



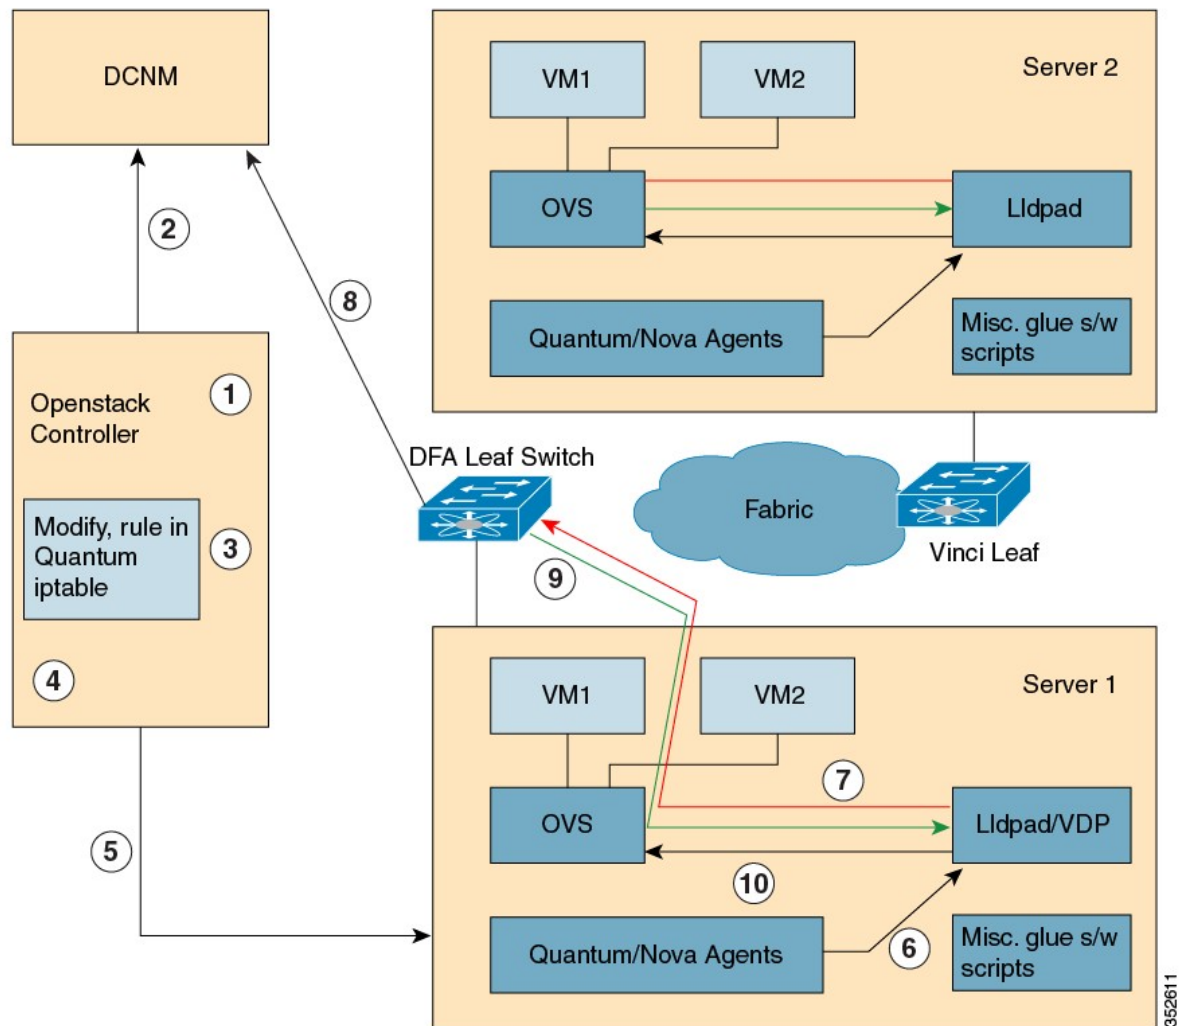
OpenStack for DFA Overview

- [Overview of OpenStack for DFA, page 2](#)

Overview of OpenStack for DFA

OpenStack serves as one of orchestrators of the cloud enabled through DFA. For this release, all the orchestration is done using OpenStack's dashboard Horizon graphic user interface.

Figure 1: OpenStack for DFA System Flow



Note

Steps 2, 4, and 6 can be disabled, along with other external DFA modules, so OpenStack will continue to function natively, without DFA.

The above diagram provides an overview of the system in the DFA context. The control flow can be summarized as follows:

- 1 A network of type "tunnel" with a tunnel ID range that is greater than 4000 is created. In OpenStack, the tunnel ID is a segment ID.
- 2 Network Information (Subnet/mask, tenant name) is sent to the Cisco Prime Data Center Network Manager (DCNM) through the Cisco Prime DCNM representational state transfer (REST) application programming interfaces (APIs).
- 3 A virtual machine (VM) instance is launched and the network for this instance is specified.
- 4 The host- specific IP table rules in the quantum agent are modified to unblock incoming Dynamic Host Configuration Protocol (DHCP) frames from the outside world and to add rules that allow only incoming frames with OpenStack-assigned source IP addresses.
- 5 Information about the network (subnet/mask, tenant name) and the VM is sent to the compute node.
- 6 Vision station interface discovery and provisioning (VDP) is notified about the VM and the segment ID for the VM.
- 7 VDP passes the information and segment ID for the VM to the Cisco DFA leaf switch.
- 8 The leaf switch contacts the Cisco Prime DCNM with the VM segment ID for retrieving the network attributes.
- 9 The leaf switch responds with the VLAN ID for tagging the traffic from the VM.
- 10 The Open vSwitch (OVS) is configured for tagging the packets from the VM to the network with the VLAN ID provided by the leaf switch. The virtual interface card (vNIC) is now operational.

