

Contents

[Introduction](#)

[Problem](#)

[Solution](#)


Introduction

This document describes a currently unsupported VxLAN inter-data center design.

Problem

Considering you have two data centers, where each has a few nexus switches acting as VxLAN leaf/spine and you configure the leaf as Anycast gateway.

Now you want to connect the two data centers together using ethernet or other data center interconnect (DCI) technology like Overlay Transport Virtualization (OTV).



If you enable ARP suppression under a VxLAN Network Identifier (VNI), you will likely experience connectivity issue when two hosts in different Vlan are trying to communicate across the data center.

Solution

This issue is caused by following sequence.

1. A local VM sends a packet to a remote VM via local Anycast gateway, which is the local nexus leaf.
2. The nexus leaf receives the packet and checks the destination IP address, which is directly connected. It then sends ARP request sourcing from the Anycast IP, which is normally configured the same on both data center.
3. Since both ARP suppression is enabled and SVI is up, the remote nexus leaf will intercept the ARP request. Due to ARP duplication check, this ARP request is considered a duplication of local IP and is dropped silently.

This logic caused inter-data center communication break when the two hosts in the different Vlan are trying to talk.

Cisco is aware of this issue and is working on a solution to resolve this in the future release. The workaround for now is to disable ARP suppression under VNI.