

WHITE PAPER

FlashBlade Direct-attach Connectivity for Cisco UCS X-Series

Deployment Guide for Cisco UCS X-Series with Intersight Managed Mode and UCS Manager

Contents

| | |
|--|----|
| Introduction | 3 |
| Considerations for Running in Direct-attach Mode | 3 |
| Design Considerations | 3 |
| Ethernet Port Configuration and Topology Constraints | 4 |
| Connectivity Diagram | 4 |
| Infrastructure Components | 5 |
| Configuration Guide: Pure Storage FlashBlade | 5 |
| Physical Port Connectivity | 7 |
| Link Aggregation Groups | 7 |
| Virtual Interface Abstraction | 7 |
| Pure Storage FlashBlade Network Configuration Details | 8 |
| Use Case: File Services over Direct-attached Ethernet (NFS/SMB) | 8 |
| Object Services over Direct-attached Ethernet (S3 Object) | 8 |
| Port Configuration: Ethernet and Virtual Interfaces | 8 |
| Configuration Guide: Cisco Intersight Managed Mode | 10 |
| Configuration Guide: Cisco UCS Manager | 12 |
| UCS Storage Port Configuration | 12 |
| Assigning VLANs for FlashBlade Storage Traffic on Cisco UCS | 15 |
| Enable Jumbo Frames for Performance (Optional) | 15 |
| Appliance Interface Configuration: Ethernet | 16 |
| Failover Test Summary: Ensuring High Availability of FlashBlade File Services | 17 |
| XFM Module Reboot (xfm1/xfm2) | 17 |
| Fabric Interconnect Reboot | 17 |
| Object Service Validation (During Failover) | 18 |
| Conclusion | 19 |
| Additional Resources | 19 |



Introduction

This white paper presents a streamlined reference architecture for deploying Pure Storage® FlashBlade® directly with Cisco UCS Fabric Interconnects. The solution supports a wide range of FlashBlade protocols, including Network File System (NFS), Server Message Block (SMB), and S3 Object, making it well-suited for diverse use cases such as data analytics, backup and restore, software development, media workflows, and artificial intelligence.

By eliminating the need for intermediary network switches between storage and compute, this direct-attach model reduces infrastructure complexity, lowers power and rack footprint, and accelerates deployment time. It provides an ideal foundation for private cloud environments that demand consolidated file and object services with simplified management and strong performance.

The purpose of this white paper is to guide organizations that have selected or are considering a direct-attach model, providing key design insights and operational considerations. The architecture leverages a direct-attach topology, in which Cisco UCS Fabric Interconnects are directly connected to a Pure Storage FlashBlade. This implementation supports both Intersight Managed Mode (IMM) and UCS Manager (UCSM), offering flexibility in how the infrastructure is managed.

Considerations for Running in Direct-attach Mode

The implementation operates a FlashStack® environment in a direct-attach configuration to validate the NFS, SMB, and S3 Object protocols on the Pure Storage FlashBlade. This setup eliminates the need for intermediate Cisco Nexus switches, providing a simplified and cost-effective alternative for connecting compute and high-performance unified storage.

The FlashBlade delivers consistent high availability, operational simplicity, and exceptional throughput for data-intensive file-based applications such as AI/ML pipelines, media processing, backup and restore, and DevOps workloads.

Design Considerations

In Cisco UCSM direct-attach environments, network policies and port configurations are centrally managed via Fabric Interconnects. Direct FlashBlade connectivity, without top-of-rack (ToR) switches, introduces a critical design consideration: **scalability and resource sharing across Unified Computing System (UCS) domains.**

- **Single UCS domain confinement:** A direct-attached FlashBlade is limited to providing storage services within a single UCS domain. Its resources cannot be directly shared with servers in separate UCS domains.
- **Scalability trade-off:** While simplifying single-domain networking, this design restricts broader sharing. Multi-domain resource sharing typically requires external ToR switches (for example, Cisco Nexus) for a common network fabric.

When considering a direct-attached FlashBlade, it is important to assess current and future multi-UCS domain sharing requirements.

For FlashBlade unified file and object services, the trade-offs are generally manageable. Adhering to validated design practices—including proper port channeling, Link Aggregation Control Protocol (LACP), virtual local-area network (VLAN) tagging, and interface-to-workload mapping—allows effective monitoring and maintenance. UCS command-line interface (CLI) tools and FlashBlade real-time metrics provide sufficient operational insight for robust performance, offering a balanced trade-off between architectural simplicity and operational control.



Ethernet Port Configuration and Topology Constraints

In a direct-attach model, Ethernet ports on UCS Fabric Interconnects must be explicitly configured as appliance ports to connect with the FlashBlade. Unlike traditional Nexus-based topologies, port channeling is supported but requires precise LACP configuration and policy alignment between the Fabric Interconnects and FlashBlade.

- VLANs must be manually defined at both ends (UCS Fabric Interconnects and FlashBlade) to ensure proper network segmentation and traffic isolation for NFS, SMB, and S3 Object services.
- While removing external switching reduces complexity, it also reduces flexibility—especially when scaling or adding more uplinks or arrays.

Connectivity Diagram

Figure 1 shows the physical connections between the FlashBlade and Cisco UCS Fabric Interconnects in the reference architecture.

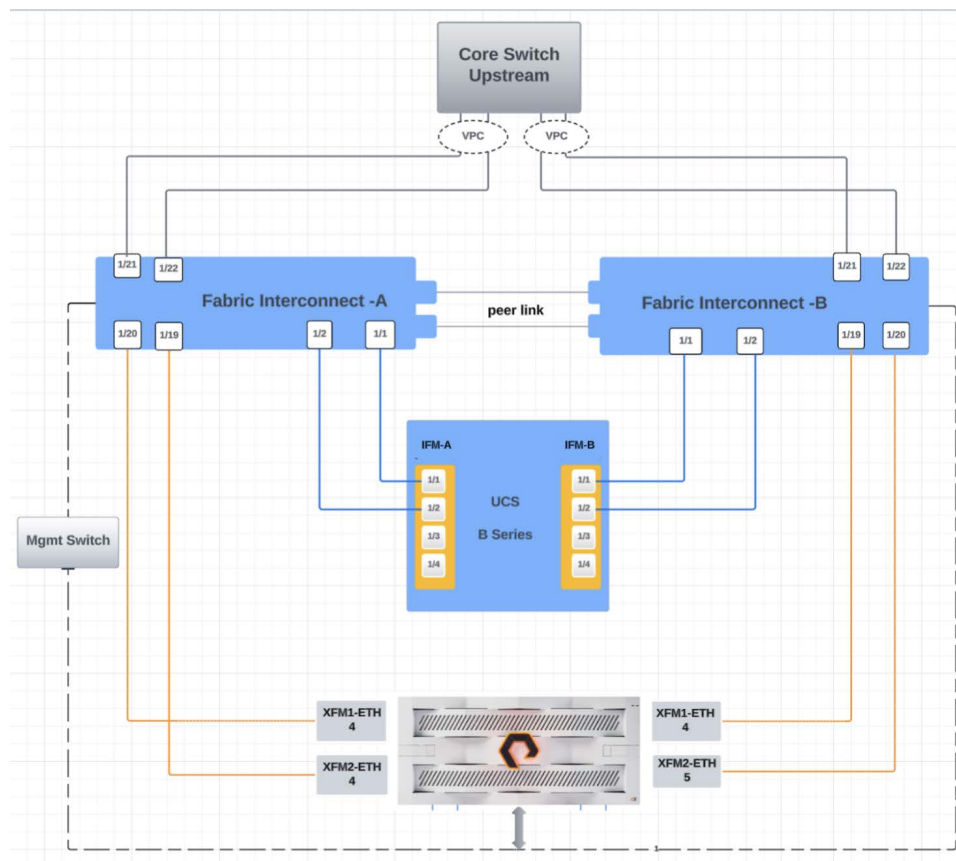


FIGURE 1 Diagram of physical connections between FlashBlade and Cisco UCS Fabric Interconnects.

Infrastructure Components

Table 1 lists the components used to build the configuration of a direct-attached FlashStack.

| Infrastructure Component | Model | Version |
|---------------------------------|--------------------------------|---------------------------|
| Storage | FlashBlade//S™, FlashBlade//E™ | Purity//FB 4.2.0 or later |
| Fabric Interconnects | Cisco UCS 6536 | - |
| Compute Chassis | Cisco UCS X9508 | - |
| Intelligent Fabric Module (IFM) | Cisco 9108 IFM | - |
| Compute Nodes | Cisco UCS X210c, X410c, X215c | M6 or later |
| Linux Operating System | Bare metal | - |
| Windows Operating System | Bare metal | - |

TABLE 1 All infrastructure components and models or versions used within the documented FlashStack deployment in direct-attach configuration

Configuration Guide: Pure Storage FlashBlade

In high-performance storage environments, multiple network paths are critical to maximizing throughput, availability, and fault tolerance. When connecting the Pure Storage FlashBlade to upstream network components—such as Cisco UCS Fabric Interconnects (FI-A and FI-B)—via multiple physical links, it’s essential to logically aggregate these links to ensure optimal performance and reliability. Technologies such as link aggregation groups (LAGs) and LACP play a key role in achieving these objectives.

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 LAGs with LACP, the architecture minimizes single points of failure and ensures consistent performance under varying workloads.

Figure 2 outlines the key infrastructure components in a direct-attached FlashBlade deployment, highlighting their roles in enabling resilient client access, load balancing, and operational efficiency across unified file and object storage services.



| Component | Description |
|---|--|
| FI-A / FI-B | Fabric Interconnects A and B (if used in UCS-managed environments) or top-of-rack switches that connect FlashBlade to the core network. They aggregate uplinks from FlashBlade's blade interfaces. |
| LAG1 / LAG2 | Link Aggregation Groups that combine multiple 10/25/40/100 GbE connections from different blades into a single logical uplink, enhancing throughput and fault tolerance. |
| Blade Interfaces (e.g., CH0.eth1, CH1.eth1) | Physical Ethernet ports on each FlashBlade chassis slot (CHx). These interfaces directly participate in client access and are often connected redundantly across fabrics. |
| Virtual Interface (VIP) | A dynamic IP address used for client connections. The VIP floats across healthy blade interfaces, ensuring continuous availability during hardware or network path failures. |
| Blades (Storage Nodes) | Each blade is an independent compute and storage node, forming a scale-out cluster. All blades contribute to I/O, metadata, and performance. |
| FlashBlade File Services | Native file service layer on FlashBlade supporting NFS and SMB protocols. It provides high-throughput, low-latency file access with built-in resilience and parallelism. |

FIGURE 2 Key infrastructure components in a direct-attached FlashBlade deployment

To deliver high availability, scalability, and performance for NFS, SMB, and S3 Object-compatible storage services, Pure Storage FlashBlade is deployed in direct-attach mode with Cisco UCS Fabric Interconnects (see Figure 3). This streamlined architecture reduces complexity while maintaining enterprise-grade resiliency and throughput for file and object workloads.

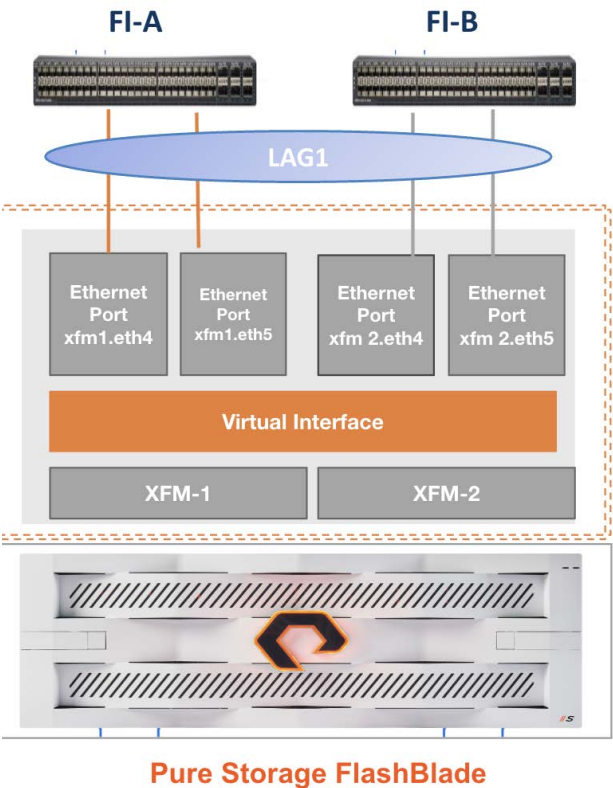


FIGURE 3 FlashBlade deployment with Cisco UCS Fabric Interconnects



Physical Port Connectivity

The Pure Storage FlashBlade//S200 is architected for performance and high availability through dual internal controllers—external fabric modules (XFM)s—each supporting multiple Ethernet interfaces:

- **XFM-1:** connected to Fabric Interconnect A (FI-A) via xfm1.eth4 and xfm1.eth5
- **XFM-2:** connected to Fabric Interconnect 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, xfm2.eth4, and xfm2.eth5 and spans FI-A and FI-B. 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.



Pure Storage FlashBlade Network Configuration Details

This section outlines the configuration of Pure Storage FlashBlade network connectivity to support high-performance file data services.

Use Case: File Services over Direct-attached Ethernet (NFS/SMB)

The Pure Storage FlashBlade is deployed in a direct-attach topology using high-throughput Ethernet links to Cisco UCS Fabric Interconnects (FI-A and FI-B). This configuration is optimized for enterprise-grade NFS and SMB workloads, offering:

- High availability through dual-path connectivity
- Enhanced bandwidth using LAGs
- Seamless client access via a virtual interface

The following subsections detail the physical port layout, LAG/LACP setup, and interface abstraction required for reliable and scalable file services deployment with FlashBlade.

Object Services over Direct-attached Ethernet (S3 Object)

The Pure Storage FlashBlade also offers robust S3 Object storage capabilities that can leverage the same direct-attach topology to Cisco UCS Fabric Interconnects (FI-A and FI-B). This configuration ensures high performance and highly available access for S3 workloads, which benefit from:

- High availability through the redundant Ethernet links
- Scalable bandwidth via LAGs
- Simplified access for S3 clients

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 UCS Fabric Interconnects. For direct connectivity from the Pure Storage FlashBlade to the UCS Fabric Interconnects, we will use subnets with VLAN interfaces that match the VLAN IDs we have defined for our environment.

We will create a subnet with a VLAN interface for each data path (A and B) and will then attach subinterfaces from each of our physical Ethernet interfaces 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 we have defined for our environment
 - **Gateway:** the gateway for your network subnet; not required in our 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**

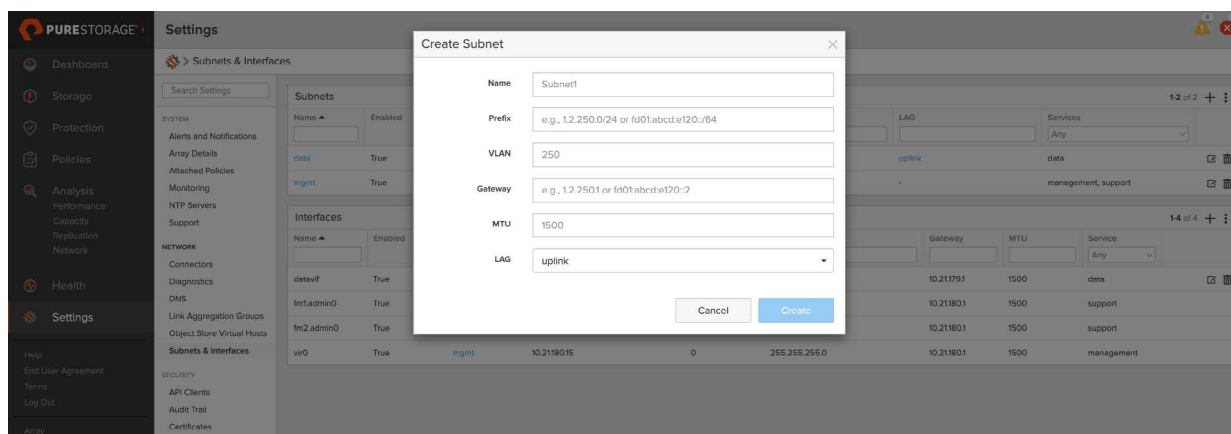


FIGURE 4 Create subnet

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 UCS Fabric Interconnects 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 UCS Fabric Interconnects.

Configuration Guide: Cisco Intersight Managed Mode

To enable direct-attach connectivity between Cisco UCS Fabric Interconnects and the Pure Storage FlashBlade in IMM, specific configurations must be performed to correctly classify the ports and ensure efficient and reliable data access.

1. **Log in to Cisco Intersight console:** Access the [Cisco Intersight console](#) using your credentials. Ensure you have the appropriate permissions to configure policies and manage Fabric Interconnects.
2. **Navigate to the Policies section:** In the left-hand navigation panel, go to **Policies > Create Policy**. Here, you'll define the connectivity behavior for the Fabric Interconnects.

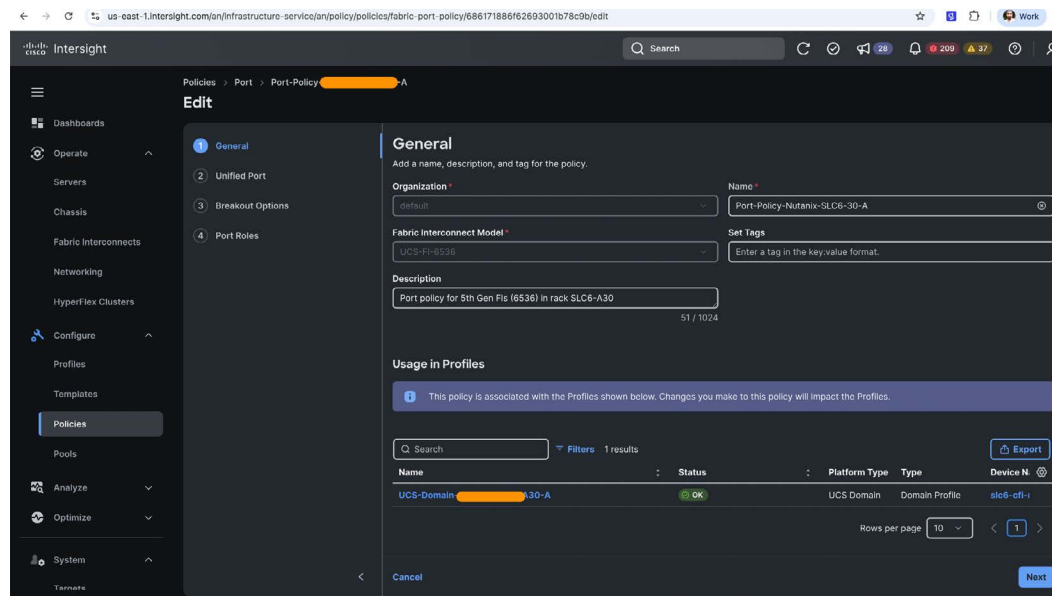


FIGURE 5 Navigate to Policies section of Cisco Intersight console

3. **Create or edit a port policy:** Under the Port Policies section, either create a new port policy or edit an existing one that is associated with your UCS domain profile. This policy will define the role and configuration of each port on the Fabric Interconnect.
4. **Configure appliance ports on the Fabric Interconnect:** Within the port policy, locate the ports on the Fabric Interconnects (typically on the uplink side) that you intend to use for connecting directly to the FlashBlade. These ports must be configured as appliance ports. To do this, select one of the relevant ports, right-click on it, and choose **Configure as Appliance Port**. Repeat for each port.



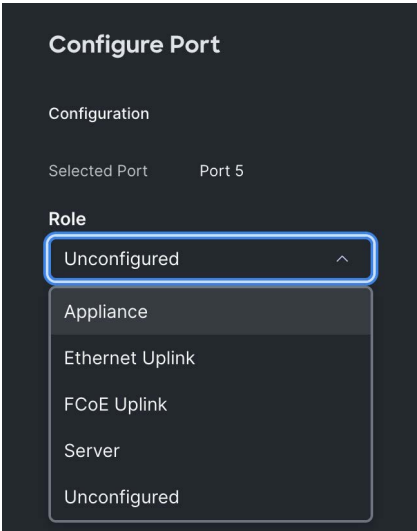
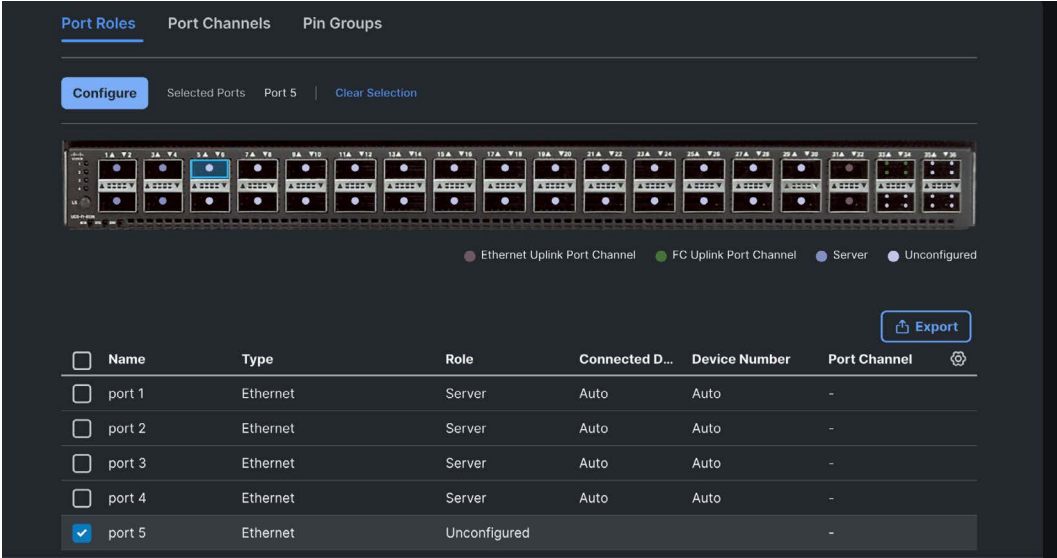


FIGURE 6 Configure appliance ports of Fabric Interconnect

Configuration Guide: Cisco UCS Manager

Because this white paper focuses on running FlashStack in direct-attach mode, there are some portions of the configuration of a UCSM environment that are not covered (see Cisco Validated Designs and [UCS documentation](#) for additional configuration details).

The UCSM configuration policies not covered in this white paper are listed in Table 2, broken down by UCSM navigation tabs.

| Configuration Tab | Policies Not Covered in this Guide |
|-------------------|--|
| Admin | Fault policies, user management, key management, communication management |
| Equipment | Firmware management, equipment policies |
| Server | Various policies—adapter, BIOS, host firmware, IPMI/Redfish access, KVM management, maintenance, power control, Serial over LAN, Server Pool, iSCSI authentication, vMedia |
| LAN | Various policies—dynamic vNIC connection, LACP, Multicast, QoS, VMQ connection |
| Storage | Storage policy |
| Chassis | Chassis maintenance policy |

TABLE 2 All configuration policies not covered in this white paper, listed by UCSM tab

This white paper does cover the other requisite policies and templates to be configured for the deployment of FlashStack in direct-attach mode, using running ESXi hosts as an example because it is a common deployment for customer environments.

UCS Storage Port Configuration

Before we can create our configuration components within the UCS environment, we must enable the physical interfaces within the UCSM environment to provide connectivity for the data paths between the Pure Storage FlashBlade and UCS Fabric Interconnects. The following steps walk through configuring the server, network uplink, and storage appliance ports.

For the configuration of Fabric Interconnects, Figure 7 shows how the filtered view of the Fabric Interconnects under the Equipment tab will appear.



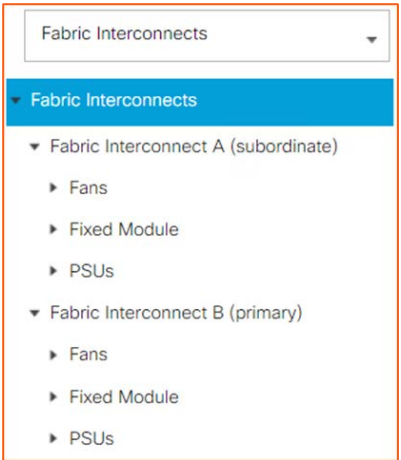


FIGURE 7 Filtering the view to Fabric Interconnects within the Equipment tab

First, the ports on the Fabric Interconnects need to be assigned as appliance ports. Select a port that needs to be assigned, right-click on it (**Port 33** is selected in Figure 8), and choose **Configure as Appliance Port**. Repeat for the remaining ports.

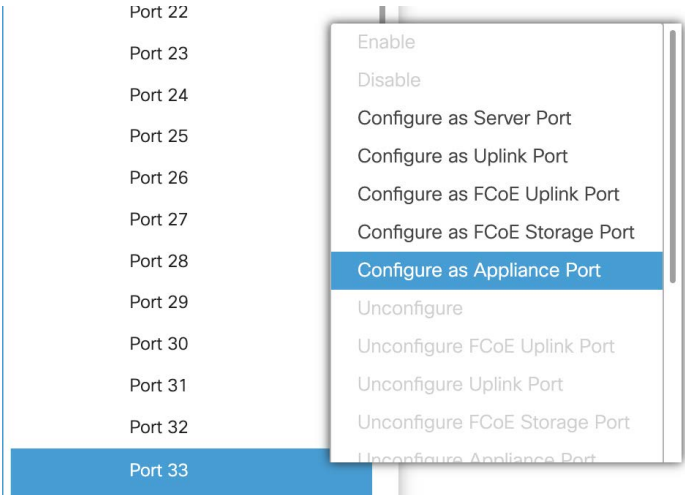


FIGURE 8 Configure ports as appliance ports

Unlike a server port, the appliance port requires some additional configuration. Appliance ports use their own special appliances VLAN cloud, which is separate from the standard LAN cloud. You can either make the clouds as you configure the ports or prepare them ahead of time.

Create individual port channel for each FI in Cisco UCS

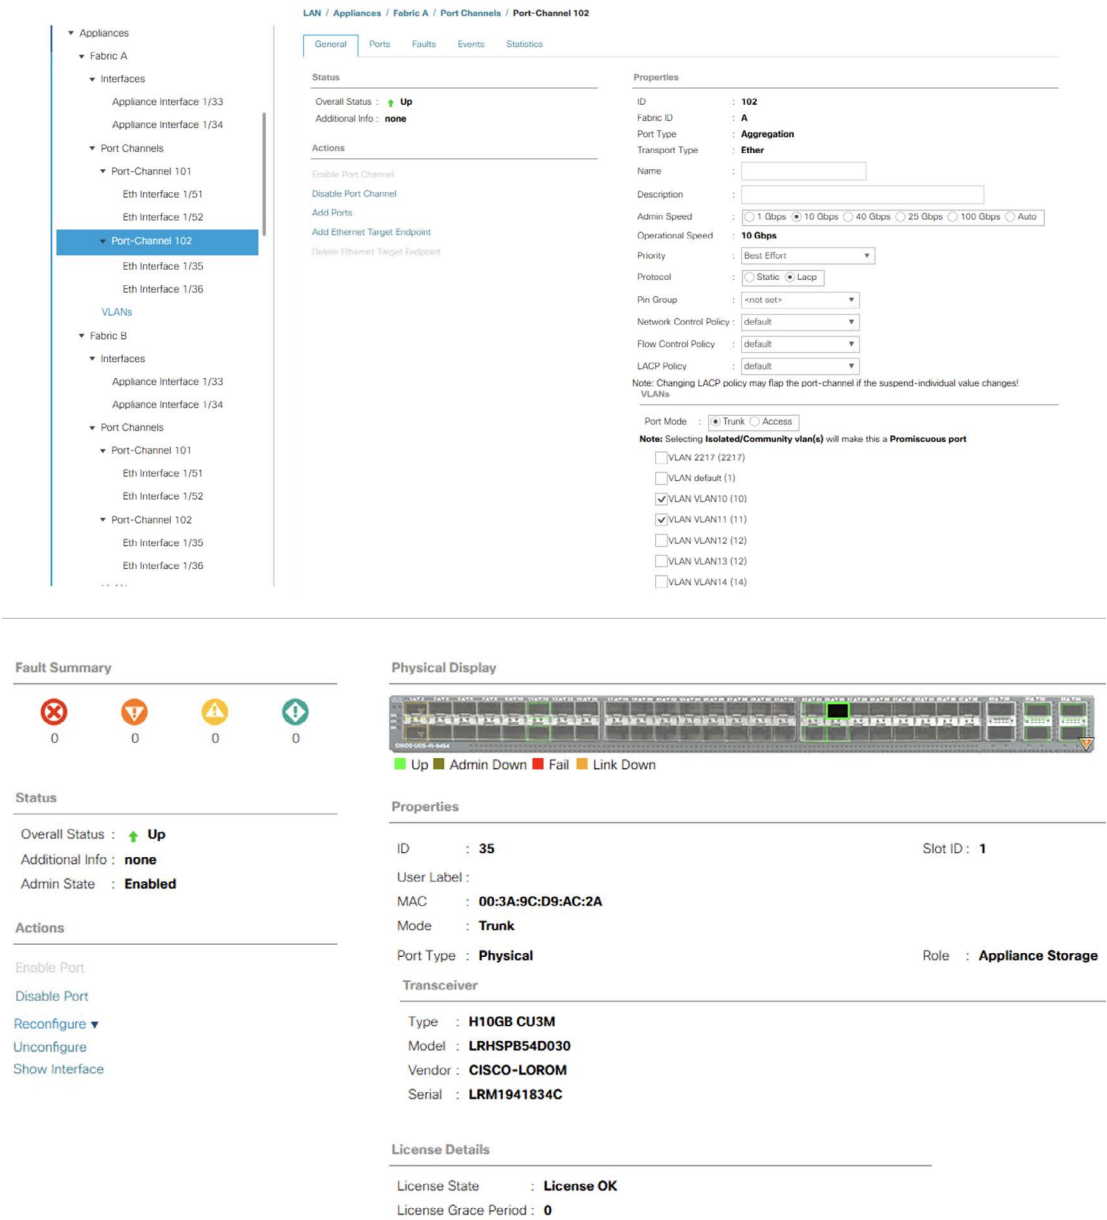


FIGURE 9 Create appliances VLAN cloud

Assigning VLANs for FlashBlade Storage Traffic on Cisco UCS

1. Log in to UCS Manager.
2. Open the UCS Manager graphical user interface (GUI)—typically via the cluster IP of the Fabric Interconnects.
3. Create VLANs for different storage protocols:
 - Navigate to **LAN > VLANs**.
 - Create separate VLANs for each FlashBlade protocol. For example:
 - VLAN 100 – NFS
 - VLAN 101 – SMB
 - VLAN 102 – Object Storage (S3)
 - Ensure that these VLANs are non-native and not configured as default VLANs.
4. Assign VLANs to the appliance port:
 - Navigate to **LAN > Appliances**.
 - Locate the previously configured appliance port (for example, Eth1/10, Eth1/11).
 - Select the appliance port.
 - In the VLANs tab, click **Add VLAN**.
 - Associate VLAN 100, 101, and 102 to the selected port.
 - Set the native VLAN correctly if needed for fallback traffic.

Note: Appliance ports bypass UCS server management and are intended for external systems (for example, network-attached storage or S3 storage). These ports do not participate in UCS pinning or server vNIC policies.

Enable Jumbo Frames for Performance (Optional)

1. Navigate to **LAN > Policies > MTU Policies**.
2. Create a new MTU policy (for example, JumboFrames-9000)
3. Set **MTU size** to **9000 bytes**.
4. Apply the MTU policy to the VLANs.
5. Navigate back to **LAN > VLANs**.
6. For each VLAN (100, 101, 102), edit the VLAN configuration. Under MTU Policy, select and apply **JumboFrames-9000**.



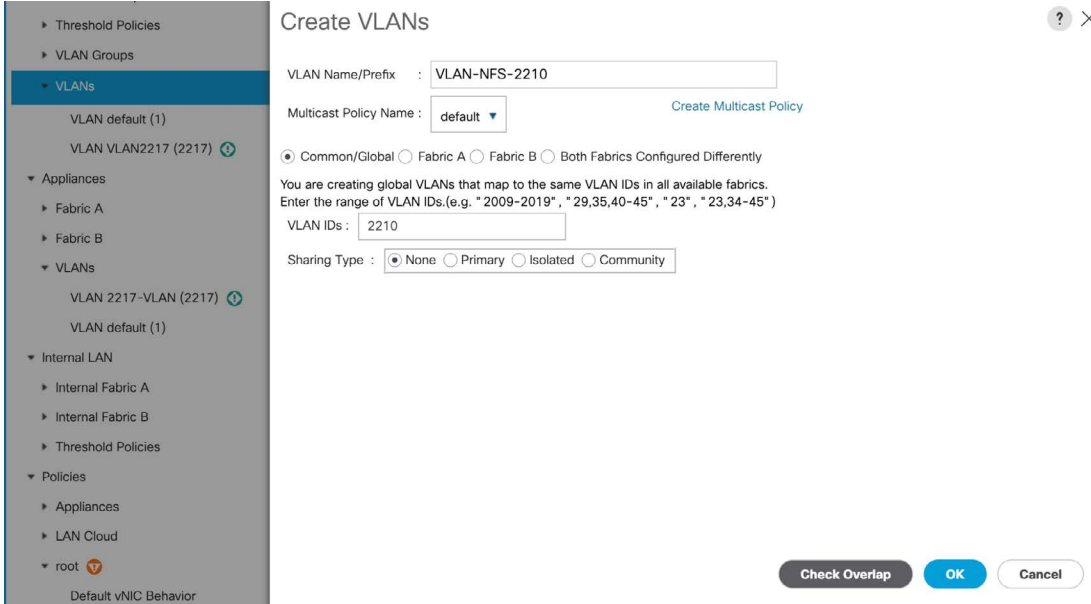


FIGURE 10 Enable jumbo frames (optional)

Appliance Interface Configuration: Ethernet

The same steps used previously will be followed to create an appliance interface that will allow for direct-attach connectivity between the FlashBlade and UCS Fabric Interconnects.

1. **Navigate to the Equipment section:** In the Cisco UCS Manager, click the **Equipment** tab in the left navigation pane.
2. **Filter to Fabric Interconnects:** Select the **Fabric Interconnects** filter at the top of the navigation pane to view only the Fabric Interconnects configuration.
3. **Locate Ethernet ports:** Once Fabric A and Fabric B are visible, expand either one to see the fixed module. Then, expand the fixed module to view the Ethernet ports.
4. **Configure appliance port:** With **Ethernet Ports** selected, locate the required port number for your environment. Right-click on the desired port and click **Configure as Appliance Port**.
5. **Confirm appliance port selection:** In the Configure as Appliance Port notification, ensure the correct port number is shown. Click **Yes** to proceed.
6. **Enter appliance port configuration details:** In the pop-up window, enter the following:
 - **Priority:** Select **Platinum** (QoS priority).
 - **Pin Group:** Leave as **<not set>**.
 - **Network Control Policy:** Select the policy created earlier (for Cisco Discovery Protocol/Link Layer Discovery Protocol).
 - **Flow Control Policy:** Leave as **default**.
 - **Admin Speed (Gbps):** Set to match your physical interface and transceiver specifications.
 - **VLANs:**
 - Set port mode to **Trunk**.
 - Select VLANs, including the dummy trunk native VLAN. The iSCSI VLAN is defined for your environment.
 - **Native VLAN:** Set to the dummy trunk native VLAN.
 - **Ethernet Target Endpoint:** Enter the name and MAC address of the FlashBlade interface this UCS Fabric Interconnect port connects to.



7. **Repeat steps for all required ports:** Repeat steps 4–6 for all necessary physical interfaces on Fabric Interconnect A. Then, repeat the same process for Fabric Interconnect B.
8. **Finalize configuration:** Click **OK** to close the NAS Appliance Manager pop-up window and save the configuration.

Failover Test Summary: Ensuring High Availability of FlashBlade File Services

To validate the resiliency and high availability of the Pure Storage FlashBlade File Services environment, a series of controlled failover scenarios was executed. These tests simulate real-world events and demonstrate the ability of the FlashBlade to maintain file service continuity through its redundant architecture and automated failover mechanisms. The following types of tests were performed:

1. XFM module reboot (xfrm1/xfrm2)
2. Fabric Interconnect reboot
3. Object service validation (during failover)

The flow of these tests and their respective outcomes are detailed in the following subsections.

XFM Module Reboot (xfrm1/xfrm2)

Purpose

To verify complete hardware-level failover from one XFM module to another.

Method

The active XFM module (for example, xfrm1) was intentionally rebooted. During this time:

- File services and LAG interfaces were automatically migrated to the standby module (xfrm2).
- The virtual interface was reassigned to the active module.

Expected Behavior

- VIP migrated automatically to the active XFM.
- Clients experienced minimal to no disruption.
- LACP ensured continued operation through the healthy paths.

Fabric Interconnect Reboot

Purpose

To test the redundancy of upstream connectivity and validate LAG/LACP resiliency between **FI-A** and **FI-B**.

Method

Either **FI-A** or **FI-B** was rebooted while the FlashBlade remained online. Each XFM module has dual Ethernet ports connected to both Fabric Interconnects. During the test:

- LACP detected the link failure.
- Interfaces tied to the failed Fabric Interconnect were removed from the LAG.
- Traffic rerouted through the alternate Fabric Interconnect.



Expected Behavior

- **Brief packet drops** were observed **before connectivity was reestablished**, which is expected.
- This behavior is due to the **Fabric Interconnect not functioning as a traditional switch** (it doesn't form a virtual port channel like Nexus switches).
- The system automatically reconverged and resumed normal operations using the surviving path.

This validates upstream resilience for environments using **direct-attached Ethernet** to Cisco UCS Fabric Interconnects, such as in FlashStack architectures.

Object Service Validation (During Failover)

Purpose

Specifically validate S3 Object storage accessibility and service continuity during hardware/network failover events.

Method

- Ran continuous S3 PUT/GET operations from a client system during XFM module reboot and Fabric Interconnect reboot
- Monitored connection stability and response times

Expected Behavior

- No S3 session termination observed
- Ongoing PUT/GET operations resumed successfully after brief network path failover
- The FlashBlade distributed service architecture ensured that the S3 service layer remained abstracted from the underlying hardware disruption

| Test Cases | Purpose | Method | Expected Behavior |
|---|---|--|--|
| • 1 • XFM Module Reboot (xfm1 / xfm2) | • To verify complete hardware-level failover from one XFM module to another. | <ul style="list-style-type: none"> • Intentionally rebooted the active XFM module (e.g., xfm1). - Observed automatic migration of file services and LAG interfaces to standby (xfm2). - VIP reassigned to active module. | <ul style="list-style-type: none"> • - VIP migrated automatically to the active XFM. - Minimal to no client disruption. - LACP maintained traffic via healthy paths. |
| • 2 • Fabric Interconnect (FI) Reboot | • To test redundancy of upstream connectivity and validate LAG/LACP resiliency between FI-A and FI-B. | <ul style="list-style-type: none"> • - Rebooted either FI-A or FI-B while FlashBlade remained online. - Each XFM module had dual connections to both FIs. - Observed LACP behavior and traffic rerouting. | <ul style="list-style-type: none"> • - Brief packet drops occurred, as expected. - Interfaces tied to failed FI removed from LAG. - Traffic rerouted via alternate FI. - System auto-reconverged and resumed normal operations. |
| • 3 • Object Store Validation During Failover | • To validate the high availability and uninterrupted access to S3 object services during component-level failover scenarios. | <ul style="list-style-type: none"> • Performed continuous S3 operations (PUT/GET) from client systems during both: <ul style="list-style-type: none"> - XFM module reboot - Fabric Interconnect reboot - Verified response time and accessibility throughout the process. | <ul style="list-style-type: none"> • Ongoing S3 operations continued with minimal to no disruption. - No session timeouts or failures observed. - S3 endpoints remained available during XFM and network failovers. - FlashBlade's distributed architecture maintained service continuity. |

FIGURE 11 Test outcomes



Conclusion

Direct connectivity of NFS, SMB, and S3 Object storage to Cisco UCS Fabric Interconnects provides a simplified and high-performance storage solution. By understanding the limitations and best practices outlined in this white paper, IT teams can design a scalable and efficient direct-attached storage architecture that meets their workload demands.

With the capability to do this in a reduced footprint for both power and cost, the FlashBlade solution truly becomes a powerful platform, delivering more potential to customers from the smallest to largest scale.

Additional Resources

- Visit the Pure Storage FlashBlade product page: [Pure Storage FlashBlade//S](#).
- Review the [Pure Storage FlashStack Compatibility Matrix](#). This interoperability list requires a support login from Pure Storage.