

Unicast Flooding in Switched Campus Networks

Document ID: 23563

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Problem Definition

Causes of Flooding

- Cause 1: Asymmetric Routing

- Cause 2: Spanning-Tree Protocol Topology Changes

- Cause 3: Forwarding Table Overflow

How to Detect Excessive Flooding

Related Information

Introduction

This document discusses possible causes and implications of unicast packet flooding in switched networks.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

For more information on document conventions, refer to Cisco Technical Tips Conventions.

Problem Definition

LAN switches use forwarding tables (Layer 2 (L2) tables, Content Addressable Memory (CAM) tables) to direct traffic to specific ports based on the VLAN number and the destination MAC address of the frame. When there is no entry corresponding to the frame's destination MAC address in the incoming VLAN, the (unicast) frame will be sent to all forwarding ports within the respective VLAN, which causes flooding.

Limited flooding is part of the normal switching process. There are situations, however, when continuous flooding can cause adverse performance effects on the network. This document explains what issues can arise due to flooding, and the most common reasons why certain traffic might constantly be flooded.


Note that most modern switches including the Catalyst 2900 XL, 3500 XL, 2940, 2950, 2970, 3550, 3750, 4500/4000, 5000, and 6500/6000 series switches maintain L2 forwarding tables per VLAN.

Causes of Flooding

The very cause of flooding is that destination MAC address of the packet is not in the L2 forwarding table of the switch. In this case the packet will be flooded out of all forwarding ports in its VLAN (except the port it was received on). Below case studies display most common reasons for destination MAC address not being known to the switch.

Cause 1: Asymmetric Routing

Large amounts of flooded traffic might saturate low-bandwidth links causing network performance issues or complete connectivity outage to devices connected across such low-bandwidth links. Consider the following diagram:



In the diagram above, server S1 in VLAN 1 is running backup (bulk data transfer) to server S2 in VLAN 2. Server S1 has its default gateway pointing to router A's VLAN 1 interface. Server S2 has its default gateway pointing to router B's VLAN 2 interface. Packets from S1 to S2 will follow this path:

- S1—VLAN 1—switch A—router A—VLAN 2—switch B—VLAN 2—S2 (blue line)

Packets from S2 to S1 go along the following path:

- S2—VLAN 2—switch B—router B—VLAN 1—switch A—flooded to VLAN 1—S1 (red line)

Note that with such an arrangement, switch A will not "see" traffic from the S2 MAC address in VLAN 2 (since the source MAC address will be rewritten by router B and the packet will only arrive in VLAN 1). This means that every time switch A needs to send the packet to the S2 MAC address, the packet will be flooded to VLAN 2. The same situation will occur with the S1 MAC address on switch B.

This behavior is called asymmetric routing. Packets follow different paths depending on the direction. Asymmetric routing is one of the two most common causes of flooding.

Impact of Unicast Flooding

Returning to the above example, the result is that packets of the data transfer between S1 and S2 will mostly be flooded to VLAN 2 on switch A and to VLAN 1 on switch B. This means every connected port (workstation W in this example) in VLAN 1 on switch B will receive all packets of conversation between S1 and S2. Suppose the server backup takes 50 Mbps of bandwidth. This amount of traffic will saturate 10 Mbps

links. This will cause a complete connectivity outage to the PCs or slow them down considerably.

This flooding is due to asymmetric routing, and may stop when server S1 sends a broadcast packet (for example Address Resolution Protocol (ARP)). Switch A will flood this packet to VLAN 1 and switch B will receive and learn the MAC address of S1. Since the switch is not receiving traffic constantly, this forwarding entry will eventually age out and flooding will resume. The same process applies to S2.

There are different approaches to limit the flooding caused by asymmetric routing. Refer to these documents for more information:

- Asymmetric Routing with Bridge Groups on Catalyst 2948G–L3 and 4908G–L3 Switches
- Asymmetric Routing and HSRP (Excessive Flooding of Unicast Traffic in Network with Routers Running HSRP)

The approach is normally to bring the router's ARP timeout and the switches' forwarding table–aging time close to each other. This will cause the ARP packets to be broadcast. Relearning must occur before the L2 forwarding table entry ages out.

A typical scenario where this kind of issue might be observed is when there are redundant Layer 3 (L3) switches (such as a Catalyst 6000 with Multilayer Switch Feature Card (MSFC)) configured to load–balance with Hot Standby Router Protocol (HSRP). In this case, one switch will be active for even VLANs and the other one will be active for odd VLANs.

Cause 2: Spanning–Tree Protocol Topology Changes

Another common issue caused by flooding is Spanning–Tree Protocol (STP) Topology Change Notification (TCN). TCN is designed to correct forwarding tables after the forwarding topology has changed. This is necessary to avoid a connectivity outage, as after a topology change some destinations previously accessible via particular ports might become accessible via different ports. TCN operates by shortening the forwarding table aging time, such that if the address is not relearned, it will age out and flooding will occur.

TCNs are triggered by a port that is transitioning to or from the forwarding state. After the TCN, even if the particular destination MAC address has aged out, flooding should not happen for long in most cases since the address will be relearned. The issue might arise when TCNs are occurring repeatedly with short intervals. The switches will constantly be fast–aging their forwarding tables so flooding will be nearly constant.

Normally, a TCN is rare in a well–configured network. When the port on a switch goes up or down, there is eventually a TCN once the STP state of the port is changing to or from forwarding. When the port is flapping, repetitive TCNs and flooding occurs.

Ports with the STP portfast feature enabled will not cause TCNs when going to or from the forwarding state. Configuration of portfast on all end–device ports (such as printers, PCs, servers, and so on) should limit TCNs to a low amount. Refer to this document for more information on TCNs:

- Understanding Spanning–Tree Protocol Topology Changes

Note: In MSFC IOS, there is an optimization that will trigger VLAN interfaces to repopulate their ARP tables when there is a TCN in the respective VLAN. This limits flooding in case of TCNs, as there will be an ARP broadcast and the host MAC address will be relearned as the hosts reply to ARP.

Cause 3: Forwarding Table Overflow

Another possible cause of flooding can be overflow of the switch forwarding table. In this case, new addresses

cannot be learned and packets destined to such addresses are flooded until some space becomes available in the forwarding table. New addresses will then be learned. This is possible but rare, since most modern switches have large enough forwarding tables to accommodate MAC addresses for most designs.

Forwarding table exhaustion can also be caused by an attack on the network where one host starts generating frames each sourced with different MAC address. This will tie up all the forwarding table resources. Once the forwarding tables become saturated, other traffic will be flooded because new learning cannot occur. This kind of attack can be detected by examining the switch forwarding table. Most of the MAC addresses will point to the same port or group of ports. Such attacks can be prevented by limiting the number of MAC addresses learned on untrusted ports by using the port security feature.

Configuration Guides for Catalyst switches running Cisco IOS® or CatOS software have a section called Configuring Port Security or Configuring Port-Based Traffic Control. Refer to the Technical Documentation for your switch on the Cisco Switches product pages for more information.

How to Detect Excessive Flooding

Most switches implement no special command to detect flooding. Catalyst 6500/6000 Supervisor Engine 2 and higher series switches running Cisco IOS System software (Native) version 12.1(14)E and higher or Cisco CatOS system software version 7.5 or higher implements '**unicast flood protection**' feature. In short, this feature allows the switch to monitor the amount of unicast flooding per VLAN and take specified action if flooding exceeds specified amount. Actions can be to syslog, limit or shutdown VLAN – the syslog being the most useful for flood detection. When flooding exceeds the configured rate and the action configured is syslog, a message similar to the following will be printed:

```
%UNICAST_FLOOD-4-DETECTED: Host 0000.0000.2100 on vlan 1 is flooding
to an unknown unicast destination at a rate greater than/equal to 1 Kfps
```

The MAC address indicated is the source MAC from which the packets are flooded on this switch. It is often needed to know the destination MAC addresses to which switch is flooding (because switch is forwarding by looking at the destination MAC address). Cisco IOS (Native) versions 12.1(20)E for catalyst 6500/6000 supervisor engine 2 and on will implement capability to display the MAC addresses to which flooding is occurring:

```
cat6000#sh mac-address-table unicast-flood
Unicast Flood Protection status: enabled

Configuration:
vlan      Kfps      action      timeout
-----+-----+-----+-----
      55          1          alert       none

Mac filters:
No.  vlan  souce mac addr.      installed on      time left (mm:ss)
-----+-----+-----+-----+-----

Flood details:
Vlan  souce mac addr.      destination mac addr.
-----+-----+-----+-----
      55  0000.2222.0000      0000.1111.0029, 0000.1111.0040, 0000.1111.0063
                                0000.1111.0018, 0000.1111.0090, 0000.1111.0046
                                0000.1111.006d
```

Further investigation can then be carried out to see if MAC address 0000.2222.0000 is supposed to be sending traffic to the MAC addresses listed in the destination MAC address section. If traffic is legitimate, then one would need to establish why destination MAC addresses are not known to the switch.

For information on the unicast flood protection configuration, refer to these documents:

- [Configuration](#)
- [Monitoring](#)

You may detect if flooding is occurring by capturing a trace of packets seen on a workstation during the time of slowdown or outage. Normally, unicast packets not involving the workstation should not be seen repeatedly on the port. If this is happening, chances are that there is flooding occurring. Packet traces may look different when there are various causes of flooding.

With asymmetric routing, there are likely to be packets to specific MAC address that will not stop flooding even after the destination replies. With TCNs, the flooding will include many different addresses, but should eventually stop and then restart.

With L2 forwarding table overflow, you are likely to see same kind of flooding as with asymmetric routing. The difference is that there will likely be a high amount of strange packets, or normal packets in abnormal quantities with a different source MAC address.

Related Information

- [LAN Technologies Index Page](#)
- [LAN and ATM Switching](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Nov 17, 2007

Document ID: 23563
