PURESTORAGE®

WHITE PAPER

Pure Storage FlashBlade and Cisco ACI Networking

Resilient connectivity and network-centric policy enforcement with FlashBlade and Cisco ACI

Contents

Introduction

This paper demonstrates and validates the deployment of Pure Storage® FlashBlade® attached to a Cisco Application Centric Infrastructure (ACI) fabric using a standards-based design. In this architecture, ACI serves as the unified system for forwarding, routing, and policy enforcement, ensuring that security and connectivity intent are centrally managed. The spine-leaf fabric provides both link- and device-level resilience, enabling fully active paths between FlashBlade and the paired-leaf Virtual Port Channel (vPC) while maintaining policy enforcement within the fabric. This design aligns with modern data center best practices, minimizes failure domains, and keeps policy control where it belongs—in the network.

This paper is intended for architects designing storage systems integrated with Cisco ACI, network engineers implementing Cisco ACI policies and vPCs, and storage administrators deploying Pure Storage FlashBlade in an ACI environment. It serves as a reference document for integrating Pure Storage FlashBlade platforms with Cisco ACI leaf pairs via vPC, that preserve ACI as the policy/control point, and validates failure states across availability scenarios and evidence outcomes via endpoint reachability, Link Aggregation Control Protocol (LACP) health, and client I/O continuity.

Objectives

The following objectives guide the validation and deployment of Pure Storage FlashBlade attached to a Cisco ACI spine-leaf fabric.

- Define best practices for connecting Pure Storage FlashBlade to Cisco ACI using a spine-leaf fabric and application-centric policy model.
- Demonstrate seamless interoperability with FlashBlade and external fabric modules (XFMs) as standard endpoints within the ACI fabric.
- Validate end-to-end resiliency under realistic disturbance scenarios, including:
  - Member link failures
  - Leaf switch outages
  - Spine-facing uplink events
- Confirm measurable outcomes such as:
  - Session continuity
  - Subsecond convergence
  - Restoration of baseline throughput
- Verify protocol behavior for Network File System (NFS), Server Message Block (SMB), and Simple Storage Service (S3) workloads under both steady-state and failover conditions (planned and unplanned), ensuring data integrity, session stability, and consistent performance across the fabric.

Solution overview

Solution components

The following components form the foundation of the solution architecture:

- **Pure Storage FlashBlade**: FlashBlade is a modular, scale-out unified file and object platform engineered for high parallel performance. A chassis expands linearly with additional blades—each adding compute, cache, and capacity—so throughput scales predictably. Native 25/40/100GbE enables high-bandwidth east–west traffic, and Evergreen® nondisruptive upgrades keep the system current without downtime.
- **Cisco Nexus N9K-C9364C Spine Switch**: A fixed-format spine with 64 × 40/100GbE (QSFP28) in a 2RU chassis, the switch is suitable for ACI spine roles with wire-rate forwarding and equal-cost multi-path (ECMP) across the fabric.
- **Cisco Nexus N9K-C93180YC-EX Switch**: A 1RU ACI leaf offering 48 × 1/10/25GbE (SFP28) and 6 × 40/100GbE (QSFP), the switch is ideal for FlashBlade vPC attachment and server access.
- **Cisco APIC M3:** This Application Policy Infrastructure Controller (APIC) serves as the centralized controller for Cisco ACI, providing policy management, automation, health monitoring, and visibility across the fabric. In this deployment, APIC M3 is used to model virtual routing and forwarding (VRF), bridge domains, endpoint groups (EPGs), and contracts as well as to manage the leaf vPC attachment for FlashBlade.

Fabric topology

The fabric follows a spine-leaf design with Cisco Nexus N9K-C9364C-FX2 spines and a dual N9K-C93180YC-EX leaf pair. The leaf pair operates as a vPC domain for downstream storage and host attachment, providing active-active forwarding, bounded failure domains, and stable policy enforcement at the access layer. This yields deterministic behavior during member link and node-level events.

Storage connectivity

FlashBlade connects to the leaf pair via LACP port-channels that are presented as a vPC across both leaf switches, enabling chassis-level redundancy and nondisruptive path failover. The ACI policy model defines:

- PC/vPC interface objects bound to an interface policy group (LACP active)
- Mirrored port blocks on both leaf switches
- Virtual local area network (VLAN)/EPG mappings for NFS and SMB traffic
- Path attachment of the EPGs to the storage port-channel

Attachment model

FlashBlade exposes dual XFM-8400 modules cabled symmetrically to the leaf pair. XFMs do not join ACI fabric membership and are treated as endpoints; they appear in the storage EPG's operational endpoint views with explicit encapsulation and leaf/interface context.

Diagrams and APIC validation artifacts

As shown in Figure 1, FlashBlade integrates with ACI as a standard endpoint in EPG **EPG-Flashblade**, learned through the vPC FB-VPC (VLAN-120). The associated 00:50:56:xx:xx:xx MAC/IP entries correspond to VMware endpoints on server access ports, demonstrating that ACI manages all network forwarding and policy enforcement for traffic to and from the FlashBlade XFM, which functions as an endpoint.

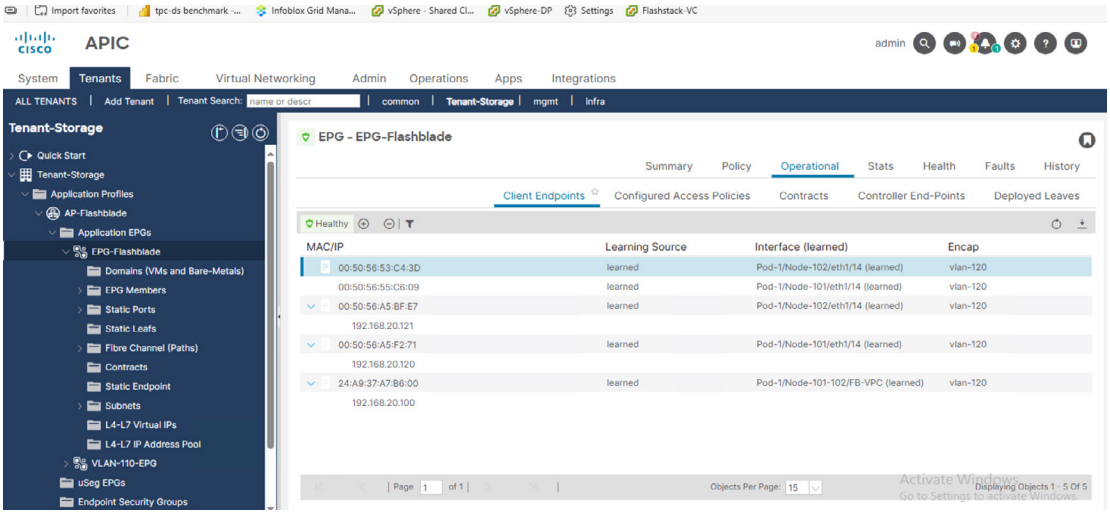

FIGURE 1 APIC tenants

Within the ACI fabric, only ACI spine and leaf switches are registered as fabric nodes; the Pure Storage FlashBlade XFM operates as an external storage point connected through the fabric rather than as part of the ACI switching infrastructure.

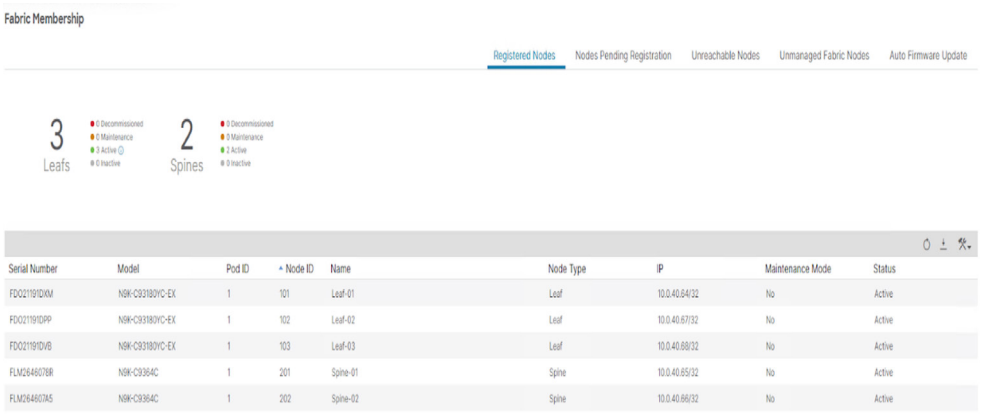

FIGURE 2 Fabric membership

Design considerations and validation goals

Figure 3 shows the high-level design for this deployment.

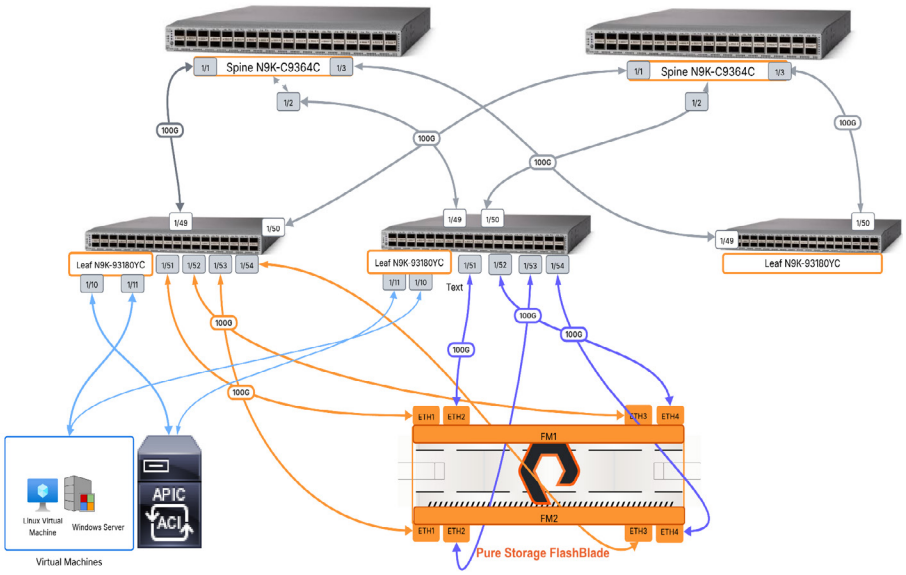

FIGURE 3 High-level design for deployment

Deployment requirements: infrastructure hardware and software matrix

Table 1 outlines the infrastructure hardware and software components required in the FlashBlade and Cisco ACI deployment.

Day 0: design component	Model/image	Version	Role/notes
ACI spines	Cisco Nexus N9K-C9364C-FX2	aci-n9000-dk9.16.0(9d)	Fabric core; deterministic north–south/east–west paths
ACI leaves	Cisco Nexus N9K-C93180YC-EX (vPC pair)	aci-n9000-dk9.16.0(9d)	vPC termination for FlashBlade link aggregation group (LAG); policy enforcement
Storage	Pure Storage FlashBlade (dual XFM-8400)	Purity//FB 4.6.0	Endpoint (not a switch); NFS/SMB/S3
Day 1: deploy component	**Model/image**	**Version**	**Role/notes**
Controllers	Cisco APIC M3	6.0(9d)	Central policy, health, visibility, automation
Virtualization host	VMware ESXi on Cisco UCS	8.0	Platform hosting APIC VM
Management plane	VMware vCenter Server Appliance	8.0	Centralized management of ESXi hosts and APIC VM deployment
Day 2: operate component	**Model/image**	**Version**	**Role/notes**
Linux client	Bare metal/VM (10/25 GbE)	—	Validation clients (NFS/S3); integrity checks
Windows client	Bare metal/VM (10/25 GbE)	—	Validation clients (SMB); integrity checks

TABLE 1 Infrastructure matrix detailing hardware and software roles across Day 0–2 of deployment.

High availability goals

The design targets continuous storage access during the following activities:

- Leaf/switch outages (planned and unplanned)
- XFM/aggregator failover on FlashBlade
- Control plane and upgrade events (APIC, including abrupt interruption)

Validated protocols and network addressing

Following are the validated protocols, addressing scheme, and segmentation strategy for this deployment.

- **Validated protocols**: NFS, SMB, and S3
- **Addressing**: IPv4 (IPv6-ready where supported)
- **Segmentation**: dedicated VLANs/EPGs per protocol (for example, VLAN 2155 for storage), with subnets aligned

Implementation guide

The following steps outline the implementation workflow for integrating FlashBlade with Cisco ACI.

1. Prepare the policy objects (tenant/VRF/bridge domain/EPG, AEP/PhysDom, and VLAN pool) and contracts (NFS/SMB/S3).
2. Configure the leaf interface policies and vPC policy group; ensure that LACP is active.
3. Bind storage EPG statically to the leaf vPC with the storage VLAN; verify endpoint learning in APIC.
4. Validate the maximum transmission unit (MTU), protocol access, and baseline throughput.
5. Execute failure tests (member, leaf, and uplink) and confirm convergence and session continuity.

Validation tests and results

Validation tests

Leaf power-off

Action: Intentional shutdown of one leaf in the vPC pair.

Expected outcome: Traffic continues via the surviving leaf, there are **no client dismounts or session resets**, reconvergence completes within **single-digit seconds**, and throughput returns to the steady-state band.

Spine-link flap (disable one leaf-to-spine uplink)

Action: Shutdown of one uplink from the active leaf toward the spines.

Expected outcome: ECMP re-route with only a **brief micro-dip**, **no file I/O timeouts**, application sessions are preserved, and leaf and fabric faults are recorded in APIC for audit.

Protocol validation

The following validation use cases were executed to assess end-to-end protocol resiliency and client session stability during planned and unplanned failover events.

NFS validation (Linux clients)

A large-object file transfer was initiated from a Linux client to a FlashBlade NFS export over the designated storage EPG/VLAN.

Observations during the validation tests were as follows:

- During a leaf or uplink event, the NFS mount persisted, the transfer resumed automatically at line rate after a subsecond dip, and no mount renegotiation or abnormal retransmissions were detected.
- Sustained throughput remained within the expected link speed range.
- Post-copy confirmed data integrity.

SMB validation (Windows clients)

A single large-file transfer was performed from a Windows client to a FlashBlade SMB share mapped through the storage EPG.

Observations during the validation tests were as follows:

- A transient link event on one leaf produced no session resets or user-visible errors, confirming seamless convergence and preservation of SMB session state.
- Throughput aligned with host network interface card (NIC) and policy expectations.
- File hash verification confirmed content integrity.

S3 validation (object access)

A multipart PUT/GET operation of a multi-gigabyte object was executed against the FlashBlade S3 endpoint within the storage EPG.

Observations during the validation tests were as follows:

- A brief vPC member event resulted in only a momentary throughput dip; client retries stayed within normal thresholds and no operation failures occurred.
- Aggregate transfer time remained within the baseline range.
- All parts completed successfully with HTTP 200 responses; the final ETag and hash matched the source.

Validation actions and results

Table 2 summarizes key validation actions and outcomes across failure scenarios, protocol resiliency, MTU handling, and storage attachment behavior.

Category	Validation actions	Results
Failure scenarios	Member link flap, full leaf failure, and spine-facing uplink events	Session continuity; convergence; throughput returns to baseline
Protocol resiliency	NFS/SMB/S3 large-object transfers + integrity checks	No session resets during controlled events
Jumbo MTU	9,000-byte frame handling across host, ACI, and FlashBlade	Zero fragmentation; stable latency; no retransmit spikes
Storage attachment	FlashBlade and XFM not in fabric membership	Fabric membership excludes FlashBlade XFM

TABLE 2 Validation testing results

Conclusion

These validations confirm that FlashBlade XFMs behave as storage endpoints rather than switching devices. They simply present standard Ethernet interfaces to the Cisco ACI fabric, with all routing, forwarding, and policy enforcement remaining in ACI. This preserves a clean separation of roles while delivering predictable, resilient connectivity for file and object workloads.

Appendix: Pure Storage FlashBlade configuration guide

In high-performance storage designs, multiple active paths are required for throughput, availability, and fault tolerance. With Pure Storage FlashBlade, two XFMs connect to an ACI leaf vPC pair (Leaf-101/102) using LACP port-channels. This logically aggregates the member links so APIC treats them as a single interface for policy/EPG binding while distributing traffic across links and surviving individual failures. Figure 4 shows an inventory and health summary for XFM modules detected on FlashBlade.

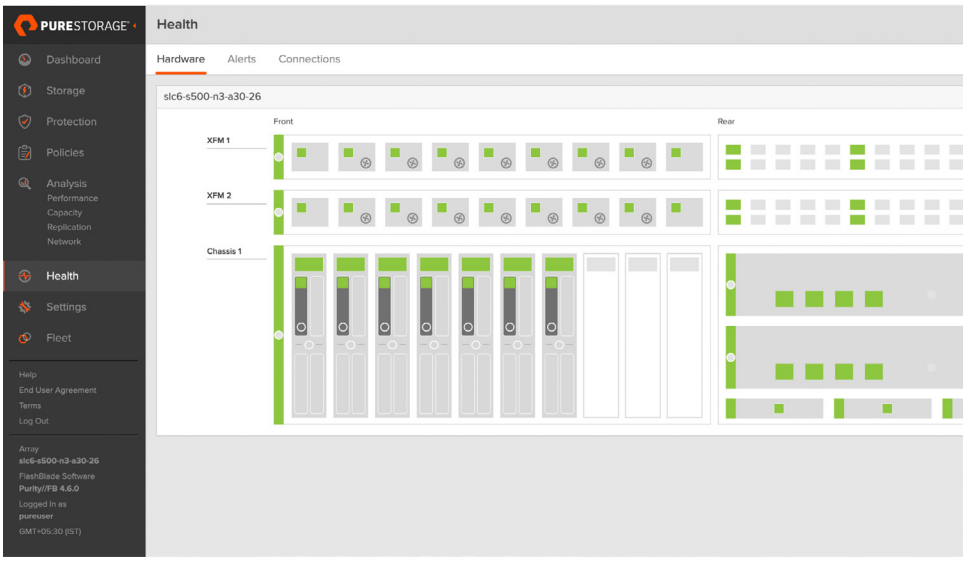

FIGURE 4 XFM module inventory and health summary

This network topology is designed to deliver high availability, redundant data paths, and enhanced aggregate bandwidth for FlashBlade file (NFS and SMB) and object (S3) services. By leveraging dual Fabric Interconnects and properly configured link aggregation groups (LAGs) with LACP, the architecture minimizes single points of failure and ensures consistent performance under varying workloads (see Figures 5 and 6).

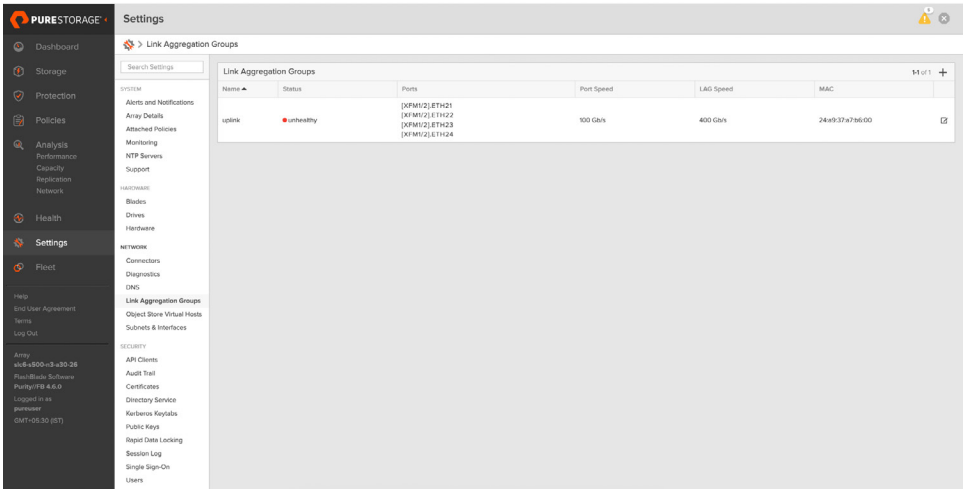

FIGURE 5 Link aggregation groups

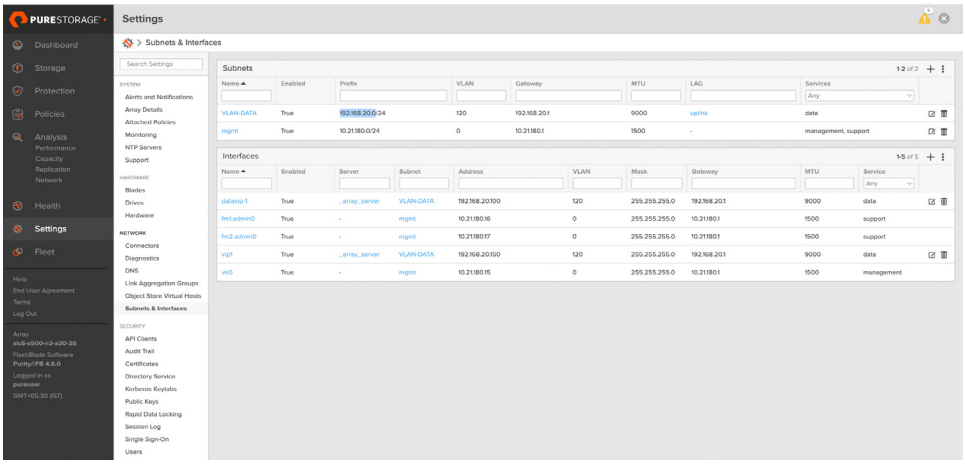

FIGURE 6 Subnets and interfaces

Port configuration: Ethernet and virtual interfaces

Each of the physical Ethernet ports on the Pure Storage FlashBlade should be enabled so that the ports are ready for connectivity once the appropriate configuration is set on the ACI leaf switches. For direct connectivity from the Pure Storage FlashBlade to the ACI leaf switches, use subnets with VLAN interfaces that match the VLAN IDs defined for your environment.

Create a subnet with a VLAN interface for each data path to connect to these subnets with VLANs.

1. To create a subnet on the Pure Storage FlashBlade, navigate to the FlashBlade Network Settings page (**Settings > Network**).
2. Click the **+** icon in the Subnets area.
3. In the Create Subnet pop-up window that appears, enter the following details:
   - **Name**: the name of the subnet used within the FlashBlade; it is recommended to include the data path (A or B)
   - **Enabled**: set by default and should remain enabled
   - **Prefix**: the prefix for the network subnet in Classless Inter-Domain Routing (CIDR) notation, which defaults to /24
   - **VLAN**: tags the VLAN to be used on the subinterface attached to this subnet; this ID will match the details defined for your environment
   - **Gateway**: the gateway for your network subnet; not required in this direct-attach configuration
   - **MTU**: sets the MTU to be used by the subinterfaces that inherit this setting from the subnet; the general recommendation is to use the standard MTU of **9000**
4. Click **Create** to finish the creation of a subnet for the FlashBlade.
5. Repeat steps 1–4 to create a second subnet with the appropriate details of the second data path.
6. Once the two subnets are created, interfaces can be added to them by clicking the **Add Interface** button under the Interfaces column of the subnet.
7. In the Add Interface of Subnet '[Subnet Name]' pop-up window that appears, click the **Name** drop-down menu and select the physical Ethernet interface that is directly connected to the ACI leaf switch for the data path of the subnet.
8. Once the correct subinterface has been picked from the menu, click **Save** to add the interface to the subnet.
9. Repeat steps 6–8 to add interfaces to each subnet so that a minimum of two subinterfaces, connected to two separate physical interfaces, are configured to provide redundant connectivity from the FlashBlade to the ACI leaf switches.

Physical port connectivity

The Pure Storage FlashBlade is architected for performance and high availability through dual internal controllers (XFMs) each supporting multiple Ethernet interfaces:

- **XFM-1**: connected to ACI Leaf A (FI-A) via xfm1.eth4 and xfm1.eth5
- **XFM-2**: connected to ACI Leaf B (FI-B) via xfm2.eth4 and xfm2.eth5

These interfaces form the physical foundation for aggregating throughput and ensuring resilience in file and object services.

Link aggregation groups

To maximize throughput and enable link-level redundancy, a LAG is configured across multiple Ethernet ports. LAG1 combines xfm1.eth4, xfm1.eth5, and xfm2.eth4 and presents as a port-channel (vPC) spanning leaf switches. It ensures balanced traffic distribution and fault tolerance across both Fabric Interconnects.

LAGs provide the following benefits:

- **Redundancy**: maintains connectivity if a physical link fails
- **Increased bandwidth**: allows load distribution across multiple active links

Virtual interface abstraction

A **virtual interface** abstracts the underlying physical and aggregated Ethernet ports, providing a single IP endpoint for client access.

Primary functions include:

- **Unified endpoint**: All NFS, SMB, and S3 Object traffic routes through the virtual IP (VIP).
- **Failover support**: The VIP dynamically migrates to an available controller in case of a path or XFM failure.

References

- FlashBlade product page
- Cisco Application Policy Infrastructure Controller (APIC)—Cisco Application Centric Infrastructure (ACI) Design Guide

ID1066-01-en-11/25