cisco® NX-OS Software, and the high-scale ACI mode using the Cisco Application Centric Infrastructure (Cisco ACI™) solution. This document describes the process of migrating from an access-layer network environment that uses NX-OS with virtual port channels (vPCs) to ACI mode using an in-place deployment model.

The goal here is to demonstrate a representative architecture and design and to provide step-by-step procedures for implementing this design that are replicable. Various permutations of server connectivity using virtualization platforms and bare-metal operating systems were tested to determine the availability ramifications of the migration process, with observations and considerations for the network, infrastructure, virtualization, and server environments reported. This document emphasizes the best architecture to allow clients to plan, implement, and benefit from the transition with little service interruption. It presents the benefits of using the Cisco Nexus 9000 Series for data center networking and notes the relevant findings.

The main audience for this document consists of analysts, consultants, managers, and interested network engineers.

## Background

### Server Uptime with Host-Based Virtual Port Channels

Application uptime for business-critical services is essential in today's IT systems. Achieving this uptime requires highly available, redundant topologies at multiple levels of the infrastructure stack: network, security, computing, and storage. Cisco enables computing uptime through host-based vPCs by allowing the industry-standard Link Aggregation Control Protocol (LACP) to bond together physical links from the server to a pair of access switches in the data center. The benefit of this approach is that during any individual access switch outage, planned or unplanned, applications running on servers attached through host-based vPC will continue to function normally.

### Virtual Port Channels for Network High Availability

Because of the shortcomings of Spanning Tree Protocol and its variants, Spanning-Tree Protocol avoidance techniques that support loop-free network topologies have become prevalent in data-link-layer networking. Network vPCs, similar to host-based vPCs, are commonly used on Cisco Nexus switches to achieve network high availability between the aggregation and access layers in the data center without the use of Spanning Tree Protocol.

The vPC uses a pair of Cisco Nexus aggregation switches to create a logical downstream LACP port channel from the aggregation switches to the access switches while maintaining a data-link-layer connection between the aggregation switches and using a high-availability forwarding protocol (such as Hot Standby Router Protocol [HSRP] or Virtual Router Redundancy Protocol [VRRP]) to allow traffic to flow over either path. vPCs are always deployed in pairs, and a single pair of aggregation switches can support multiple vPCs, each to an individual access switch. This approach is congruent with traditional data center network design and is a scalable model: each pair of aggregation switches can be interconnected through a routed core network if greater scale is required. This model creates a traditional three-tier data center network.

### Spine-and-Leaf Topologies

Spine-and-leaf topologies have become more common because of the shift in application traffic patterns in the data center from predominantly north-south flows to more complex device-to-device, or east-west, flows. Spine-and-leaf topologies support greater scalability at the aggregation layer. This scalability is achieved by requiring all aggregation switches to be connected to all access switches without the need for interconnection at the aggregation layer. Instead, traffic passes from any access switch through any aggregation switch for receipt by another access switch using Equal-Cost Multipath (ECMP) routing, facilitating the use of all paths concurrently. Cisco ACI fabric uses the Intermediate system–to–Intermediate System (IS-IS) routing protocol to forward Virtual Extensible LAN (VXLAN) encapsulated frames from leaf to leaf through the spine switches.

### Cisco Nexus 9000 Series Switches in NX-OS Mode

Cisco Nexus 9000 Series Switches can operate either in a traditional data center model with vPC or in a spine-and-leaf model using VXLAN as the tunneling protocol to achieve great scale and redundancy without the need for a core network. They can function effectively in either model in NX-OS mode, much like earlier Cisco Nexus Family switches (such as Cisco Nexus 3000, 5000, 6000, 7000 Series Switches). The main factor in determining which data center model to use with Cisco Nexus 9000 Series Switches is the pattern of traffic flows.

### Traffic Flow Patterns

In traditional data center networks, traffic is primarily transferred from servers in the data center to clients outside the data center, and returned to the servers from the clients. This pattern is referred to as north-south traffic.

In spine-and-leaf topologies, traffic is primarily passed between servers within the data center, with servers collaborating closely to collect, process, analyze, and store large amounts of data for either network storage or network throughput. This approach is particularly pertinent in cloud, enterprise resource planning (ERP), and big data applications. This pattern is referred to as east-west traffic.

### Scale and Manageability

When north-south or east-west traffic requirements increase significantly within a data center, scalability requirements also increase, as does network complexity to support such scale. Although a well-designed architecture can reduce complexity significantly, network models that use decentralized device management create a human-intensive operational burden. In such models, much of the intelligence in network management needs to be implemented manually by network administrators, often in mission-critical or high-business value environments. This approach creates business risk at scale, with an exponentially increasing data networking load placed on a linearly increasing or static human workforce.

Traditional switched networks are limited in their capability to support high-bandwidth, highly scalable traffic throughput predominantly because they are difficult to manage and optimize at scale.

### Cisco Application Centric Infrastructure with Cisco Nexus 9000 Series Switches

Cisco ACI is a policy-based network automation engine that uses the Cisco Nexus 9000 Series to mitigate scalability and manageability challenges. It addresses the challenges of limited human workforce scale and an exponentially increasing amount of data to allow organizations to effectively empower their human workforce. It provides businesses with an intelligent platform that can implement whatever the network administrator intends for the network using a declarative network policy automation model.

### Planning for the Future

Organizations plan for business growth, and to do so they must also plan for data growth. Cisco Nexus 9000 Series Switches in the data center are crucial elements in providing high-performance network services that meet today's needs and tomorrow's aspirations. Many organizations are choosing the Cisco Nexus 9000 Series operating in NX-OS mode to support the present requirements of their north-south traffic flows and build on the honed NX-OS administration skills of their workforce to meet today's needs. They are confident in the knowledge that as their business grows, and east-west traffic increases, and their data center network scales exponentially, they can easily convert their Cisco Nexus 9000 Series Switches from NX-OS mode to ACI mode. They need only add Cisco ACI policy controllers (called Cisco Application Infrastructure Policy Controllers [APICs]), border leaf switches, and spine proxy switches, with little need for business downtime or excessive future capital expenditures.

Note that this document provides a guide only to what **can** be done; it does not recommend what **should** be done. Any organization intending to move through a full infrastructure lifecycle incorporating Cisco ACI is likely to reduce its network complexity by implementing ACI mode from the start rather than after first implementing NX-OS.

## Architecture

### Initial Architecture

Figure 1 shows the initial architecture for a recommended implementation of NX-OS mode with vPC from the aggregation layer to the access layer and host vPC.

**Figure 1.** Initial NX-OS Mode Architecture

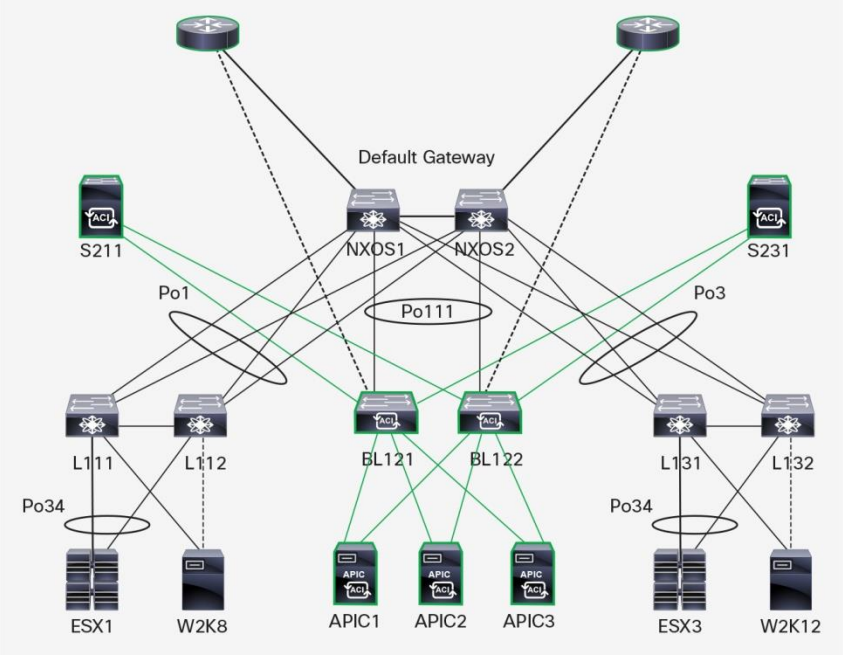

This architecture uses a preconfigured pair of border leaf switches, three APICs, and a minimum of two Cisco ACI spine switches. By adding this equipment to the preexisting pair of aggregation switches and pods of leaf switches, organizations can transparently migrate to ACI mode with little need for additional hardware (other than the new border leaf switches, the APICs, and the spine switches) and with a reduced outage period for migration activities.

This document focuses on the Cisco ACI topology. A similar procedure could also be used to migrate from a VXLAN topology to Cisco ACI.

The uplink interfaces used in this topology to interconnect the Cisco ACI border leaf switches and the aggregation switches that will be migrated to ACI mode must be front-facing ports rather than uplink ports. In most models of leaf switches (except the Cisco Nexus 9332PQ Switch), they are 10-Gbps ports. For this reason, multiple links can be bundled together into an EtherChannel when connecting border leaf switches and aggregation switches, to avoid a bottleneck in throughput during the migration process.

In a strict technical sense, migration can be performed by upgrading an in-place pair of access switches to make them the initial border leaf switches using a modified version of the process described in this document. However, this approach would incur more downtime for vPC setup between the aggregation switches and the new border leaf switches because of LACP negotiation with the initial vPC (this approach is discussed further in the "Conclusion" section of this document). This approach may be suitable for midsize enterprises, but it is less likely to be suited for larger enterprises, which are more likely to be expanding the access layer and which have a higher associated cost of downtime.

### Target Architecture

Figure 2 shows the target architecture for a recommended implementation of an ACI mode deployment.

**Figure 2.** Target ACI Mode Architecture

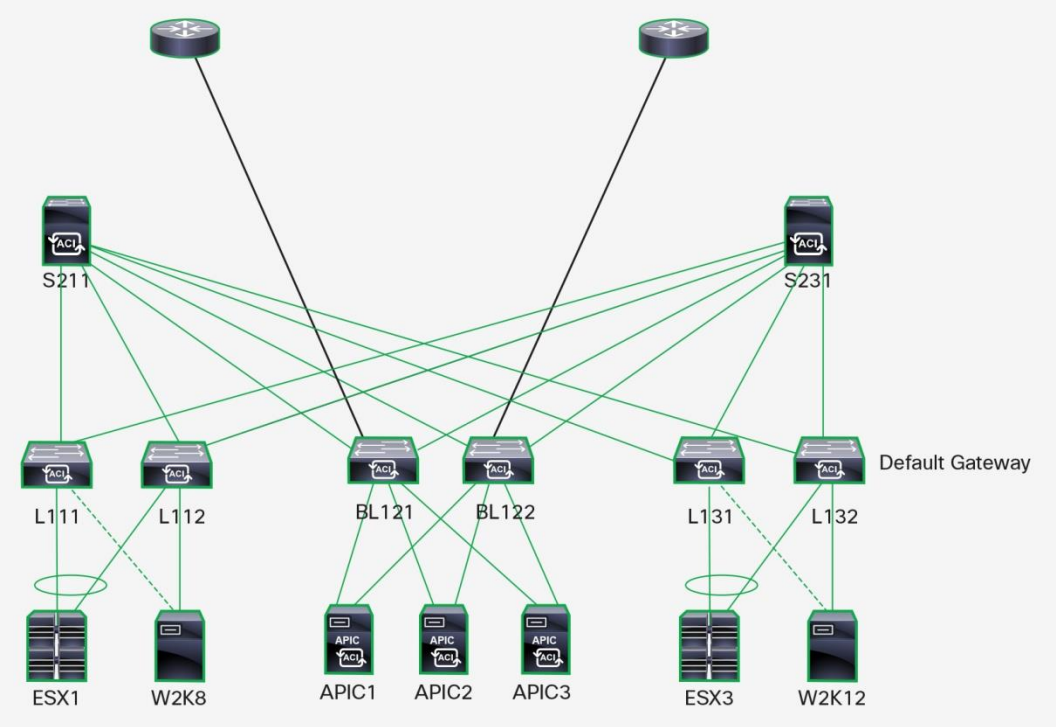

This model uses a leaf-and-spine topology with a minimum of two spine switches. This topology is horizontally scalable as throughput requirements increase.

Leaf switches are divided into pairs to create switch pods, with one leaf in each pod being primary (denoted as xx1) and one leaf being secondary (denoted xx2) for purposes of both the migration process and ongoing availability during switch maintenance windows. Each VMware ESX server is connected to a pod through a host vPC to help ensure link redundancy, and each bare-metal Microsoft Windows server uses active-backup network interface card (NIC) teaming.

A pair of leaf switches is used as border leaf switches for external Layer 2 and 3 connectivity and for APIC management connectivity. In smaller environments, these switches may also connect hosts. However, because border leaf switches may use a large portion of their ternary content-addressable memory (TCAM) for routing entries, you generally should not connect virtualized hosts with large numbers of virtual machines to the border leaf switches.

Additionally, although the NX-OS aggregation switches (either Cisco Nexus 9300 or 9500 platform switches) are shown as removed at the end of this procedure, in practice they would be repurposed either as leaf switches (Cisco Nexus 9300 platform switches) or as Cisco ACI spine switches (Cisco Nexus 9500 platform switches with Cisco Nexus 9700 platform line cards).

## Equipment List

The lab setup used for the example in this document was composed of Cisco Nexus 9300 platform switches, Cisco Unified Computing System™ (Cisco UCS®) servers, and Microsoft and VMware operating environments.

### Hardware

This design uses the following Cisco Nexus 9000 Series Switches:

- Border leaf switches: Two Cisco Nexus 9396PX Switches
- Access switches (to become Cisco ACI leaf switches): Four Cisco Nexus 9396PX Switches
- NX-OS aggregation switches: Two Cisco Nexus 9372TX Switches

**Note:** Any combination of Cisco Nexus 9300 platform switches can be used at the access and aggregation layers provided that they support the server access requirements of the topology.

The following Cisco ACI equipment is added to support the conversion to and ongoing use of ACI mode:

- Cisco ACI spine switches: Two Cisco Nexus 9336PQ Switches
- Cisco APIC servers: Three APIC Enterprise Modules (APIC-EMs)

### Software Versions

This design uses the following software:

- NX-OS mode switches: Cisco NX-OS Release 7.0(3)I2(2)
- ACI mode switches: Cisco NX-OS Release 11.2.1(i) Release Build, Firmware Version 7.41
- Cisco APIC: APIC 1.2(1i) Release Build

**Servers**

The following server models were used for testing:

- Cisco UCS C-Series Rack Servers: Four Cisco UCS C220 M4 Rack Servers

**Server Operating Systems**

The following server operating systems were tested in this architecture:

- VMware ESX 5.5
  - Two ESXi 5.5 hosts
  - Active-active VMware vSphere vNetwork Distributed Switch (vDS) LACP basic support (active [IP hash])
- Microsoft Windows Server 2008 R2
  - Active-backup (no failback): Cisco NIC-teaming miniport driver (enictool.exe), with dual 10-Gbps connectivity
- Microsoft Windows Server 2012 R2
  - Active-backup (no failback): Microsoft Network Adapter Multiplexor driver, with dual 10-Gbps connectivity

## Migration Process

As with any network migration activity, the first rule is to fully understand the impact of the changes you are making. This document assumes that you are familiar with both NX-OS and Cisco ACI configuration and management concepts. Some of the preparatory tasks (for example, fabric discovery) required are well documented and are referenced with links to procedure documents or videos available on Cisco.com or YouTube.

This document focuses on the discovery, planning, and Cisco ACI configuration tasks needed to migrate traditional VLANs to endpoint groups (EPGs) configured as VLANs. This migration is commonly referred to as VLAN–to–EPG–to–bridge domain migration in Cisco ACI terminology. In ACI mode, Cisco ACI essentially replicates or emulates a traditional switching infrastructure. The Cisco ACI deployment model provides the following benefits:

- Automated VLAN provisioning
- Single point of management for fabric configuration tasks and firmware
- Telemetry information
- Host performance enhancements for broadcast, unknown unicast, and multicast traffic
- Simplification of VMware networking configuration tasks using vDS

An application-centric deployment model generally requires a detailed understanding of how applications interact with each other so that an “allowed list,” or explicit-permit, security model can be implemented. This model is a generally accepted means of reducing overall threat vectors into data center networks. This model is beyond the scope of this document.

The phases detailed here are recommended to achieve the least amount of server communication disruption. In some cases, hosts will not know that the network infrastructure has changed except that they will experience the performance enhancements that the Cisco ACI fabric switching mode brings to the data center.

General instructions for the prebuild and build phases for the Cisco ACI environment pertinent to this document are presented here. The discussion is limited to only what is necessary for this document. Readers are expected to configure their NX-OS, server, and hypervisor equipment appropriately, using vendor guidelines, to integrate with the Cisco ACI environment. A general guideline is that any servers and hypervisors should not need to be reconfigured to support Cisco ACI, but rather the Cisco ACI topology should be configured to match that used already by the servers and hypervisors and hence the traditional Cisco Nexus 9000 Series network.

### Prebuild Phase

On first receiving the equipment from the supplier, or as soon as possible, perform these tasks to help ensure that the migration can be performed.

#### Switches

- Copy the Cisco ACI software onto the switches.
- Copy the erasable programmable logic device (EPLD) images onto the leaf switches.
- Upgrade the EPLD on the leaf switches (where required).

#### APIC

- Upgrade the APIC firmware to the current code release.

#### Servers

- Install the Cisco miniport driver on Windows Server 2008 R2 and enable active-backup teaming (Option 1 in enictool.exe setup).
- Set Windows Server 2012 R2 NIC teaming load balancing to dynamic mode with active-backup teaming.

### Cisco ACI Base Build Phase

Complete the following tasks prior to migrating to Cisco ACI to help ensure that the migration process is as transparent and automated interactive as possible.

#### Switches

- Precable the Cisco ACI border leaf switches to the NX-OS mode switches using 40-Gbps cabling. You can perform this step at any time prior to the migration.
- Connect the management ports of the Cisco ACI spine switches to the out-of-band (OOB) management network at this time as a best practice.
- Connect the APICs to the Cisco ACI border leaf switches and allow the Cisco ACI network to be discovered.
- Connect the APICs to the OOB management network through their management ports so that they can perform virtual machine manager (VMM) integration before the servers are redeployed to the Cisco ACI fabric.

#### Cisco ACI Fabric

- Fabric configuration
  - Preprovision the leaf and spine switches with their specific names, node IDs, and serial numbers.
  - Set Network Time Protocol (NTP) on the Cisco ACI fabric to point to an external time source through OOB management.
  - Set the Border Gateway Protocol (BGP) route reflectors on the Cisco ACI fabric to point to the spine nodes.

- Set OOB IP addressing for all spine and leaf switches that will be connected to Cisco ACI using the same management IP addresses currently in use for each switch wherever possible (to reduce the need for modification of management policy).
- Perform OOB VMM integration between the APIC and the VMM instance so that the appropriate VLAN-backed port groups can be created on the VMM platform in the form of an APIC-managed distributed virtual switch (DVS).

- Interface configuration
  - vPC
    - Configure border leaf switches with back-to-back vPC connections to the Cisco Nexus 9000 Series aggregation switches using front-facing ports on the aggregation switches (not the 40-Gbps uplink ports).
    - Configure ESX servers as either host vPC servers (as shown in this document) or active-standby servers to replicate the model used prior to the migration to the ACI mode vDS.
  - Access
    - Configure bare-metal servers as active-backup NIC teams using an access port configuration that matches their configuration prior to migration to Cisco ACI.
- Tenant configuration
  - Bridge domains
    - Configure the bridge domain configured in Address Resolution Protocol (ARP) flooding mode.
    - Configure the custom MAC address for each bridge domain so that it is identical to the HSRP MAC address used externally on the VLAN that the traffic is currently using to reduce the downtime associated with migrating the default gateway.
    - Verify that IP routing is turned on for each bridge domain (it is on by default), and do not modify the system MAC address or configure an IP subnet at the bridge domain level. This step will enable troubleshooting features that rely on IP discovery to work, but it will not affect the flow of traffic within the subnets.
  - EPGs
    - Domains: Carefully specify the correct VMM domain or physical domain to enable traffic to flow for the EPG without errors.
    - Static paths: Create physical bare-metal servers with static paths to their ports pointing to the specific node or port to which they are connected.
  - External bridge domains
    - Configure the external connections to the outside VLANs as well as east-west VLAN extension and VMM using a bridge domain–to–EPG–to–VLAN model at the data-link layer and connected through back-to-back vPCs.
  - External routed networks
    - Physically cable the external connections from the border leaf switches to the outside Layer 3 north-south networks through the core routers through a separate physical interface to that used to connect to the Cisco Nexus 9000 Series aggregation switches.

## Migration Procedure

### Pod Migration

When migrating the pods from standalone NX-OS mode to ACI mode, be sure to migrate the secondary access switch in each pod first before migrating the primary switch. This sequence helps prevent problems with traffic flow. The primary switch is the switch in each access switch pair in a pod that is the vPC primary switch.

**Step 1: Shut Down the Port to the Secondary Leaf ESX1 Uplink**

On the secondary leaf switch in Pod1 (LEAF112), shut down the interface connected to the ESXi server. This step will prevent packet loss during the ESXi uplink migration.

Commands

```
LEAF112(config)# interface port-channel 34
LEAF112(config-if)# shutdown
```

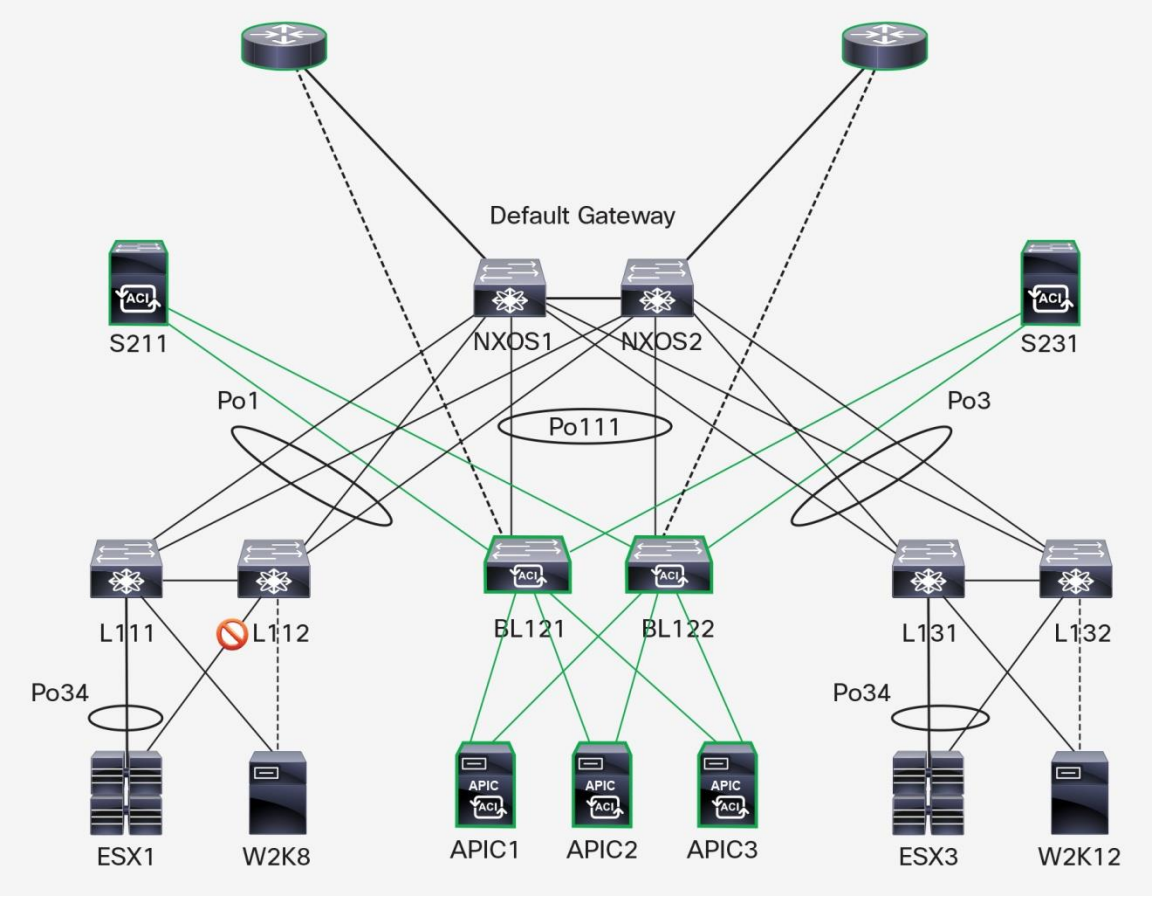


**Step 2: Migrate the ESX1 Uplink on the Secondary Leaf to the ACI Mode vDS from the NX-OS Mode vDS**

Using the vSphere web client, migrate the ESXi uplink connected to the secondary leaf switch in Pod1 (LEAF112) to use the ACI mode vDS that was preprovisioned during the build process. This uplink will remain dormant until LEAF112 is migrated to Cisco ACI and the virtual machines are migrated to use that link, helping prevent any outage.

### Process

Use the VMware vCenter Manage Host Wizard (Networking) to migrate the uplink from the NX-OS mode vDS to the ACI mode vDS.

Refer to “Using and Migrating Virtual Protection Groups on a VMware vDS in VMware vSphere” at the end of this document for detailed instructions.

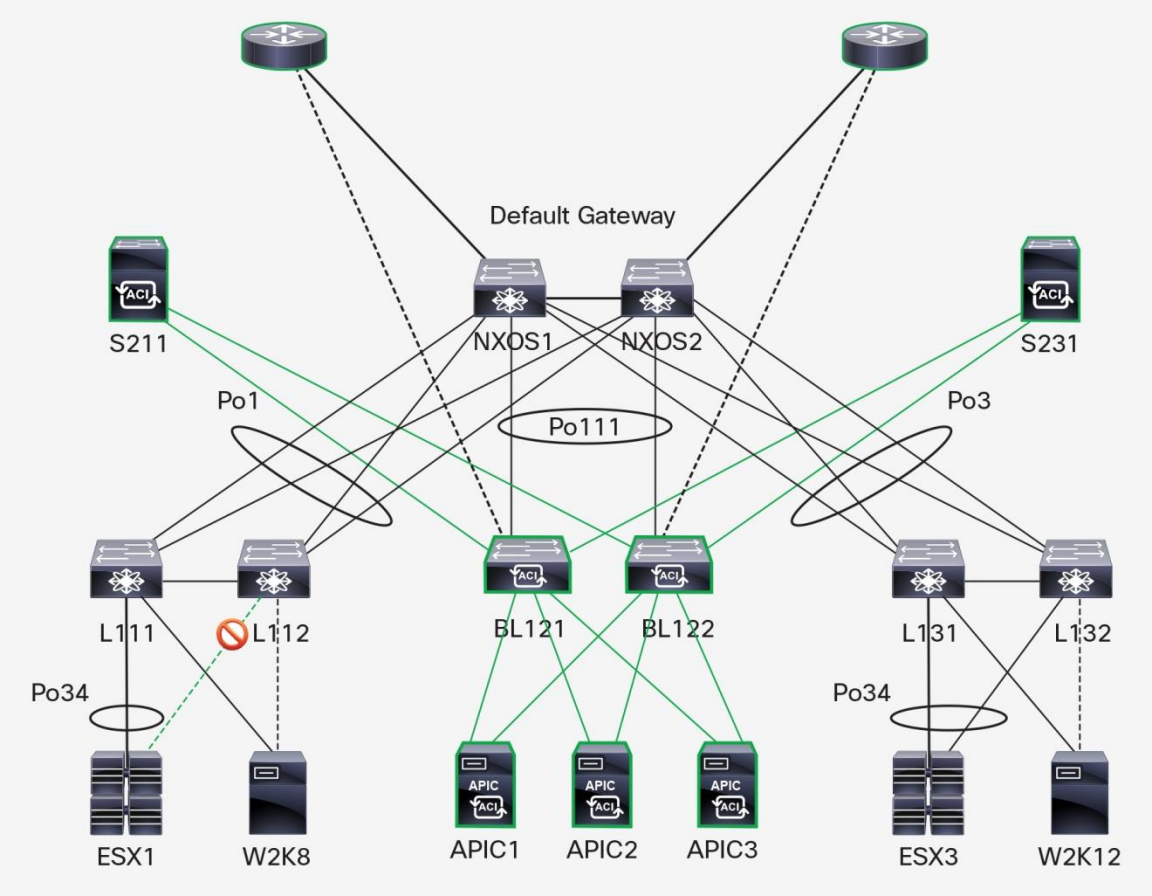


**Step 3: Reload the Secondary Leaf in Pod1 and Boot the Cisco ACI Image**

Reload the secondary switch in Pod1 and boot the leaf in ACI mode. This process should be nondisruptive to both ESX and bare-metal servers provided that all uplinks from bare-metal servers use the secondary leaf switch as the backup path.

**Note:** If a LACP active-mode host vPC to the servers is being used, a loss window of one to two seconds may be experienced in traffic to bare-metal servers while traffic is being repinned to the available vPC link.

### Commands

```
LEAF112(config)# no boot nxos
LEAF112(config)# copy running-config startup-config
LEAF112(config)# boot aci bootflash:///aci-n9000-dk9.11.2.1i.bin
LEAF112(config)# reload
```

Refer to “Converting from Cisco NX-OS to Cisco ACI Boot Mode” at the end of this document for detailed instructions.

**Note:** The bootflash memory is not reformatted regardless of the prompt, so you can roll back to NX-OS mode from ACI mode without having to replace the original NX-OS .bin file.

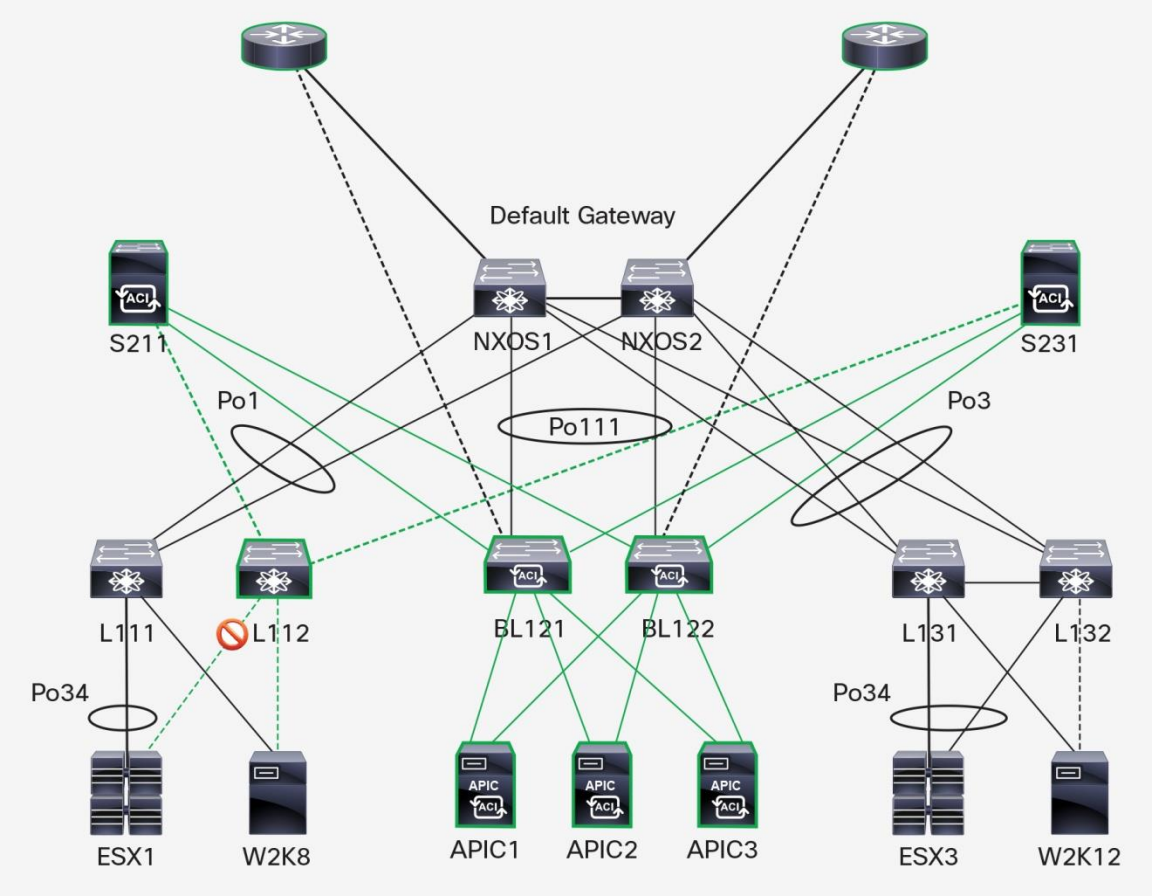


**Step 4: Cisco ACI Fabric Discovers and Initializes the Secondary Leaf**

In this noninteractive step, the Cisco ACI fabric discovers and initializes the secondary leaf switch after about four to six minutes. The discovery of the first (of two) switches in each pod is nondisruptive to traffic flows provided that the preceding steps and the preconfiguration setup are completed successfully. Navigate through the Faults pane for the leaf switch to verify that there are no pertinent faults.

**Note:** Because you are migrating the leaf switch from use as a vPC member in NX-OS mode, you may experience wiring faults. You can temporarily ignore these faults because they do not affect the functions of the leaf switch. You can remove extra cables during the cleanup phase after the migration process is complete.

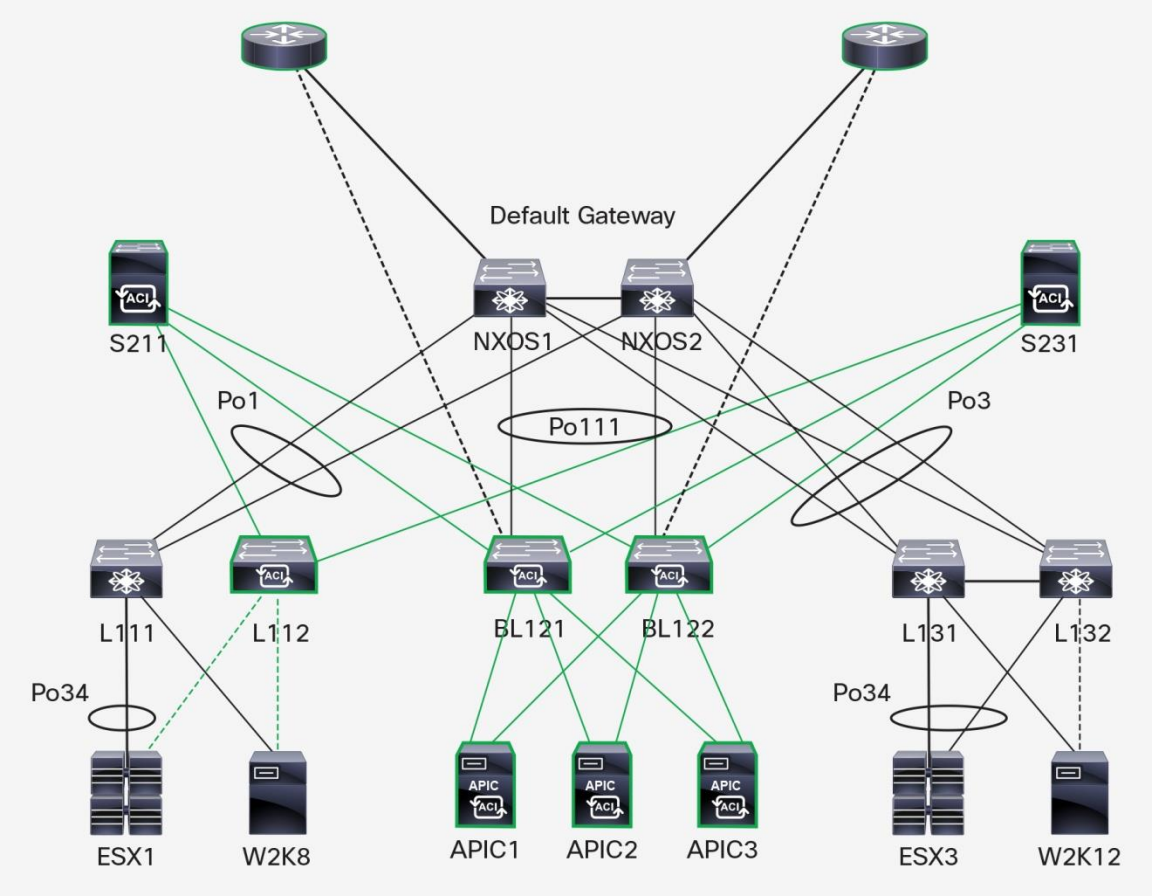


**Step 5: Migrate VMware ESX1 Virtual Machines from the NX-OS Mode vDS to the ACI Mode vDS**

Migrate virtual machines on the ESXi host in Pod1 to use the Cisco ACI network rather than the NX-OS network for uplink connectivity. This process will incur an unavoidable, but brief, one or two seconds of outage on virtual machines while they are migrated between virtual distributed switches.

### Process

Use the vCenter Manage Host Wizard (Networking) to migrate the virtual machines from the original Port Group on the NX-OS mode vDS to the corresponding Port Group on the ACI mode vDS in the thick client. Alternatively, use the vSphere web client to perform this migration.

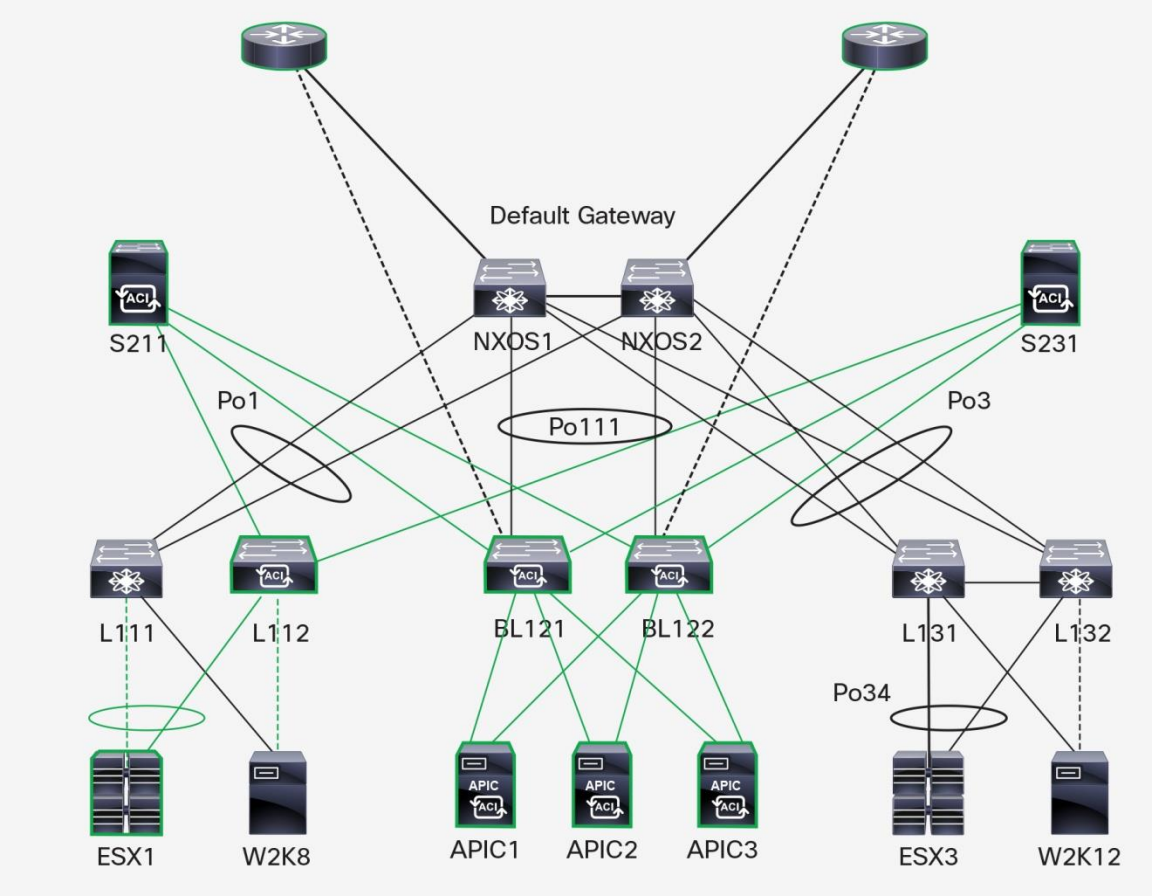


**Step 6: Shut Down the Port to the Primary Leaf ESX1 Uplink**

On the primary leaf switch in Pod1 (LEAF111), shut down the interface connected to the ESXi server. This step prevents packet loss during the ESXi uplink migration and is necessary to avoid traffic flow pinning problems from the network to ESXi. No virtual machines are using this link to forward traffic at this point.

Commands

```
LEAF111(config)# interface port-channel 34
LEAF111(config-if)# shutdown
```

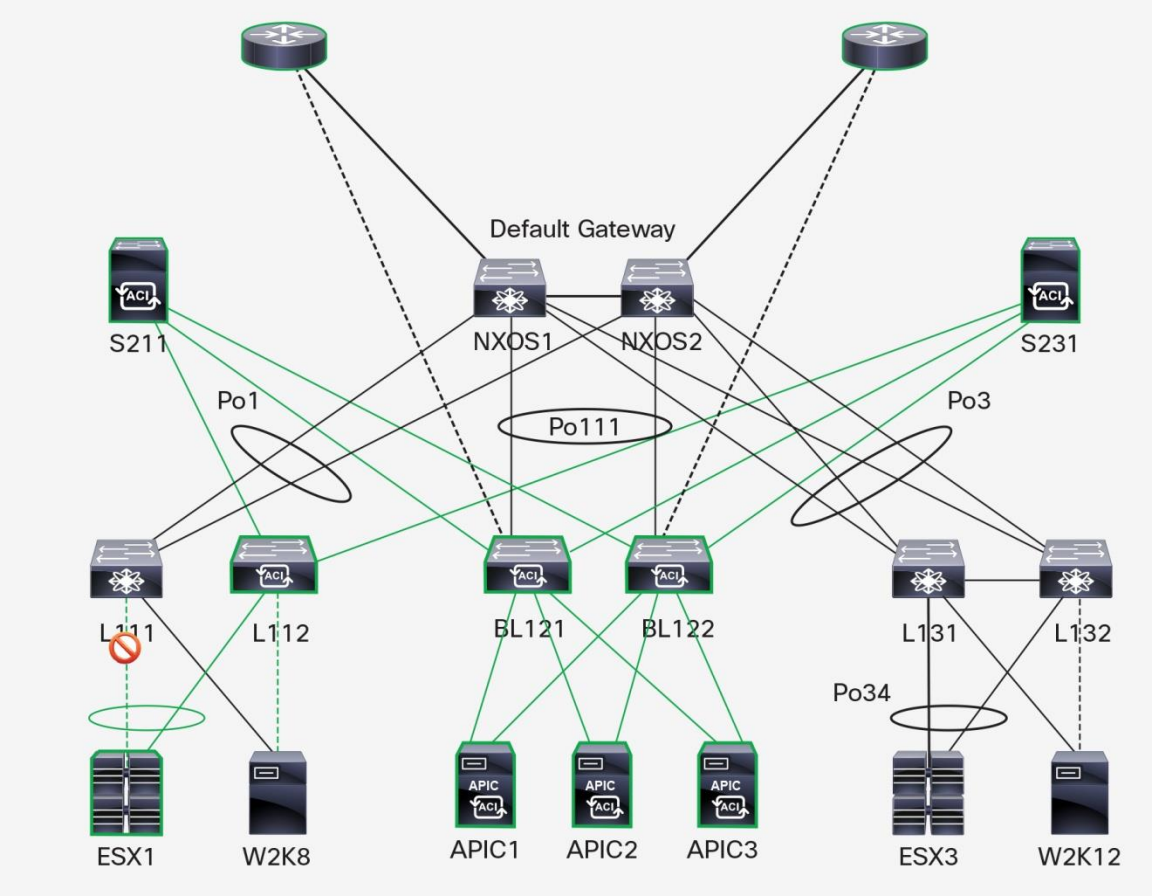


**Step 7: Migrate the ESX1 Uplink on the Primary Leaf to the ACI Mode vDS from the NX-OS Mode vDS**

Using the vSphere web client, migrate the ESXi uplink connected to the primary leaf switch in Pod1 (LEAF111) to use the ACI mode vDS that was preprovisioned during the build process. This uplink will remain dormant until LEAF111 is migrated to Cisco ACI, so no outage should be experienced.

### Process

Use the vCenter Manage Host Wizard (Networking) to migrate the uplink from the NX-OS mode vDS to the ACI mode vDS.

Refer to “Converting from Cisco NX-OS to Cisco ACI Boot Mode” at the end of this document for detailed instructions.

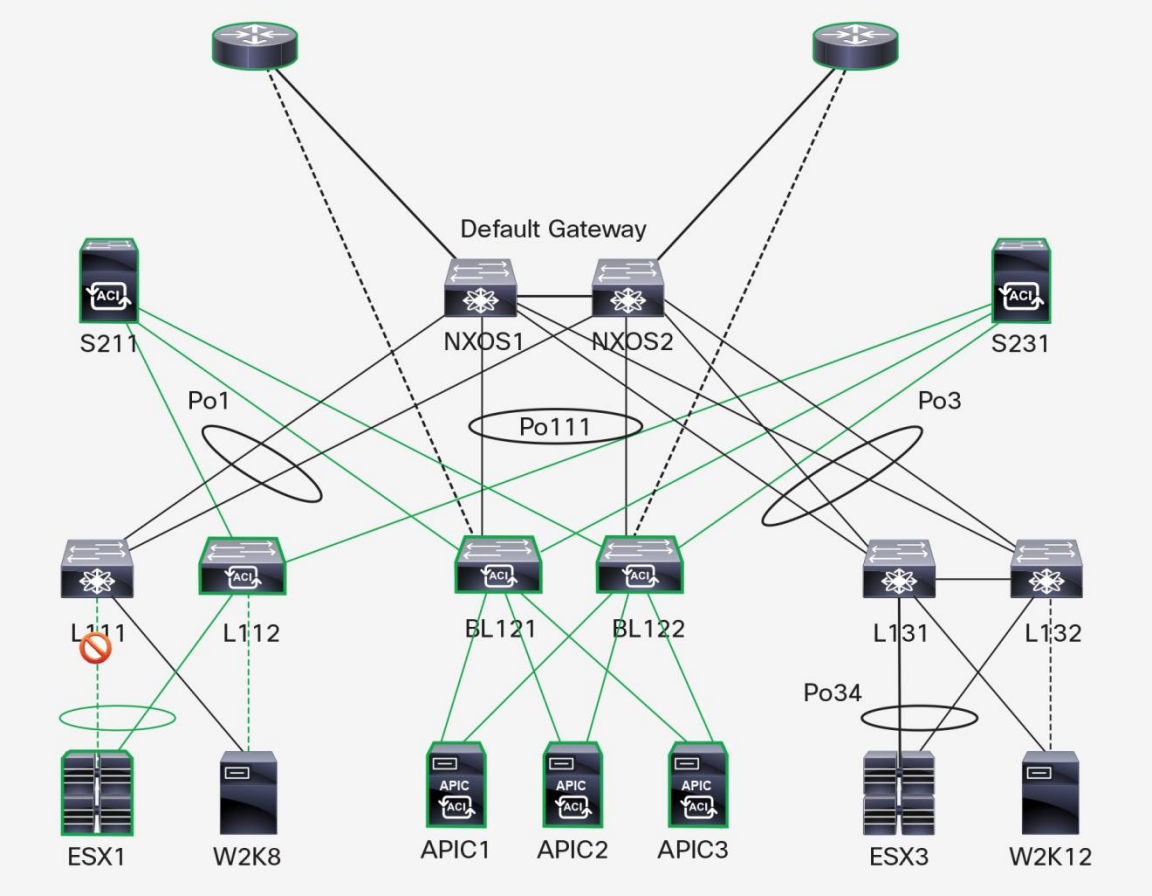


**Step 8: Reload the Primary Leaf in Pod1 and Boot the Cisco ACI Image**

Reload the primary switch in Pod1 and boot the leaf switch in ACI mode. This process should be nondisruptive to ESXi servers. The bare-metal server will switch its active link to the Cisco ACI fabric, resulting in a loss of about one second of traffic to all destinations. The build instructions specified the use of active-backup (no failback) in the Windows 2008 and 2012 R2 driver configurations, so no failback occurs during the switch reload process. The end result is that the Windows NIC that was active prior to the procedure becomes the backup at the end of the procedure, and the backup becomes the active link.

**Note:** If this is an undesirable behavior, you can configure failback at the Windows driver side. However, this configuration was not tested because each failover incurs a small (about one second) switchover packet loss.

Commands

```
LEAF111(config)# no boot nxos
LEAF111(config)# copy running-config startup-config
LEAF111(config)# boot aci bootflash:///aci-n9000-dk9.11.2.1i.bin
LEAF111(config)# reload
```

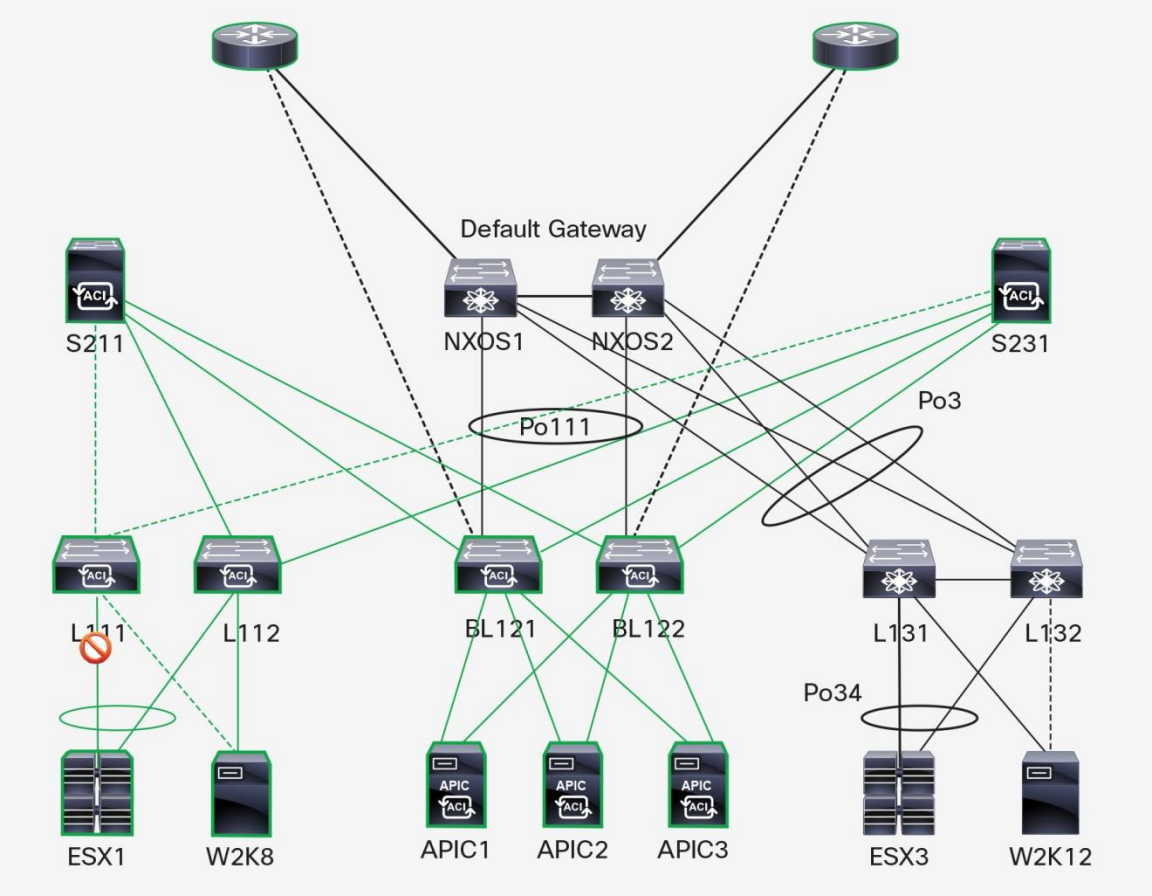


**Step 9: Cisco ACI Fabric Discovers and Initializes the Primary Leaf**

In this noninteractive step, the Cisco ACI fabric discovers and initializes the primary leaf switch after about four to six minutes. The discovery of the second (of two) switches in each pod is disruptive to traffic flows for servers with vPC connections. This disruption occurs because the Cisco ACI fabric needs to create and initialize the vPC connections to servers, and it is unavoidable. The outage was observed to be 60 to 90 seconds over a 150-second period during testing. Because of this impact, if at all possible you should perform steps 8 and 9 during a maintenance window, as stated in the “Conclusion” section, as for any production switch upgrade.

Before proceeding, navigate through the Faults pane for the leaf switch to verify that there are no pertinent faults after convergence has occurred.

**Note:** Because you are migrating the leaf switch from use as a vPC member in NX-OS mode, you may experience wiring faults. You can temporarily ignore these faults because they do not affect the functions of the leaf switch. You can remove extra cables during the cleanup phase, after the migration process is complete.

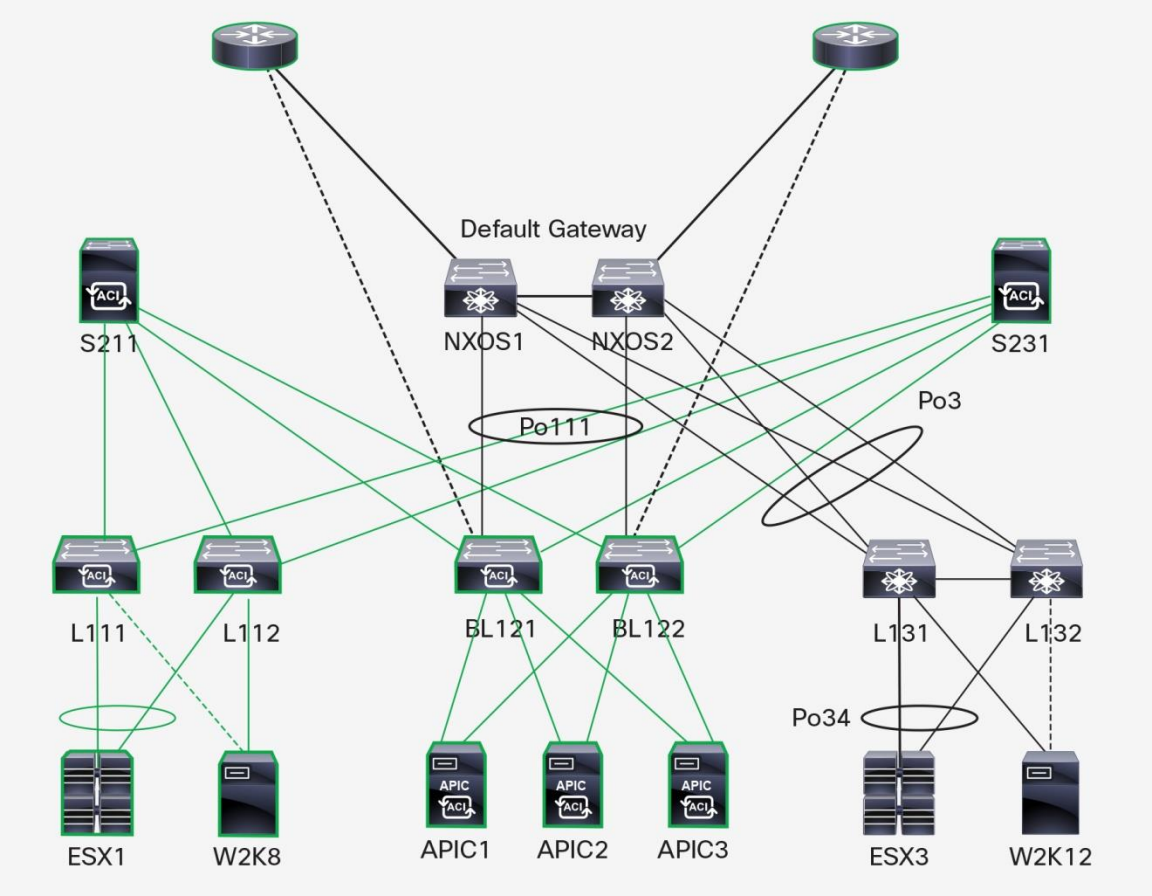


**Step 10: Migrate Other Pods**

Repeat steps 1 through 9 for each additional pod during maintenance windows. For the purposes of this document, a second pod was deployed running Windows 2012 R2 to provide an external point of connectivity during testing. Its migration process is not shown here to avoid repetition.

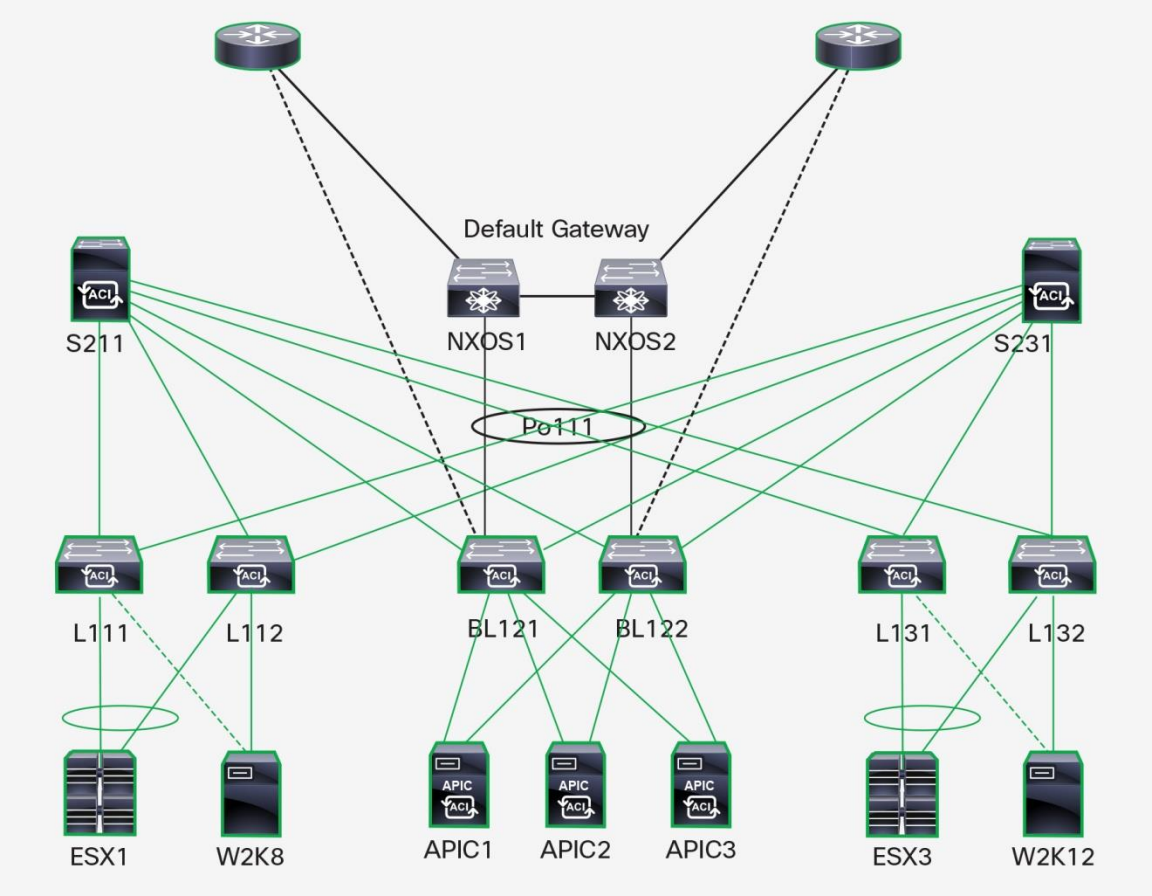


### Migrate the Default Gateway to ACI Mode

Now migrate the default gateway to ACI mode.

#### Step 1: Enable L3Out from the Border Leaf Switches to the Core Routers

Before you enable IP subnets for the bridge domains, configure a Layer 3 Outside (L3Out) network (also known as an external routed network) to point Cisco ACI to the core routers as their routing next hop. This step allows routes for subnets converted to Cisco ACI to be advertised to the core network and to the aggregation switches to avoid routing topology inconsistencies when you migrate each subnet in the next step.

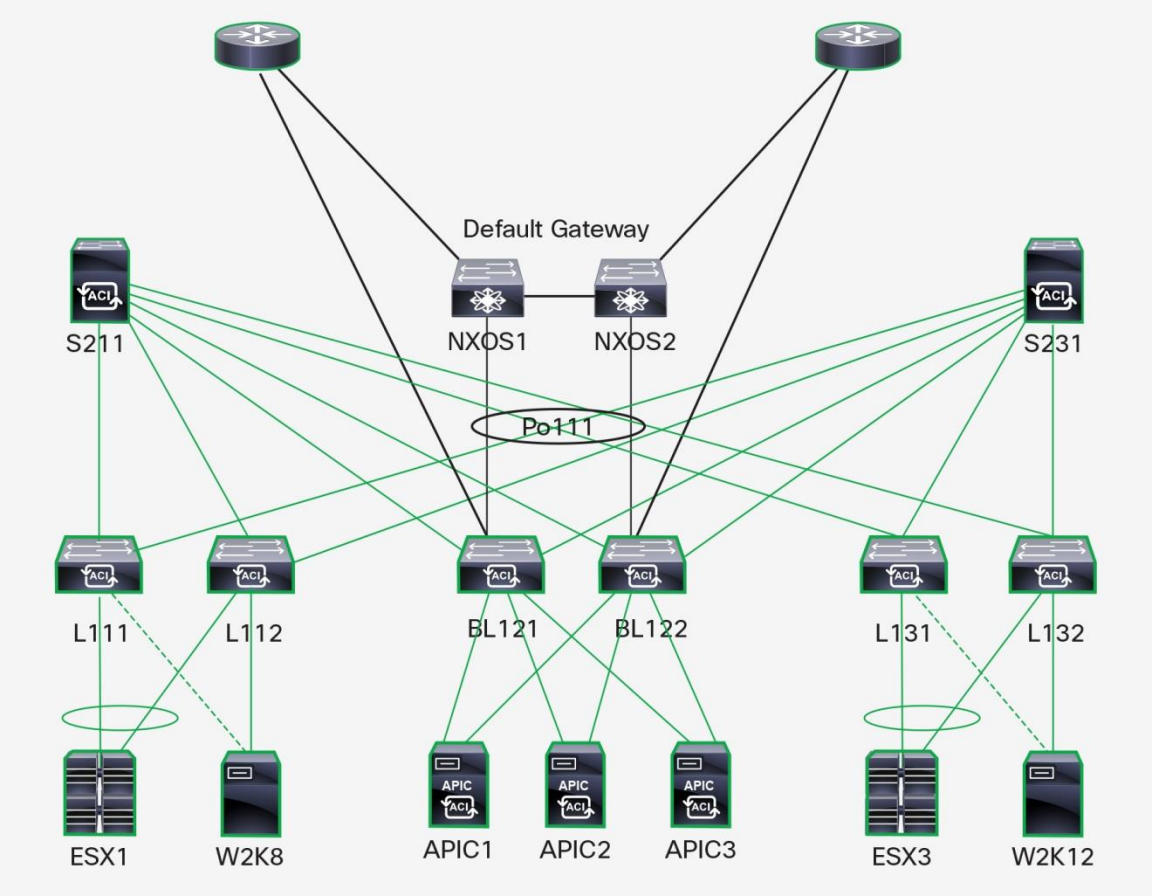


**Step 2: Enable IP Subnets on Bridge Domains with the HSRP Virtual MAC Address**

To convert the Cisco ACI bridge domains to use the persistent default gateway in the Cisco ACI fabric rather than the external NX-OS HSRP address as the default gateway, the unicast routing function must be enabled for each bridge domain in the fabric (this function is enabled by default). In addition, you must configure an IP subnet for Cisco ACI using the HSRP virtual MAC (vMAC) address and IP address to route the traffic within Cisco ACI without clients having to use ARP again for the default gateway IP address. You can perform this step through either the APIC CLI or the GUI.

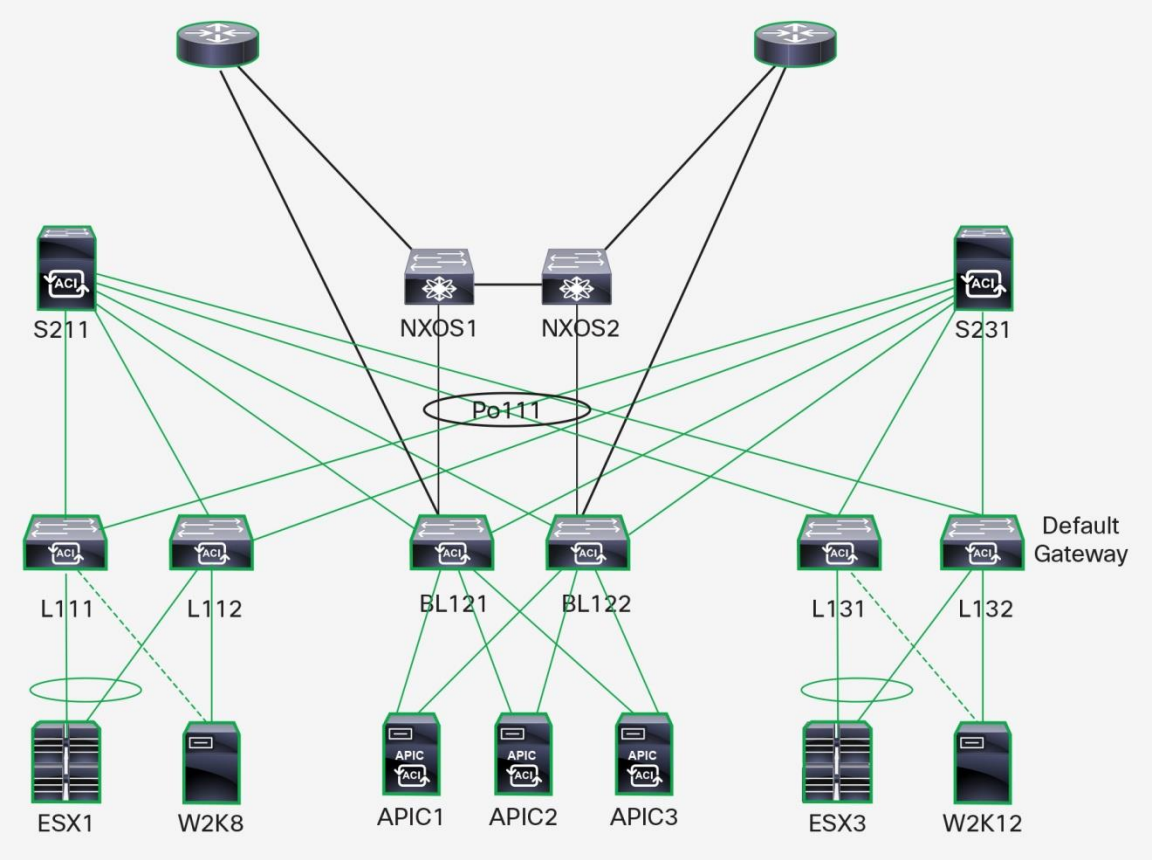


**Step 3: Shut Down the vPC Connection to Cisco ACI on Cisco NX-OS Aggregation Switches**

To verify the location of the default gateway, shut down the vPC connection between the NX-OS aggregation switches and Cisco ACI. This step is easiest to accomplish from the NX-OS side before you recommission the aggregation switches as additional Cisco ACI leaf switches in the fabric where possible.

### Commands

On NX-OS Switch 2, in the global configuration mode, enter these commands:

```
NXOS-S2(config)# interface port-channel 111
NXOS-S2(config-if)# shutdown
```

On NX-OS Switch 1, in global configuration mode, enter these commands:

```
NXOS-S1(config)# interface port-channel 111
NXOS-S1(config-if)# shutdown
```

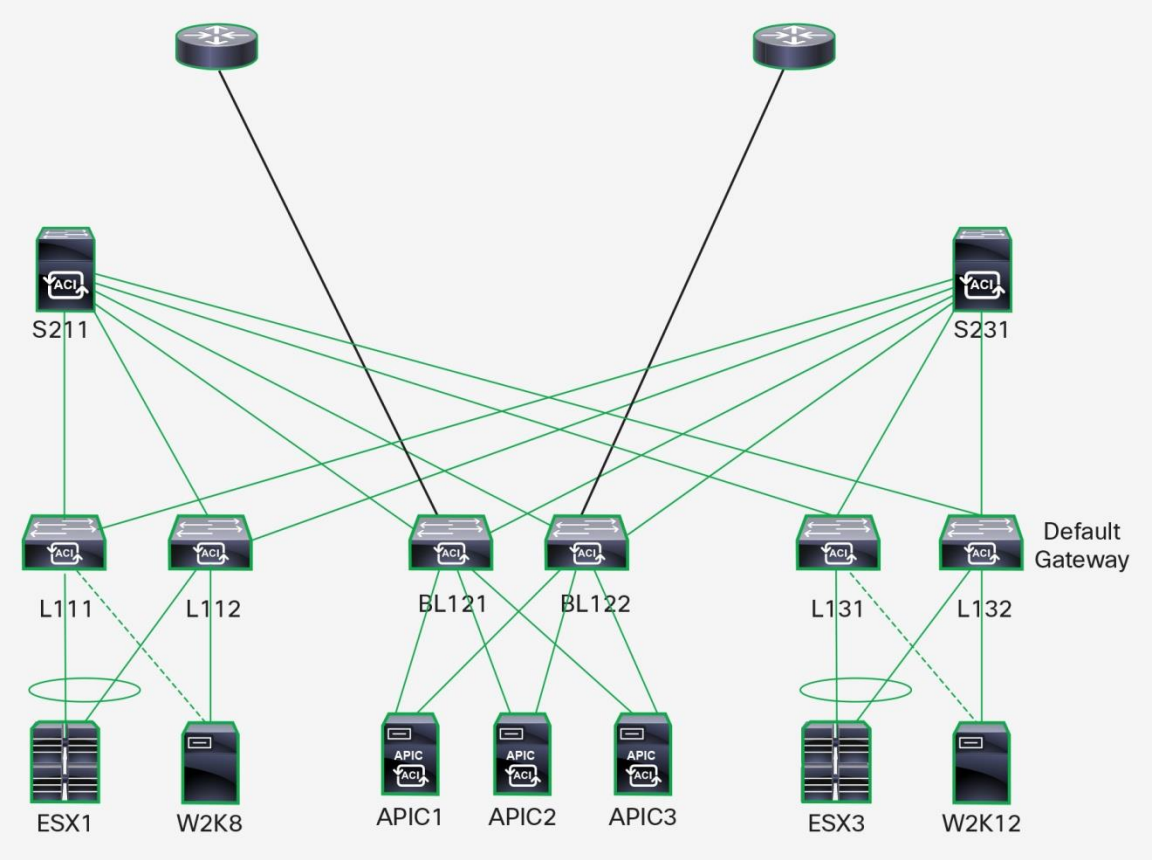


## Conclusion

If you configure your NX-OS and Cisco ACI environments ahead of time, you can easily replicate and automate migration from NX-OS mode to ACI mode on Cisco Nexus switches, and you can perform this migration in place, with little need for manual interaction. And most interactions can be automated through API calls to vSphere, NX-OS, or Cisco ACI; however, you would need to write scripts to configure this automation.

The greatest challenge encountered in the tests of both complete and isolated-phase migrations reported here was implementation of optimized protocol settings to reduce downtime while ensuring availability if a migration needs to be held in state or rolled back. The complexity of implementing optimal settings for LACP, vPC, Link-Layer Directory Protocol (LLDP), HSRP, and Spanning Tree Protocol demonstrates the need to define standards within the organization for server connectivity and to completely understand server failover characteristics.

This document identifies a process that can be run within a 60-minute outage window per pod with less than 3 minutes of total system outage per pod over that period, with the outage time occurring as a result of LACP negotiation during vPC establishment. Alternative setup options that potentially can reduce downtime further include the use of active-backup links on the ESX servers, the use of active-active (non-LACP) links on the ESX servers with manually configured load distribution, and the use of a single port channel rather than a vPC to any particular server. All these permutations were excluded during the testing reported here because of the suboptimal architecture or operational complexity entailed and the effect on the final state of the migration.

Note that the testing described here was performed using an isolated Cisco ACI fabric without upstream routers. Although separate isolated tests demonstrate that this integrated process should work for any environment, the customer should perform a thorough lab test using this document as a guide prior to performing a significant migration as discussed here.

Note also that the purpose of this document is not to detail what **should** be done: we recommend configuring Cisco ACI in a completely new ("greenfield") environment wherever possible. The purpose of this document is to help customers who have already deployed a Cisco Nexus 9000 Series standalone environment move forward with Cisco ACI while mitigating much of the risk of the transition.

## References and Additional Reading

**Converting from Cisco NX-OS to Cisco ACI Boot Mode**

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/upgrade/guide/b_Cisco_Nexus_9000_Series_NX-OS_Software_Upgrade_and_Downgrade_Guide_Release_7x/Converting_from_Cisco_NX_OS_to_ACI_Boot_Mode.html

**Using and Migrating Virtual Protection Groups on a VMware vDS in VMware vSphere**

https://www.vmware.com/files/pdf/vsphere-vnetwork-ds-migration-configuration-wp.pdf

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.



Printed in USA C11-736866-00 03/16