



Cisco ACI and OpFlex Connectivity for Orchestrators

New and Changed Information	2
OpFlex Connectivity for Orchestrators	2
Hardware Requirements	2
OpFlex Connectivity for Blade Systems	2
OpFlex Connectivity for Rack-mount Servers	3
OpFlex Connectivity for Nested Blade Systems	4
OpFlex Connectivity for Nested Rack-mount Servers	5

Revised: June 21, 2022

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Cisco APIC Release Version	Feature	Description
N/A	N/A	Added Hardware Requirements section.
3.1(2)	--	This document was made available to describe network connectivity models supported for the Cisco ACI integration of Cloud Foundry, Kubernetes, OpenShift, and OpenStack.

OpFlex Connectivity for Orchestrators

Cisco ACI supports integration with the following orchestrators through the OpFlex protocol:

- Cloud Foundry
- Kubernetes
- OpenShift
- OpenStack

This section describes the network connectivity models supported for the Cisco ACI integration of these orchestrators.

Hardware Requirements

This section provides the hardware requirements:

- Connecting the servers to Gen1 hardware or Cisco Fabric Extenders (FEXes) is not supported and results in a nonworking cluster.
- The use of symmetric policy-based routing (PBR) feature for load balancing external services requires the use of Cisco Nexus 9300-EX or -FX leaf switches.

For this reason, the Cisco ACI CNI Plug-in is only supported for clusters that are connected to switches of those models.



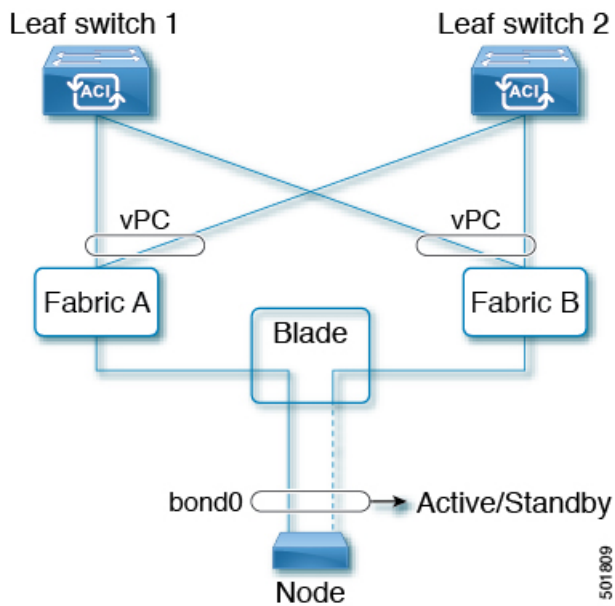
Note UCS-B is supported as long as the UCS Fabric Interconnects are connected to Cisco Nexus 9300-EX or -FX leaf switches.

OpFlex Connectivity for Blade Systems

Follow the guidelines and illustration in this section when integrating blade systems into the Cisco ACI fabric.

- Configure port channels on interfaces in Fabric A and Fabric B for uplink connectivity to Cisco ACI leaf switches. Each fabric connects to both leaves in vPC.
- On the blade server, configure two vnics, one pinned to Fabric A and one to Fabric B.
- On the Linux server, configure the vnics with bonding active/standby.

Figure 1: OpFlex Connectivity for Blade Systems

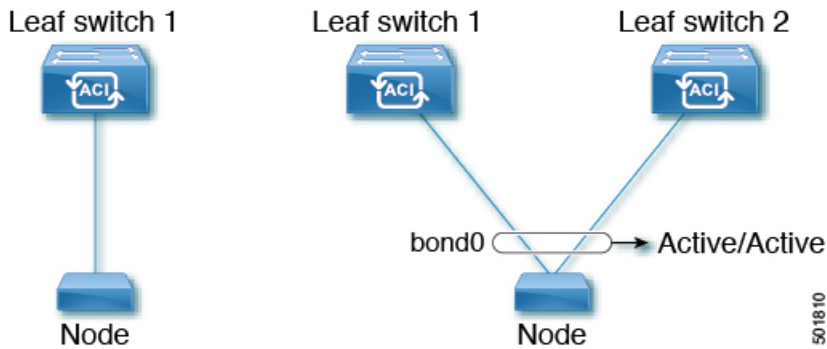


OpFlex Connectivity for Rack-mount Servers

Follow the guidelines and illustration in this section when integrating rack-mount servers into the Cisco ACI fabric.

- For a host connected to a single leaf:
 - Do not make the leaf part of a vPC pair.
 - Configure a single uplink; port channel is not supported.
- For a host connected to two leaf switches:
 - Configure vPC with LACP Active/ON.
 - Configure active/active Linux bonding.

Figure 2: OpFlex Connectivity for Rack-mount Servers



OpFlex Connectivity for Nested Blade Systems

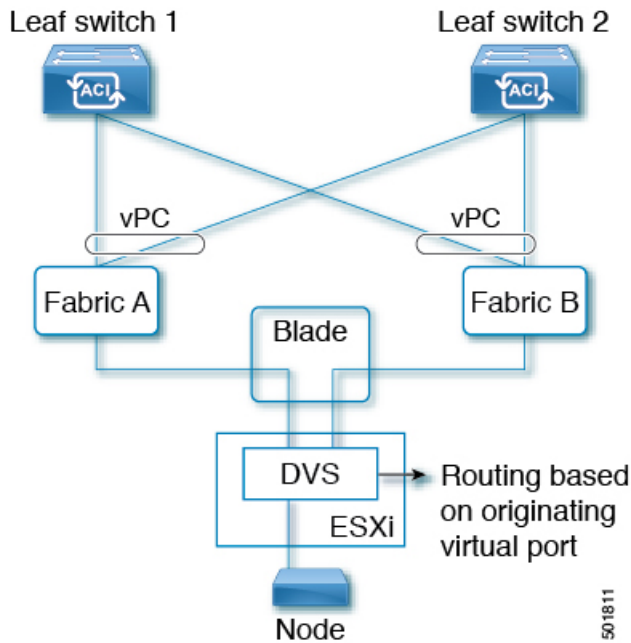
Follow the guidelines and illustration in this section when integrating nested blade systems into the Cisco ACI fabric.

- Configure port channels on interfaces in Fabric A and Fabric B for uplink connectivity to Cisco ACI leaf switches. Each fabric connects to both leaves in a vPC.
- On the blade server, configure two vnics, one pinned to Fabric A and one pinned to Fabric B.
- Route the VDS based on the originating virtual port.



Note We recommended that when using ESXi nested for Kubernetes hosts, you provision one Kubernetes host for each Kubernetes cluster for each ESXi server. Doing so ensures—in case of an ESXi host failure—that only a single Kubernetes node is affected for each Kubernetes cluster.

Figure 3: OpFlex Connectivity for Nested Blade Systems



OpFlex Connectivity for Nested Rack-mount Servers

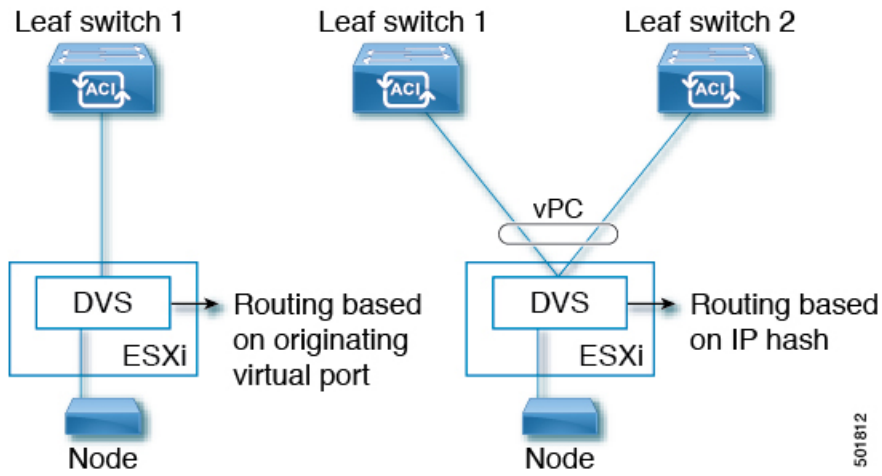
Follow the guidelines and illustration in this section when integrating nested rack-mount systems into the Cisco ACI fabric.

- For a host connected to a single leaf switch:
 - Do not make the leaf part of a vPC pair.
 - Configure a single uplink; port channel is not supported.
 - Route the VDS based on IP hashing.
- For a host connected to two leaf switches:
 - Configure the vPC with LACP Active/ON
 - Route the VDS based on IP hashing.



Note We recommend that when using ESXi nested for Kubernetes hosts, you provision one Kubernetes host for each Kubernetes cluster for each ESXi server. Doing so ensures—in case of an ESXi host failure—that only a single Kubernetes node is affected for each Kubernetes cluster.

Figure 4: OpFlex Connectivity for Nested Rack-mount Servers



501812

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2018-2020 Cisco Systems, Inc. All rights reserved.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

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.