Cisco Nexus 1000V VDP Configuration Guide, Release 4.2(1)SV2(2.2)

First Published: January 29, 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-31102-01

CONTENTS

Preface

This preface contains the following sections:

Audience

This publication is for experienced network administrators who configure and maintain Cisco Nexus devices.

This guide is for network and server administrators with the following experience and knowledge:

Note Knowledge of VMware vNetwork Distributed Switch is not required.

- An understanding of virtualization
- An understanding of the corresponding hypervisor management software for your switch, such as VMware vSwitch, Microsoft System Center Virtual Machine Manager (SCVMM), or OpenStack

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.

Convention	Description
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation for Nexus 1000V Series NX-OS Software for VMware vSphere

This section lists the documents used with the Cisco Nexus 1000V and available on Cisco.com at the following URL:

http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html

General Information

Cisco Nexus 1000V Documentation Roadmap

Cisco Nexus 1000V Release Notes

Cisco Nexus 1000V and VMware Compatibility Information

Install and Upgrade

Cisco Nexus 1000V Installation and Upgrade Guide

Configuration Guides

Cisco Nexus 1000V High Availability and Redundancy Configuration Guide

Cisco Nexus 1000V Interface Configuration Guide

Cisco Nexus 1000V Layer 2 Switching Configuration Guide

Cisco Nexus 1000V License Configuration Guide

Cisco Nexus 1000V Network Segmentation Manager Configuration Guide

Cisco Nexus 1000V Port Profile Configuration Guide

Cisco Nexus 1000V Quality of Service Configuration Guide

Cisco Nexus 1000V REST API Plug-In Configuration Guide

Cisco Nexus 1000V Security Configuration Guide

Cisco Nexus 1000V System Management Configuration Guide

Cisco Nexus 1000V vCenter Plugin Configuration Guide

Cisco Nexus 1000V VXLAN Configuration Guide

Cisco Nexus 1000V VDP Configuration Guide

Cisco Nexus 1000V DFA Configuration Guide

Programming Guide

Cisco Nexus 1000V XML API Configuration Guide

Reference Guides

Cisco Nexus 1000V Command Reference

Cisco Nexus 1000V Resource Availability Reference

Troubleshooting and Alerts

Cisco Nexus 1000V Troubleshooting Guide

Cisco Nexus 1000V Password Recovery Procedure

Cisco NX-OS System Messages Reference

Cloud Services Platform Documentation

The *Cisco Cloud Services Platform* documentation is available at http://www.cisco.com/en/US/products/ps12752/tsd_products_support_series_home.html.

Virtual Security Gateway Documentation

The *Cisco Virtual Security Gateway for Nexus 1000V Series Switch* documentation is available at http://www.cisco.com/en/US/products/ps13095/tsd_products_support_series_home.html.

Virtual Wide Area Application Services (vWAAS) Documentation

The *Virtual Wide Area Application Services* documentation is available at http://www.cisco.com/en/US/products/ps6870/tsd_products_support_series_home.html.

ASA 1000V Cloud Firewall Documentation

The *ASA 1000V Cloud Firewall* documentation is available at http://www.cisco.com/en/US/products/ps12233/tsd_products_support_series_home.html.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to one of the following:

- nexus1k-docfeedback@cisco.com

We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation*, at: http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

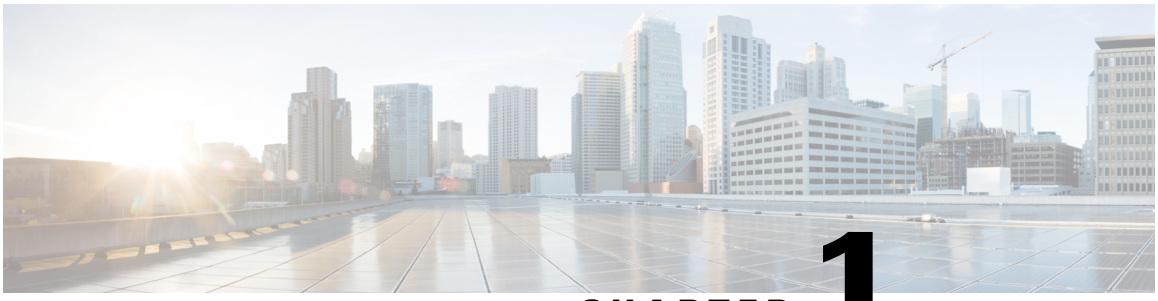

CHAPTER 1

Overview

This chapter contains the following sections :

Information about the VSI Discovery and Configuration Protocol

The Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) on the Cisco Nexus 1000V is part of the IEEE standard 802.1Qbg (Edge Virtual Bridging - [EVB]) that can detect and signal the presence of end hosts and exchange capability with an adjacent VDP-capable bridge. The VDP serves as a reliable first-hop protocol that communicates the presence of end-host Virtual Machines (VMs) to adjacent leaf nodes on the Cisco Dynamic Fabric Automation (DFA) architecture. In addition to detecting the MAC and IP addresses of the end-host VMs when a host comes up, or during VM mobility events, VDP also triggers auto-configuration of leaf nodes on the DFA architecture to make them ready for further VM traffic.

VDP enables network-based overlays that are a more scalable alternative compared to the host-based overlays for segmentation and enables access to more than 4000 vlans in a multi tenant network. When you configure VDP on the Cisco Nexus 1000V, segmentation support for bridge domains is extended to native encapsulated bridge domains. The original Virtual Extensible Local Area Network (VXLAN) based bridge domains can also coexist with these bridge domains.

For more information about the Cisco DFA architecture, see the *Cisco DFA Solutions Guide*.

Features of VDP

The VSI Discovery Protocol (VDP) provides the following features:

- VDP facilitates end-to-end segmentation enabled in the Cisco Dynamic Fabric Automation (DFA) architecture and removes the disadvantages of the host-based overlays.
- Serves as an end-host registration protocol for the Cisco DFA leaf switches that use the registration information to automatically configure the network information.
- Uses Edge Control Protocol (ECP , also part of the IEEE 802.1Qbg standard) as the transport protocol for the protocol data units (PDUs).

- Facilitates constant migration of a VM and its network state by enabling the association and de-association of VSI types and VSI instances.
- Enables segmentation through native encapsulation and other Cisco DFA-based configuration between the leaf nodes and the Cisco Nexus 1000V Virtual Ethernet Modules (VEMs).
- Defines message exchanges between the following communicating entities:
  - VDP station : End system that initiates the VDP exchange to signal the presence of a VM and the needed connection. This station could be a vSwitch on the hypervisor that runs in a physical server that supports the deployment of one or more VMs.
  - VDP bridge : Edge bridge that directly attaches to the VDP station. A VDP bridge can have multiple ports that face different VDP stations, where each port forms an independent VDP communication between its corresponding stations.

VDP Components in the Cisco Dynamic Fabric Automation Network

The VSI Discovery and Configuration Protocol (VDP) in the Cisco DFA network runs on the leaf switches and the Cisco Nexus 1000V (end stations) as shown in the following figure:

Figure 1: Components of VDP in the Cisco Dynamic Fabric Automation Network

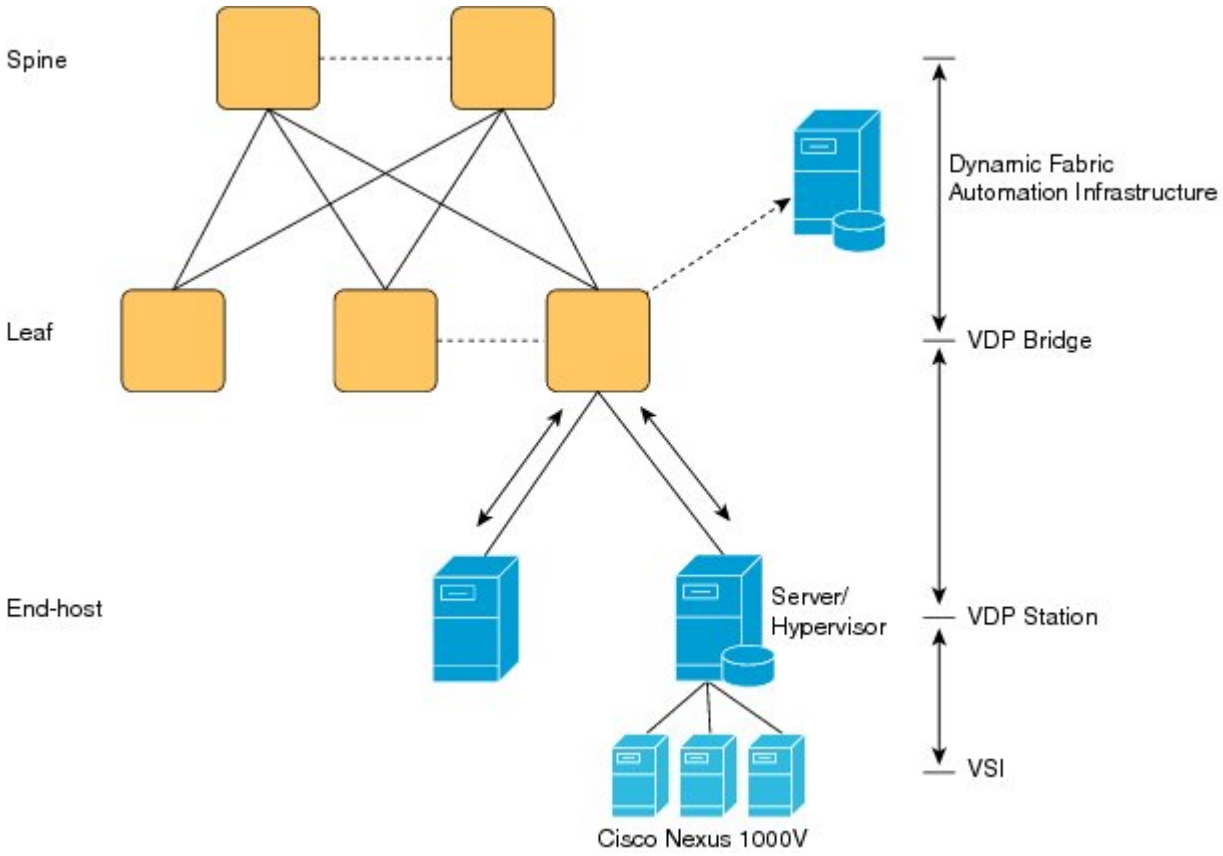

The components and functioning of the VDP Exchange in the Cisco Dynamic Fabric Automation architecture are described below:

- Leaf Switch : A DFA leaf node operates as the bridge for the VSI Discovery and Configuration Protocol (VDP) exchange that handles requests from end hosts. The leaf node also communicates with the configuration profile databases to retrieve and apply the previously defined port profiles to each attached end host.
- End-Station : An end station in Cisco DFA can be VDP capable or incapable. A VDP capable end station operates as the primary station for the VDP exchange and registers or deregisters its resident VMs to the attached leaf switch. A VDP-incapable end-station is a normal server node that does not participate in the VDP message exchange. The VEM on the Cisco Nexus 1000V acts as an end station in the Cisco DFA and the VDP implementation on the Cisco Nexus 1000V is called the station side VDP.
- Profile database : This database is a standalone server or a local configuration storage in the leaf-switch that maps each end-host to its predefined port profile. This profile can be VLAN, ACL or QoS settings.

VDP Sequence

When an end host (VM) is instantiated , the Cisco Nexus 1000V on the VDP station (host server) registers its presence with the VDP bridge and passes the network information to the Cisco DFA leaf switch using VDP. The DFA leaf switch then retrieves and applies the corresponding port profile to the end host to provide an automatic provisioning mechanism for reachability and network control.

The VDP implementation on the Cisco Nexus 1000V (station side VDP) uses the following sequence to facilitate a VDP exchange:

1. When a VM is activated, VDP passes the network information to the Cisco DFA leaf switch through a VDP request. The network information for a VM is carried in the form of TLVs (Type Length Values) that are exchanged between the station (Cisco Nexus 1000V) and the leaf. The TLVs consists of filter formats that indicate the network information parameters for a VM. The Cisco Nexus 1000V passes the IP addresses and VM names to the Cisco DFA leaf switch, using a Cisco OUI TLV.
2. After receiving the request, VDP on the leaf extracts the network information and automatically configures and attaches a VLAN value to the segment ID.
3. VDP on the leaf switch sends a response to the Cisco Nexus 1000V after the TLV's filters are modified to the new VLAN. The Cisco Nexus 1000V applies the VLAN in the dot1q encapsulation of packets for that VM.
4. After a VM is successfully associated, VDP on the station periodically sends the network information to the leaf switch for a state refresh. If there is a failure on the leaf switch or if the leaf switch becomes unresponsive, the station retries to send the request after a configurable interval.

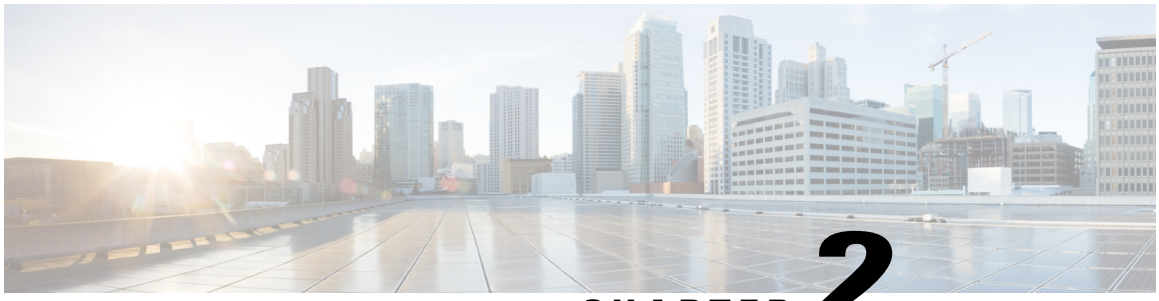

CHAPTER 2

Configuring VDP

This chapter contains the following sections:

Information about VDP for Blade-Chassis Deployment

VDP on a Cisco DFA network architecture runs the Edge Control Protocol (ECP) to forward packets upstream to the DFA leaf switch. ECP is a Layer 2 protocol that uses the nearest bridge MAC address 01:80:C2:00:00:01 as the destination MAC address to forward data traffic. In a blade-chassis deployment, blade switches such as the Cisco UCS Fabric Interconnect (UCS FI) that interface with the Cisco Nexus 1000V VEMs and the leaf switches terminate packets with the specified MAC address, because they are the same packets used for the bridge protocol data unit (BPDU) frames. Consequently, the VDP exchange between the Cisco Nexus 1000V VEMs and upstream leafs fail. To enable the VDP packets to get transported upstream to the DFA

leaf, the destination MAC address for the ECP packets must be allowed to pass through the blade switches to forward the packets upstream to the Cisco DFA leaf.

To enable VDP communication and to avoid changes in the functioning of the blade switches such as the Cisco UCS fabric interconnect and their compatibility with other network devices, you can configure the destination MAC address that originates from the VDP station. See Configuring a DMAC from the VDP Station, on page 12.

For more information about blade-chassis deployment, see http://www.cisco.com/en/US/products/ps10279/index.html

Unsupported Topology

In this release, VDP on the Cisco Nexus 1000Vdoes not support an un-clustered topology where an upstream Leaf or bridge nodes are not configured as a VPC/VPC+pair, independent of the devices such as the UCS Fabric InterConnect (UCS FI) interfacing between the Cisco Nexus 1000V VEM and the Cisco DFA Leaf.

Note

1. VDP is supported only on the Cisco Nexus 6000 Series switches in release 4.2(1)SV2(2.2).
2. VDP supports connectivity to multiple bridges that are clustered to one bridge through a virtual port channel (vPC).

Prerequisites

Configuring VDP for the Cisco Nexus 1000V has the following prerequisites:

- You have installed and configured the Cisco Nexus 1000V for VMware vSphere software using the *Cisco Nexus 1000V Installation and Upgrade Guide*.
- Ensure that the Virtual Supervisor Module (VSM) has an active SVS connection.
- Ensure that the Virtual Supervisor Module (VSM) and Virtual Ethernet Module (VEM) connectivity is functioning.
- You have added hosts to the Cisco Nexus 1000V.
- You have disabled the segmentation feature.

Guidelines and Limitations

Implementing VDP on the Cisco Nexus 1000V has the following guidelines and limitations:

- The Cisco Nexus 1000V supports the Cisco DFA capable VDP based on the IEEE Standard 802.1 Qbg, Draft 2.2, and does not support the Link Layer Discovery Protocol (LLDP). Therefore, the EVB TLVs will not be originated or processed by the Cisco Nexus 1000V.
- The VDP implementation in the current release supports a matching LLDP-less implementation on the bridge side, which is delivered as part of the Cisco DFA solution. For more information on the Cisco DFA, see *Cisco DFA Solutions Guide.*

- Timer-related parameters are individually configurable in the station and in the leaf.
- Connectivity to multiple unclustered bridges is not supported in this release. For more information about unsupported topologies, see Unsupported Topology , on page 6
- IPv6 addresses in filter format are not supported in this release.
- VDP is supported for only segmentation based port-profiles. VDP for VLAN based port-profiles is not supported in this release.
- The dynamic VLANs allocated by VDP are local to the VEM, and they should not be configured on the Cisco Nexus 1000V VSM.
- VDP is supported on VMware ESX releases 5.0. 5.1, and 5.5 in the current release.

Default Settings

The following table lists the default settings for VDP parameters:

Parameter	Default
Feature Segmentation	Disabled

Configuring VDP

This section includes the following topics:

Enabling Edge Virtual Bridging

Edge Virtual Bridging (EVB) is an IEEE 802.1Qbg standard that enables coordinated configuration and management of bridge services for virtual stations in a network. VDP is a part of the EVB standard that is used to detect the presence of end hosts and exchange VDP capability with an adjacent VDP bridge. For more information about the EVB Standard, see Standards, on page 15.

To configure VDP on the Cisco Nexus 1000V, you must enable the EVB feature.

Before You Begin

- You have installed and configured the Cisco Nexus 1000V for VMware vSphere software using the *Cisco Nexus 1000V Installation and Upgrade Guide*.

- Ensure that the Virtual Supervisor Module (VSM) and Virtual Ethernet Module (VEM) connectivity is functioning.
- Log in to the CLI in EXEC mode.

Procedure

	Command or Action	Purpose
Step 1	switch # **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature evb**	Enables EVB.
Step 3	switch(config)# **show feature**	(Optional) Displays the enabled status for the Cisco Nexus 1000V for features such as EVB.
Step 4	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

show feature

This example shows how to display the features after evb is enabled.

```
switch # show feature
Feature Name Instance State
-------------------- -------- --------
cts 1 disabled
dhcp-snooping 1 disabled
evb 1 enabled
http-server 1 enabled
lacp 1 disabled
netflow 1 disabled
network-segmentation 1 disabled
port-profile-roles 1 disabled
private-vlan 1 disabled
segmentation 1 enabled
sshServer 1 enabled
tacacs 1 disabled
telnetServer 1 enabled
vff 1 enabled
vtracker 1 disabled
vxlan-gateway 1 disabled
```

Modifying a Port Profile

You can the modify the Cisco Nexus 1000V port profile to configure the vEthernet interfaces or a port channel as VDP-capable links.

Before You Begin

- Log in to the CLI in EXEC mode.
- Configure the interface must be configured as a trunk mode interface.
- Enable the EVB feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **port-profile type ether vdp-capable-uplink**	Specify the port profile configuration mode for the VDP-capable uplink. If the port profile does not already exist, it is created using the following parameter: • *name*— Port profile name that can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V. **Note** If a port profile is configured as an Ethernet type, it cannot be used to configure VMware virtual ports.
Step 3	switch(config-port-prof)# **switchport mode trunk**	Designates that the interfaces are to be used as a trunking ports. A trunk port transmits untagged packets for the native VLAN and transmits encapsulated, tagged packets for all other VLANs.
Step 4	switch(config-port-prof)# **switchport trunk dynamic**	Designates that the interfaces are to be used as dynamic trunking ports.
Step 5	switch(config-port-prof)# **channel-group auto mode active**	(Optional) Configures the port profile for a port channel. **Note** If more than one physical uplink port or port channels inherit the port profile information from the original configuration, only one of them is chosen as the designated uplink port over which the VDP communication is enabled. The selected port functions in active mode and the other ports move to the standby mode.
Step 6	switch (config)# **show running-config port-prof vdp-capable uplink**	(Optional) Displays a list of interfaces that inherited a port profile.
Step 7	switch (config)# **show running interface port-channel**	(Optional) Displays the port channel that has inherited a port profile.
Step 8	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

show running-config port-prof

This example shows how to display the port-profile configuration of a VDP capable uplink.

```
switch # show running-config port-prof vdp-capable uplink
```

```
port-profile type ethernet uplink-vdp-capable uplink
vmware port-group
switchport mode trunk
switchport trunk allowed vlan 2-3967,4048-4093
switchport trunk dynamic
no shutdown
state enabled
```

Configuring Global Mode

At a global configuration level, you can set the transport mode to a native (VDP) state to employ the network-based overlays.

Before You Begin

- Log in to the CLI in EXEC mode.
- You have previously enabled the EVB feature.

Procedure

	Command or Action	Purpose
Step 1	switch # **configure terminal**	Enters global configuration mode.
Step 2	switch (config)# **feature segmentation**	Enables the segmentation feature.
Step 3	switch (config)# **segment transport-mode native**	Sets the default transport mode to VXLAN. Specify native to set it to VDP global configuration mode.
Step 4	switch (config)# **show running-config bridge-domain**	(Optional) Displays the segmentation configuration for all bridge domains.
Step 5	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

show run bridge-domain

This example shows how to display the segmentation configuration for all bridge domains.

```
switch # show running-config bridge-domain
bridge-domain seg22222
segment id 22222
group 239.1.1.1
segment transport-mode native
fabric forwarding mode proxy-gateway
```

Configuring a VDP Segment Bridge Domain

The transport mode that you configure under a bridge domain always overrides the segment transport mode that you can set globally. Use this procedure to configure a VDP segment bridge domain.

Before You Begin

- Log in to the CLI in EXEC mode.

Procedure

	Command or Action	Purpose
Step 1	switch # **configure terminal**	Enters global configuration mode.
Step 2	switch (config)# **bridge-domain** *name-string*	Creates a bridge domain and associates an identifying name to it.
Step 3	switch (config-bd)# **segment id** *number*	Specifies the bridge domain segment ID. Only one bridge domain can use a particular segment ID value. Valid values are from 4096 to 16000000. (1 to 4095 are reserved for VLANs.)
Step 4	switch (config-bd)# **group** *name*	Specifies the multicast group name for broadcasts and floods. Reserved multicast addresses are not allowed. **Note** If you enable native encapsulation, the group name is not used in data packet forwarding or in the control plane associated with the VDP segments. The group name is used only for VXLAN segments.
Step 5	switch (config-bd)# **segment transport-mode {native | vxlan}**	Specifies the default transport mode. The default transport mode is set to VXLAN. If you specify native, sets it to VDP global configuration mode.
Step 6	switch (config-bd)# **show running-config bridge-domain**	(Optional) Displays the segmentation configuration.
Step 7	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

show running -config bridge-domain

This example shows how to display the segmentation configuration.

```
switch # show running-config bridge-domain
version 4.2(1)SV2(2.2)
feature segmentation
no segment mode unicast-only
bridge-domain seg22222
segment id 22222
group 239.1.1.1
segment transport-mode native
fabric forwarding mode proxy-gateway
```

Configuring a DMAC from the VDP Station

To avoid any changes to the blade switches such as the Cisco UCS fabric interconnect and other network devices, you must manually configure the destination MAC address for ECP packets that originates from the VDP station, to enable forwarding data traffic upstream to the Cisco DFA leaf.

Note You must ensure that the same MAC configuration is present at the upstream Cisco DFA leaf.

Before You Begin

- You have installed and configured the Cisco Nexus 1000V for VMware vSphere software using the *Cisco Nexus 1000V Installation and Upgrade Guide*.
- Ensure that the Virtual Supervisor Module (VSM) and Virtual Ethernet Module (VEM) connectivity is functioning.
- Log in to the CLI in EXEC mode.

Procedure

	Command or Action	Purpose
Step 1	switch # **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **[no] evb mac**	Adds the destination MAC address information for the ECP packets originating from the VDP station to blade switches such as the Cisco UCS fabric interconnect and other similar network devices.
Step 3	switch # **show evb**	Displays the configured MAC addresses.
Step 4	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

show evb

This example shows how to display the evb information:

```
switch # show evb
Edge Virtual Bridging
Role : VDP Station
VDP Mac Address : 0180.0000.0000
VDP Resource Wait Delay : 22(66 secs)
VDP Reinit Keep Alive : 21(20 secs)
```

Specifying EVB TLV Parameters

Because the Cisco Nexus 1000V does not support the Link Layer Discovery Protocol (LLDP), VDP uses the EVB TLV communicated through the LLDP payloads to negotiate the VDP/ECP parameters. Use the following commands to configure the EVB TLV parameters:

Before You Begin

- Log in to the CLI in the EXEC mode.
- Configure the EVB feature to enable VDP on the Cisco Nexus 1000V .

Procedure

	Command or Action	Purpose
Step 1	switch **configure terminal**	Enters global configuration mode.
Step 2	switch (config)# **[no] ecp max-retries** *<1-7>*	(Optional) Configures the number of times ECP retries to send an upper layer protocol message. This parameter corresponds to the R value in the EVB TLV. If you specify a value of zero, the standard default value is used.
Step 3	switch (config)# **[no] ecp retransmission-timer-exponent** *<10-20>*	(Optional) Configures the exponential value of the interval for which ECP waits before trying to retransmit the packet. This parameter corresponds to the RTE value in the EVB TLV.
Step 4	switch (config)# **[no] evb resource-wait-delay** *<20-31>*	Configures the resource wait delay used by VDP to calculate the time it waits before concluding that a request has timed out. VDP will retry its request after the timeout.
Step 5	switch (config)# **[no] evb reinit-keep-alive** *<20-31>*	Configures the interval at which VDP refreshes the VSI state in the bridge by sending a VDP associate refresh.
Step 6	switch(config) # **show evb**	(Optional) Displays the configured VDP/ECP information.
Step 7	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to display details of the configured EVB TLV parameters on the Cisco Nexus 1000V:

```
 switch # show evb
Edge Virtual Bridging
Role : VDP Station
VDP Mac Address : 0000.1111.2222
VDP Resource Wait Delay : 20(17 secs)
VDP Reinit Keep Alive : 20(10 secs)
```

Verifying VDP Configuration

To display the VDP configuration information, use the following commands:

Command	Purpose
show evb	Displays the EVB segmentation information. See Example Example 1 - show evb, on page 14
show run evb	Displays the running configuration for the EVB segmentation. See Example Example 2 - show running-config evb, on page 14
show evb vsi *interface veth*	Displays the VDP VSI information from the Cisco Nexus 1000V VEMs. See Example Example 3 - show evb vsi interface , on page 14
show evb *module*	Displays EVB information for a module. See Example Example 4 - show evb module, on page 15
show ecp	Displays the ECP information. See Example Example 5 - show ecp, on page 15
show ecp [module *modid***]**	Displays the state information and statistics for ECP. See Example Example 6 - show ecp module, on page 15

Example 1 - show evb

This example shows how to display the EVB segmentation information.

```
switch # show evb
Edge Virtual Bridging
Role : VDP Station
VDP Mac Address : 0180.0000.0000
VDP Resource Wait Delay : 22(66 secs)
VDP Reinit Keep Alive : 21(20 secs)
```

Example 2 - show running-config evb

This example shows how to display the EVB segmentation configuration:

```
switch #: show running-config evb
evb resource-wait-delay 24
evb reinit-keep-alive 25
ecp retransmission-timer-exponent 15
ecp max-retries 6
```

Example 3 - show evb vsi interface

This example shows how to display the EVB vsi information from the Cisco Nexus 1000V VEMs:

```
switch# show evb vsi interface vethernet 15
LTL : 50 [module: 4]
```

```
Segment : 33333
MAC : 0050.5693.7D25
IP : 222.222.221.100
VSI State : 3
State Machine State : 7
Rwd Expiry Count : 37
Last CMD Time : 24
Last RSP Time : 21
```

Example 4 - show evb module

This example shows how to display EVB information for a module.

```
switch # show evb module 4
Edge Virtual Bridging
Role : VDP Station
VDP Mac Address : 0180.C200.0000
VDP Resource Wait Delay : 20(22 secs)
VDP Reinit Keep Alive : 25(335 secs)
n1kv-vsm#
```

Example 5 - show ecp

This example shows how to display the configuration information for ECP.

```
switch # show ecp
ECP Max ReTries : 3
ECP Retransmition Timer Exp : 14(163840 micro seconds)
```

Example 6 - show ecp module

This example shows how to display the statistics and state information for a module.

```
switch # show ecp mod 4
ECP Max ReTries : 3
ECP Retransmition Timer Exp : 14(163840 micro seconds)
TX Sequence No : 127
Retry Count : 0
TX Count : 0
TX Count Errors : 0
In TX Queue : 0
RX Count : 0
RX Sequence : 42634
```

Standards

The following table lists the standards supported in this release:

Standards	Title
IEEE 802.1Qbg	Edge Virtual Bridging (EVB) http://www.ieee802.org/1/pages/802.1bg.html

Feature History for Configuring VDP

Feature	Release	Feature information
VSI Discovery and Configuration protocol	4.2(1)SV2(2.2)	This feature was introduced.