How Much Broadcast and Multicast Traffic Should I Allow in My Network?

This Application Note relates to the following Dell PowerConnect™ products:

- PowerConnect 33xx
- PowerConnect 52xx

Abstract

Broadcast and multicast traffic perform valuable roles in terms of network discovery and content delivery, but too much of either can have an adverse effect on performance. Excessive amounts of broadcast or multicast traffic not only waste bandwidth, but also degrade the performance of every device attached to the network. This application note discusses the steps network managers can take to minimize the effects of broadcast and multicast traffic without compromising application functionality or performance.

Applicable Network Scenarios

The diagram below depicts a corporate network carrying multiprotocol traffic. While there is no universal threshold that determines how much broadcast or multicast traffic is too much, this is definitely a network at risk. Each of the various protocols on the network potentially can degrade the performance of all stations attached to the network.

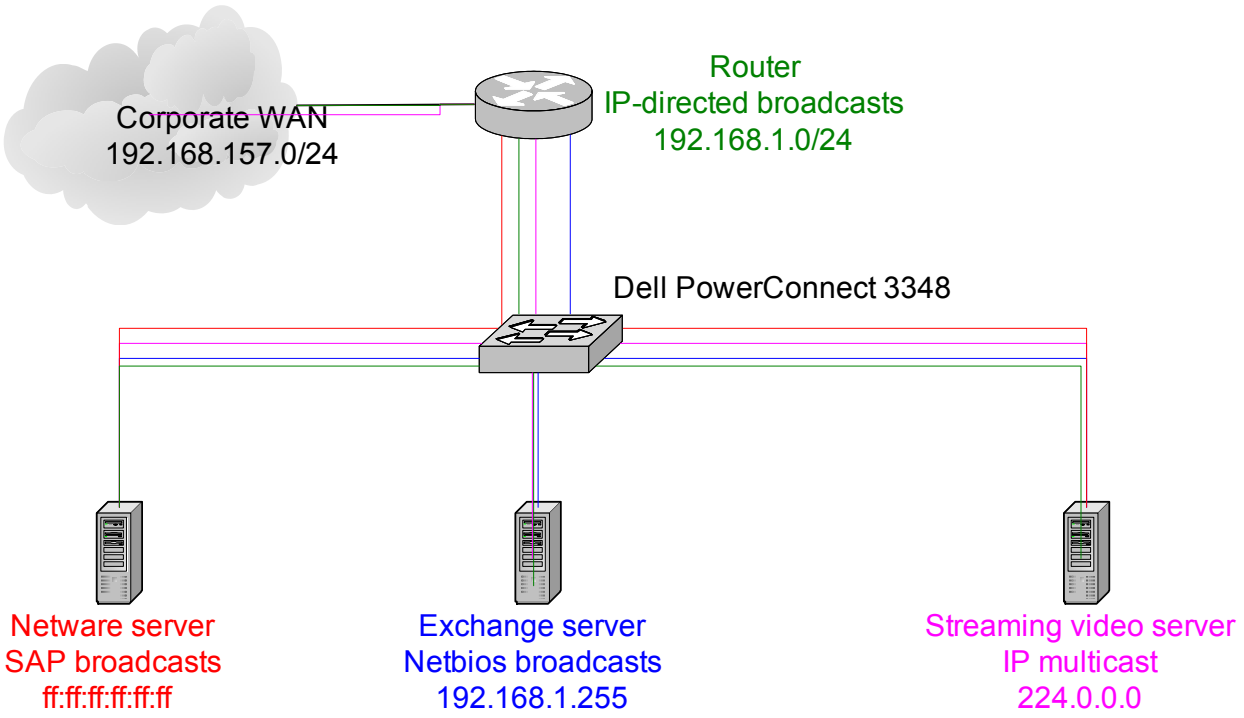

The Microsoft® Exchange server uses Netbios broadcasts for host discovery, while the Novell® Netware® server transmits IPX/SPX service access point (SAP) broadcasts. Both Netbios and IPX/SPX are "chatty" protocols, meaning each transmits large amounts of broadcast traffic.

A router connects this segment to other networks. Routers may forward IP-directed broadcasts, again consuming bandwidth.

All these broadcasts tie up system resources as well as consuming network bandwidth. Every node (PC, server, printer, or any other network-attached device) in a given broadcast domain must process each broadcast frame it receives. When a node receives a broadcast, it generates an interrupt. In turn, each interrupt consumes some amount of processing time by the node.

Multicast traffic is somewhat more benign, though it too can affect network performance. The streaming video server in this diagram uses IP multicast for traffic delivery. While multicast reduces bandwidth consumption by sending traffic only to subscribers of a given group, multicast-enabled applications tend to be bandwidth-intensive.

Measure first, then manage

Before applying any traffic management commands to switches, it is essential to understand how much unicast, multicast, and broadcast traffic exists in the network. The amount of broadcast and multicast traffic on a given network may be minuscule, or it may overload the existing infrastructure. The adage "you can't manage what you can't measure" applies here.

Because Dell PowerConnect switches keep track of all traffic they handle, it is relatively simple to determine the amount of multicast vs. non-multicast traffic in the network.

Every interface of every Dell PowerConnect switch has counters that measure the number of frames, bytes, and errors it sees. These counters are broken down by transmit and receive traffic, and also by the type of traffic handled – multicast, unicast, or broadcast.

Here is the command-line interface sequence to view counters for interface 23 of a Dell PowerConnect 3348. The same command applies on a PowerConnect 3324:

```
console# show interface counters ethernet 1/e23
      Port        InOctets  InUcastPkts InMcastPkts InBcastPkts
---------------- ---------- ----------- ----------- -----------
     1/e23          1488159           1     1488095

      Port       OutOctets  OutUcastPkts OutMcastPkts OutBcastPkts
---------------- ---------- ------------ ------------ ------------
     1/e23             6400            0          100

Alignment Errors: 0
FCS Errors: 0
Single Collision Frames: 0
Multiple Collision Frames: 0
SQE Test Errors: 0
Deferred Transmissions: 0
Late Collisions: 0
Excessive Collisions: 0
Internal MAC Tx Errors: 0
Carrier Sense Errors: 0
Oversize Packets: 0
Internal MAC Rx Errors: 0
Symbol Errors: 0
Received Pause Frames: 0
Transmitted Pause Frames: 0
```

In the instance above, the interface has received 1 unicast frame and 1,488,095 multicast (what type of frames? Multicast?) frames. Therefore, in this case the ratio of unicast to inbound multicast traffic is 1:1,488,095.

Here is a similar example for viewing interface statistics from interface 24 of a Dell PowerConnect 5224:

```
Vty-0#show interfaces counters ethernet 1/12
Ethernet 1/12
 Iftable stats:
  Octets input: 1564024419, Octets output: 649097
  Unicast input: 7429, Unicast output: 4762
  Discard input: 0, Discard output: 0
  Error input: 0, Error output: 0
  Unknown protos input: 0, QLen output: 0
 Extended iftable stats:
  Multi-cast input: 75630, Multi-cast output: 3981
  Broadcast input: 1115761, Broadcast output: 0
 Ether-like stats:
  Alignment errors: 0, FCS errors: 0
  Single Collision frames: 0, Multiple collision frames: 0
  SQE Test errors: 0, Deferred transmissions: 0
  Late collisions: 0, Excessive collisions: 0
  Internal mac transmit errors: 0, Internal mac receive errors: 0
  Frame too longs: 0, Carrier sense errors: 0
  Symbol errors: 0
 RMON stats:
  Drop events: 0, Octets: 1564673940, Packets: 1207574
  Broadcast pkts: 1115763, Multi-cast pkts: 79614
  Undersize pkts: 0, Oversize pkts: 0
  Fragments: 0, Jabbers: 0
  CRC align errors: 0, Collisions: 0
  Packet size <= 64 octets: 124645, Packet size 65 to 127 octets: 37922
  Packet size 128 to 255 octets: 3740, Packet size 256 to 511 octets: 20599
  Packet size 512 to 1023 octets: 29, Packet size 1024 to 1518 octets: 1020639
```

In this example with a PowerConnect 5224, the interface has received 7,429 unicast frames, 75,630 multicast frames, and 1,115,761 broadcast frames. The ratio of unicast to multicast in this case is about 1:10, but the broadcast frame count dwarfs both the unicast and multicast totals.

Interface counters have two limitations when it comes to determining multicast-vs.-non-multicast traffic breakdowns.

First, counters are cumulative. They present information about traffic over time, not real-time rates. So, for example, it may appear that a switch interface handled about the same amounts of multicast and unicast traffic, even though hours or days have passed between the times when multicast and unicast traffic were active.

The best way to keep track of traffic statistics is to import them into a network management system such as one using simple network management protocol (SNMP). Dell PowerConnect switches support standards-based SNMP and remote monitoring (RMON), an extension to SNMP intended specifically for device and traffic management. By tracking traffic over time, it is possible to get a sense of longer-term trends of multicast vs. non-multicast usage.

This sort of long-term observation is essential in capacity planning – the science of determining network resource requirements based on measurements of past and current trends.

A second problem with counters is that they are accurate only as of the last system reset, the last time the counters were cleared, or the last time the counter “wrapped” (reached its maximum value, at which time the counter reverts to 0 and starts counting up again). Here again, importing counter statistics into

an external network management system will provide the most accurate picture of traffic statistics over time.

Technology Background

Dell PowerConnect switches offer several mechanisms for reducing the overhead associated with broadcast and multicast traffic. These mechanisms include IEEE 802.1Q VLANs, storm control, and Internet Group Management Protocol (IGMP) snooping.

VLANs logically segment network nodes into a common broadcast domain isolated from other VLANs and physical segments. Thus, even if broadcast or multicast traffic is heavy on one VLAN, it can be isolated from the rest of the network.

If isolating a segment within a VLAN is not an option, network managers can implement another mechanism known as storm control. Storm control regulates the rate at which switches forward broadcast and multicast traffic. Dell PowerConnect switches allow network managers to specify a maximum rate (given in packets per second) for broadcast traffic.

To control rates of multicast traffic, the PowerConnect switches support a feature called IGMP snooping. Normally, a switch receiving a multicast frame will “flood” it – that is, transmit the frame on all ports in the VLAN associated with the interface that injected it. With IGMP snooping enabled, the switch floods ports only once. After the initial flood, the switch intercepts IGMP traffic sent between an IGMP router and hosts wishing to join a multicast group. Once a switch determines which hosts joined the multicast group, it forwards all subsequent multicast traffic only to those ports to which the hosts are attached.

Proposed Solution

Overview

Network managers can address the performance degradation associated with the excessive broadcast and multicast traffic by taking the following steps:

1. Disable IP-directed broadcasts on the router
2. Enable storm control on Dell PowerConnect switches
3. Verify that storm control is operating correctly
4. Enable IGMP snooping on Dell PowerConnect switches

Typical Network Designs

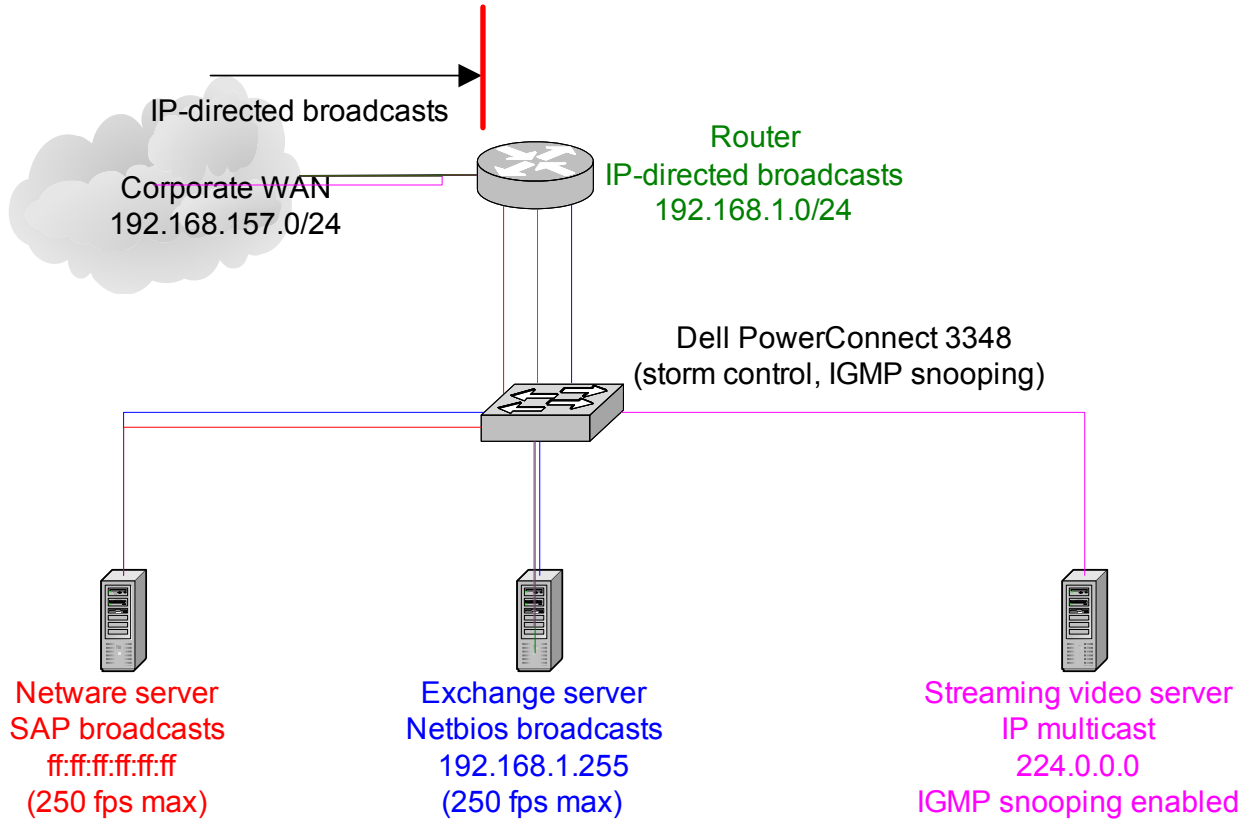

With IP-directed broadcasts disabled on the router, all external traffic destined for 192.168.1.255 will be filtered and dropped, keeping out broadcast storms. IGMP snooping restricts all multicast traffic to ports participating in an IGMP group. Storm control limits high-overhead protocols to an arbitrary rate; in this case, broadcasts are getting throttled at 250 packets per second.

Step-By-Step Instructions

The following configuration guidelines are compatible with Dell PowerConnect 33xx and 52xx series switches. The Cisco router configurations are compatible with most Cisco routers running IOS.

1. Disable IP-directed broadcasts on the router.

```
Router1> enable
Router1# configure terminal
Router1(config)# no ip directed-broadcast
Router1(config)# exit
```

2. Enable storm control on Dell PowerConnect switches.

For PowerConnect 33xx series:

```
console# configure
console(config)# port storm-control enable broadcast fastethernet
console(config)# port storm-control rate fastethernet 250
```

NOTE: Storm control is not supported in PowerConnect 52xx switches.

3. Verify that storm control is operating correctly.

For PowerConnect 33xx series:

```
console# show port storm-control

|      Port      | |Unknown| |Broadcast| |Multicast| |Rate[Packets/sec]|
------------------ --------- ----------- ----------- -------------------
FastEthernet 1/e1  Disabled    Enabled     Enabled           250
FastEthernet 1/e2  Disabled    Enabled     Enabled           250
FastEthernet 1/e3  Disabled    Enabled     Enabled           250
FastEthernet 1/e4  Disabled    Enabled     Enabled           250
FastEthernet 1/e5  Disabled    Enabled     Enabled           250
FastEthernet 1/e6  Disabled    Enabled     Enabled           250
FastEthernet 1/e7  Disabled    Enabled     Enabled           250
FastEthernet 1/e8  Disabled    Enabled     Enabled           250
FastEthernet 1/e9  Disabled    Enabled     Enabled           250
FastEthernet 1/e10 Disabled    Enabled     Enabled           250
FastEthernet 1/e11 Disabled    Enabled     Enabled           250
FastEthernet 1/e12 Disabled    Enabled     Enabled           250
FastEthernet 1/e13 Disabled    Enabled     Enabled           250
FastEthernet 1/e14 Disabled    Enabled     Enabled           250
FastEthernet 1/e15 Disabled    Enabled     Enabled           250
FastEthernet 1/e16 Disabled    Enabled     Enabled           250
FastEthernet 1/e17 Disabled    Enabled     Enabled           250
FastEthernet 1/e18 Disabled    Enabled     Enabled           250
FastEthernet 1/e19 Disabled    Enabled     Enabled           250
FastEthernet 1/e20 Disabled    Enabled     Enabled           250
FastEthernet 1/e21 Disabled    Enabled     Enabled           250
FastEthernet 1/e22 Disabled    Enabled     Enabled           250
FastEthernet 1/e23 Disabled    Enabled     Enabled           250
FastEthernet 1/e24 Disabled    Enabled     Enabled           250
```

4. Enable IGMP snooping on Dell PowerConnect switches.

For PowerConnect 33xx and 52xx series:

```
console(config)# ip igmp snooping
```

To discover which multicast addresses a PowerConnect switch has learned, use this command:

```
console# show ip igmp snooping groups
```

Conclusion

Excessive broadcast and multicast traffic can seriously degrade network performance. By following the configuration examples given here and disabling IP-directed broadcasts on edge routers, it is possible to mitigate the effects of excessive broadcast and multicast traffic.