UCS L2 Multicast with Nexus 5000 and 1000V Series Switches Configuration Example

Document ID: 117360

Contributed by Vishal Mehta, Cisco TAC Engineer.
Jan 27, 2014

Contents

Introduction

This document describes how to configure and troubleshoot Layer 2 (L2) multicast for Virtual Machines (VMs) upon setup of the Cisco Unified Computing System (UCS), Cisco Nexus 1000V Series switches (N1kV), and Cisco Nexus 5000 Series switches (N5k).

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Basics of Multicast
- Cisco UCS
- N1kV
- N5k

Components Used

The information in this document is based on these software and hardware versions:

- Cisco Nexus 5020 Series Switch Version 5.0(3)N2(2a)
- Cisco UCS Version 2.1(1d)
- Cisco UCS B200 M3 Blade Server with Cisco Virtual Interface Card (VIC) 1240
- vSphere 5.1 (ESXi and vCenter)

- Cisco N1kV Version 4.2(1)SV2(1.1a)

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command or packet capture setup.

Background Information

Multicast was initially designed to use Layer 3 (L3) functionality, where multiple hosts from a network subscribe to a multicast address. The new trend is to use L2 multicast functionality, where traffic flows between VMs that participate in a multicast application across hosts on the same VLAN. Such multicast traffic stays within the same L2 domain and does not need a router.

When there is no multicast router in the VLAN that originates the queries, you must configure an Internet Group Management Protocol (IGMP) snooping querier in order to send membership queries. IGMP snooping is enabled by default on the UCS, N1kV, and N5k. You can enable IGMP snooping querier on either the UCS or a N5k, dependent upon the scope of the L2 multicast. If there are multicast receivers outside of the UCS, configure the snooping querier on the N5k.

When an IGMP snooping querier is enabled, it sends out periodic IGMP queries that trigger IGMP report messages from hosts that want to receive IP multicast traffic. IGMP snooping listens to these IGMP reports in order to establish appropriate forwarding.

The IGMP snooping software examines IGMP protocol messages within a VLAN in order to discover the interfaces that are connected to hosts or other devices interested in receiving this traffic. With the interface information, IGMP snooping can reduce bandwidth consumption in a multi−access LAN environment in order to avoid a flood of the entire VLAN. The IGMP snooping feature tracks the ports that are attached to multicast−capable routers in order to help manage the forwarding of IGMP membership reports. Also, the IGMP snooping software responds to topology change notifications.

Configure

Use this section in order to configure L2 multicast for VMs.

Network Setup

Here are some important notes about the network setup in this example:

- The UCS is connected to a N5k through a Virtual Port Channel (vPC).
- The Operating System (OS) that is installed on both of the hosts is VMware ESXi 5.1. Each host has VMs with Microsoft Windows 2012 Guest−OSs.
- The source of the multicast is ***MCAST VM*** (IP address 172.16.16.226) on host IP address 172.16.16.222 (UCS Blade 1/5), that sends traffic to multicast IP address 239.14.14.14.
- The multicast receivers are ***AD−1 VM*** (IP address 172.16.16.224) on host IP address 172.16.16.220 (UCS Blade 1/6), and ***TEST VM*** (IP address 172.16.16.228) on host IP address 172.16.16.222 (UCS Blade 1/5).
- The IGMP snooping querier is configured on the N5k with an IP address of 172.16.16.2, and also on the UCS with an IP address of 172.16.16.233.

There is no need to configure two queriers in the same VLAN (16). If there are multicast receivers outside of the UCS, configure the snooping querier on the N5k. If the multicast traffic is within the UCS domain, then create the snooping querier on the Cisco Unified Computing System Manager (UCSM).

Note: The N5k IGMP querier is elected per ***RFC 4605***, which explains the querier election process.

N5k IGMP Querier Configuration

Here is an example configuration of an IGMP querier on a N5k:

```
vlan 16

  ip igmp snooping querier 172.16.16.2

!

int vlan 16

  ip address 172.16.16.2/24

  no shut
```

The querier IP address does not need to be for a switched−virtual interface, and it can be a different IP address within the same subnet of VLAN 16.

Note: Refer to the Configuring IGMP Snooping section of the ***Cisco Nexus 5000 Series NX−OS Software Configuration Guide*** for information about how to configure the IGMP querier for your specific version.

UCS IGMP Querier Configuration

Complete these steps in order to configure the IGMP querier for UCS:

1. Create a new multicast policy under the ***LAN*** tab of the UCSM, as shown here:

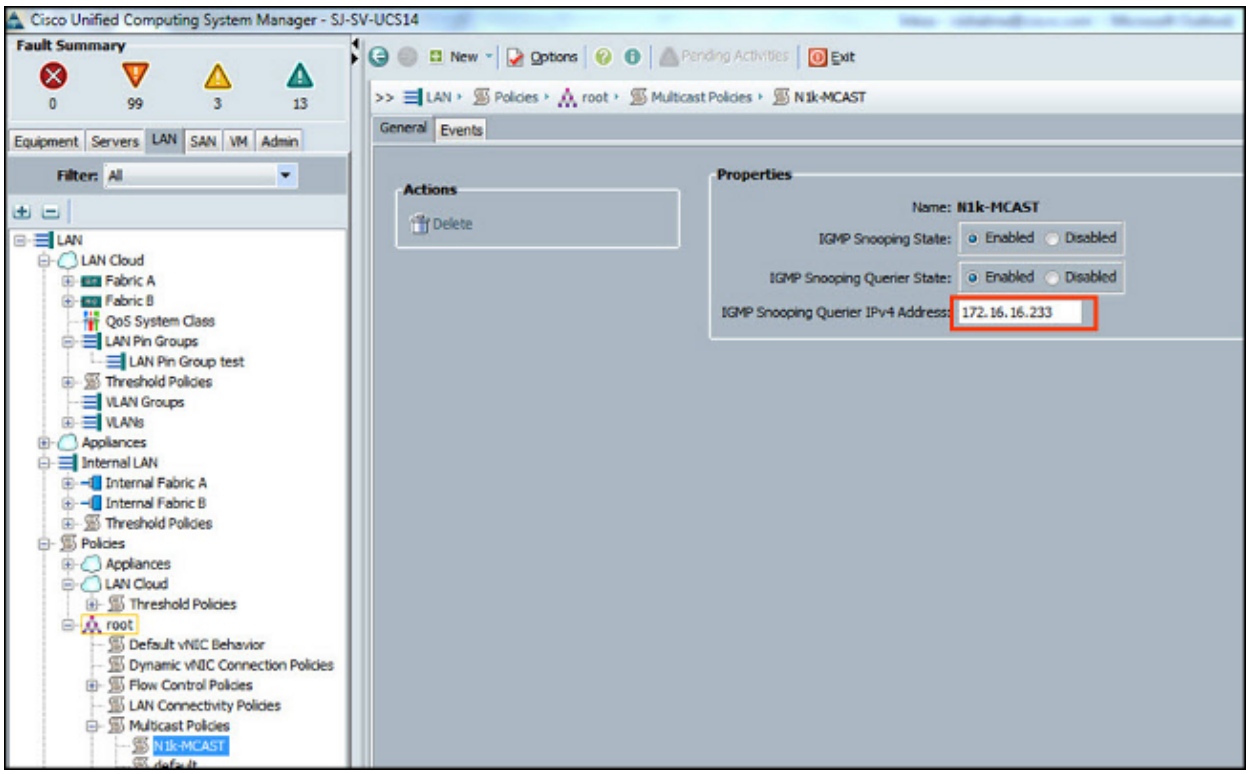

2. Apply multicast policy ***N1k−MCAST*** to VLAN 16:

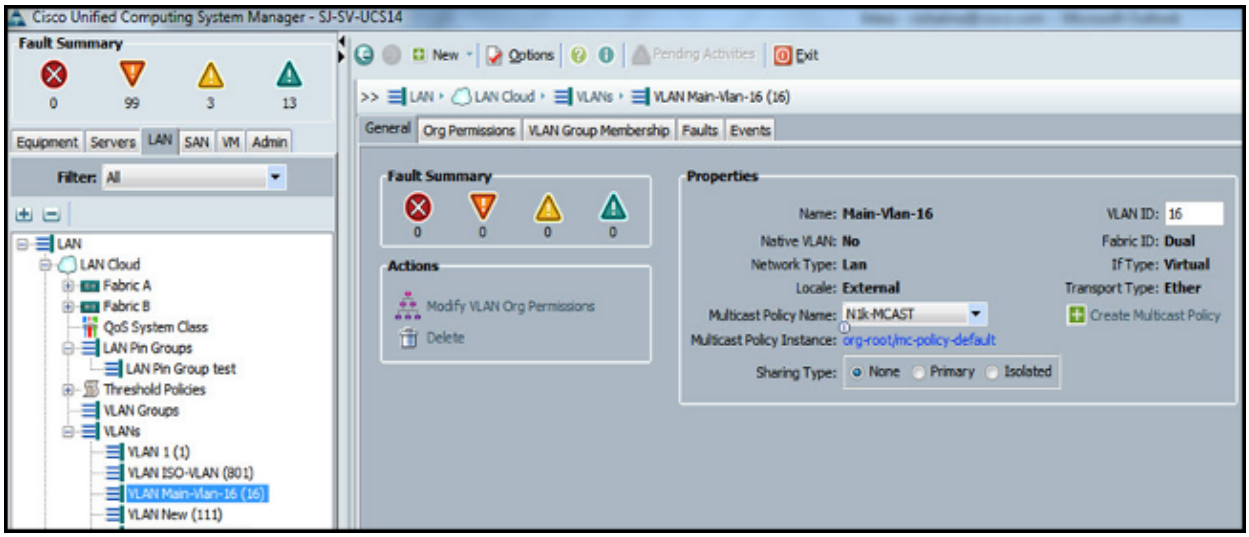

3. For the N1kV, confirm that IGMP snooping is enabled on VLAN 16 (which is enabled by default). No configuration must be done on an N1kV in order to support basic L2 multicast.

Note: A VideoLAN Client (VLC) media player is used in order to demonstrate multicast. For more details on how to use a VLC player for multicast streaming, refer to the How to use VLC Media player to stream multicast video article.

Verify

Use this section in order to verify that your configuration works correctly.

Verification on the N1kV

Verify that the multicast receivers ***TEST VM*** and ***AD−1 VM*** have joined multicast stream ***239.14.14.14***, from which ***MCAST VM*** sources traffic. This image shows that multicast receiver ***TEST VM*** receives the stream:

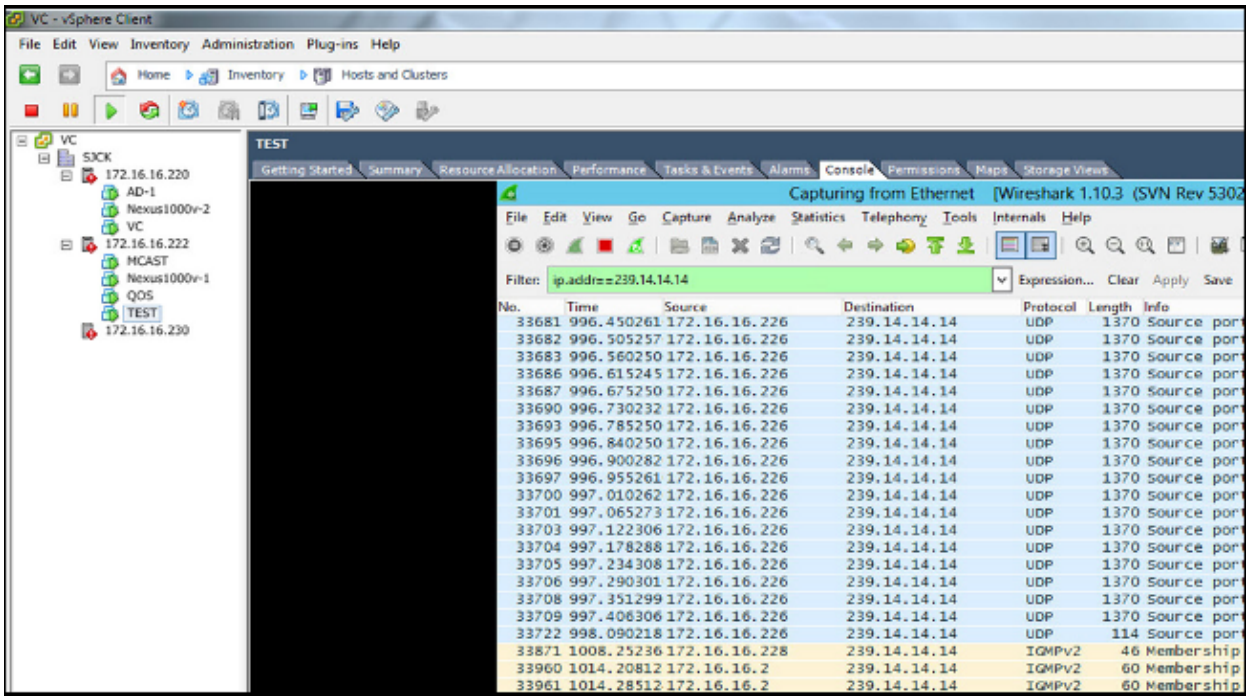

The N1kV snooping output shows the Group Address and the Veths of the multicast receiver, not the Veth of the VM that sources the multicast traffic (as expected):

```
Nexus1000v# sh ip igmp snooping groups

Type: S - Static, D - Dynamic, R - Router port

Vlan  Group Address       Ver   Type   Port list
16    */*                 -     R      Eth3/2 Eth4/2
16    239.14.14.14        v2    D      Veth3 Veth6
```

This N1kV output shows the active ports for multicast and the IGMP querier:

```
Nexus1000v# sh ip igmp snooping groups vlan 16
IGMP Snooping information for vlan 16
  IGMP snooping enabled
  IGMP querier present, address: 172.16.16.2, version: 2, interface Ethernet4/2
  Switch-querier disabled
  IGMPv3 Explicit tracking enabled
  IGMPv2 Fast leave disabled
  IGMPv1/v2 Report suppression disabled
  IGMPv3 Report suppression disabled
  Link Local Groups suppression enabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 2
  Number of groups: 1
  Active ports:
    Veth1       Eth3/2  Veth2   Eth4/2
    Veth3       Veth4   Veth5   Veth6
```

At the host level, you can verify that multicast traffic is received by the VMs that participate. This output shows the VM ***AD−1***, which is on ***Module 3*** of the Virtual Supervisor Module (VSM):

```
Nexus1000v#  module vem 3 execute vemcmd show bd

BD 7, vdc 1, vlan 16, swbd 16, 3 ports, ""

Portlist:
     18  vmnic1
     49  vmk0
     50  AD-1 ethernet0

Multicast Group Table:
Group 239.14.14.14 Multicast LTL 4672
    18
    50
Group 0.0.0.0 Multicast LTL 4671
    18
```

This output shows the VM ***TEST***, which is on ***Module 4*** of the VSM:

```
Nexus1000v#  module vem 4 execute vemcmd show bd

BD 7, vdc 1, vlan 16, swbd 16, 6 ports, ""

Portlist:
     18   vmnic1
     49   vmk0
     50   TEST.eth0
     51   QOS.eth0
     52   MCAST.eth0 <-- Source
    561

Multicast Group Table:
Group 239.14.14.14 Multicast LTL 4672
    50
    561
Group 0.0.0.0 Multicast LTL 4671
    561
```

Verification on the UCS

This UCS output shows the active ports for multicast and the ***Group Address***:

```
SJ-SV-UCS14-B(nxos)# sh ip igmp snooping group
Type: S - Static, D - Dynamic, R - Router port

Vlan  Group Address      Ver  Type  Port list
1     */*                -    R     Po1
11    */*                -    R     Po1
15    */*                -    R     Po1
16    */*                -    R     Po1
16    239.14.14.14       v2   D     Veth1257 Veth1255
30    */*                -    R     Po1
111   */*                -    R     Po1
172   */*                -    R     Po1
800   */*                -    R     Po1
```

This UCS snooping output for VLAN 16 verifies that the querier is configured on the UCSM and the N5k, and it shows that only the querier on the N5k is currently active (as expected):

```
SJ-SV-UCS14-B(nxos)# sh ip igmp snooping vlan 16
IGMP Snooping information for vlan 16
  IGMP snooping enabled
  Optimised Multicast Flood (OMF) disabled
  IGMP querier present, address: 172.16.16.2, version: 2, interface port-channel1
  Switch-querier enabled, address 172.16.16.233, currently not running
  IGMPv3 Explicit tracking enabled
  IGMPv2 Fast leave disabled
  IGMPv1/v2 Report suppression enabled
  IGMPv3 Report suppression disabled
  Link Local Groups suppression enabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 1
  Number of groups: 1
  Active ports:
    Po1 Veth1257        Veth1251        Veth1255
    Veth1279    Veth1281
```

Verification on the N5k

On the N5k, confirm that multicast group address ***239.14.14.14*** and the active port−channel is connected to the UCS Fabric Interconnects (FIs):

```
n5k-Rack18-1# sh ip igmp snooping groups
Type: S - Static, D - Dynamic, R - Router port, F - Fabricpath core port

Vlan  Group Address      Ver  Type  Port list
1     */*                -    R     Po40
15    */*                -    R     Po40 Po1110 Po1111
15    239.255.255.253    v2   D     Po10 Po11 Po12
                                     Po13 Po40
16    */*                -    R     Po3 Po40
16    239.14.14.14       v2   D     Po15 Po16
17    */*                -    R     Po40
18    */*                -    R     Po40
```

Troubleshoot

This section provides information that you can use in order to troubleshoot your configuration.

Here is a list of basic caveats about multicast in the L2 domain:

- If IGMP snooping is not enabled on the switch, then multicast traffic is broadcast within the L2 domain.
- If IGMP snooping is enabled, a querier must run on the uplink switches on the VLAN that contain multicast sources and receivers.
- If there is no IGMP querier in the VLAN, the N1kV and the UCS do not forward the multicast. This is the most common misconfiguration seen in Cisco Technical Assistance Center (TAC) cases.
- By default, IGMP snooping is enabled on both the N1kV and the UCS.
- With UCS Versions 2.1 and later, IGMP snooping can be enabled or disabled per−VLAN, and the IGMP querier can be configured at the UCS level.

Updated: Jan 27, 2014 Document ID: 117360