



Overview

This chapter contains the following sections :

- [Information About the VSI Discovery and Configuration Protocol, page 1](#)
- [Features of VDP, page 1](#)
- [VDP Components in the Cisco Dynamic Fabric Automation Network, page 3](#)
- [VDP Sequence, page 4](#)

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). VDP can detect and signal the presence of end hosts and define message exchanges with an adjacent VDP-capable bridge. VDP is 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 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 multitenant 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:

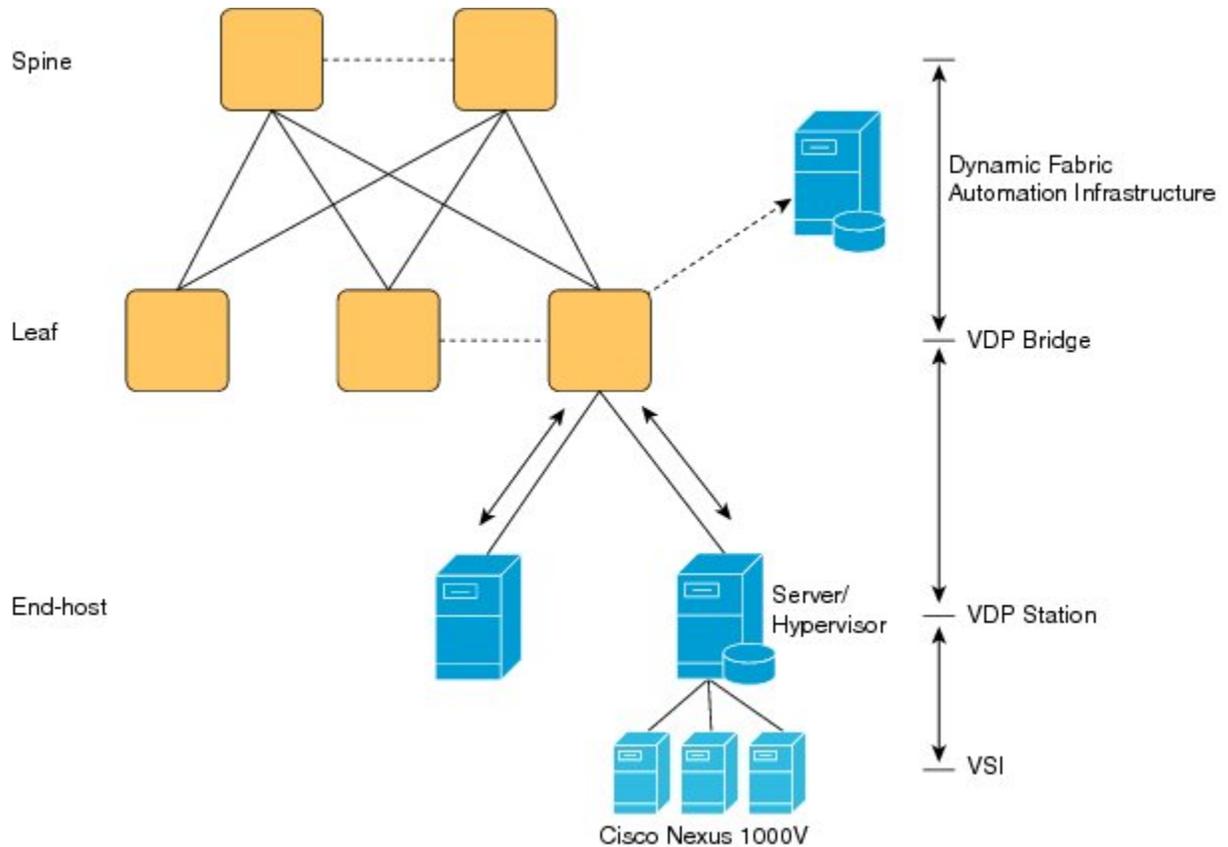
- Facilitates end-to-end segmentation enabled in the Cisco Dynamic Fabric Automation (DFA) architecture and removes the disadvantages of the host-based overlays.
- Serves as an end-host registration protocol for the Cisco DFA leaf switches that use the registration information to automatically configure 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 as follows:

- 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 Virtual Ethernet Module (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—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, Access Control List (ACL) or Quality of Service (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 Type Length Values (TLVs) 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 Organizationally Unique Identifier (OUI) TLV.
- 2 After receiving the request, VDP on the leaf switch 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 tries to send the request again after a configurable interval.