Technical Guide

How To | Catering for Special Considerations in an IP Video Surveillance Network

Introduction

IP multicast networks for video surveillance have characteristics that significantly differentiate them from the video distribution networks that have traditionally been the common users of IP multicast.

The key differentiating characteristics of video surveillance networks are:

- They invert the traditional model of 'few sources, many listeners'.
- They can be quite sensitive to the flooding of IP multicast data and of IGMP signalling packets.

What information will you find in this How To Note?

In this How To Note, we will:

- Describe these key differences in more detail.
- Look at what implications these differences have for the operation of IGMP snooping devices in the network.
- Learn about the commands provided in AlliedWare Plus to configure Allied Telesis x-series switches to operate in a manner that is optimised for video surveillance.

Related How To Notes

Allied Telesis offers How To Notes with a wide range of video surveillance solutions.

For a complete list of video surveillance How To Notes:

- Go to: http://www.alliedtelesis.com/support/documentation
- Enter key word: **video**

Which products and software version does this apply to?

This How To Note applies to all **x-Series** switches, running AlliedWare Plus software version **5.4.3-3.13** or later.

Contents

Inversion of the 'few sources, many listeners'-model

IP multicasting was designed around a model in which one, or a few, source devices transmit multicast streams into a network, and many hosts in the network listen to one or more of those streams. So, the standard image of a multicast network has a few sources acting as the root of forwarding trees that have multiple listening hosts as the leaves.

In a traditional multicast network, a few multicast servers send streams (TV channels or video) into the network which are received by multicast clients (set-top boxes/decoders) at the edge of the network.

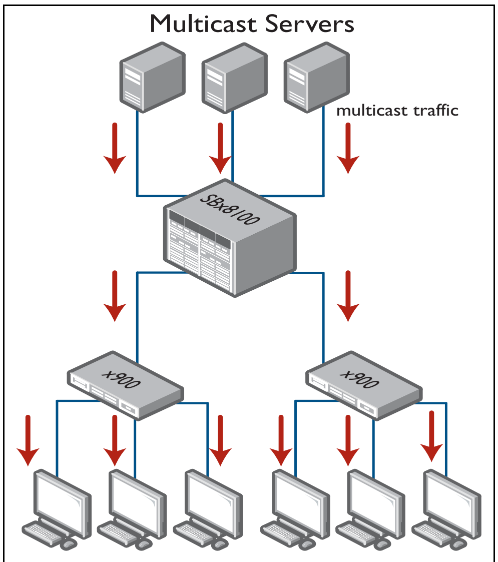

A traditional multicast network

An IP surveillance multicast network effectively turns that traditional model on its head. With IP video surveillance, multiple cameras at the edge of the network are multicasting streams (video) which are viewed by a monitoring/recording control center usually located at the center of the network. A small subset of all the cameras is viewed by the control center at any one time, usually cycling through all available cameras. And, frequently, all the cameras' streams are being recorded all the time.

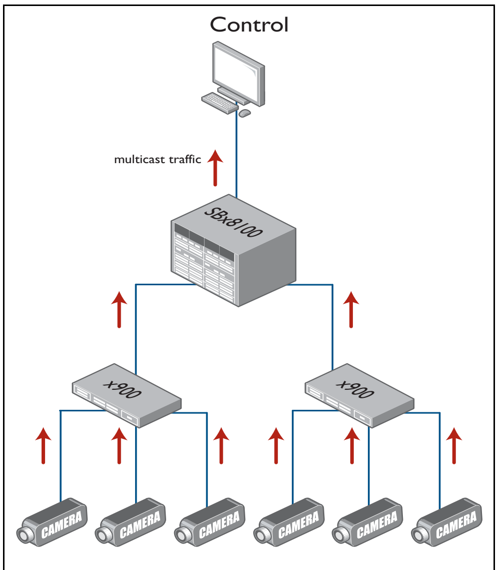

An IP video surveillance network

Preserving camera CPU processor capacity

By necessity, to keep costs and power consumption down, video surveillance cameras use CPU processors of modest capacity. The task of processing the incoming image data, and pumping it out as IP data, will typically consume most of the power of a camera's CPU. If it has the additional load of examining and processing unnecessary IP data coming into it off the network, then that can overload the CPU, causing deterioration of the quality of the video it is feeding into the network.

So, in an IP video surveillance network, care must be taken not to overload the cameras' CPU processors.

Design criteria for IGMP snooping devices

Certain design criteria have been quite standard in IGMP snooping switches for some years. The resulting three characteristic behaviors are:

1. **Flood queries right through a VLAN**: You never know where the listeners might be located, so ensure that queries are sent everywhere, so they can hunt out all the listeners.
2. **Try to free up hardware forwarding entries as soon as possible**: The switching chips that perform hardware forwarding of multicast have a limited number of Layer 2 multicast forwarding entries. Potentially, there may well be more multicast groups being transmitted into a LAN than there are Layer 2 forwarding entries in a switch's switching chip. So, the IGMP snooping process in switches is designed to free up any forwarding table entries that are not currently in use, so that if downstream listeners start requesting different groups, there are entries available to be used for forwarding those groups.
    - Hence, IGMP snooping switches that have not recently received IGMP reports for a given group will periodically remove the hardware forwarding entry for that group, so the CPU can check if any data is currently being received for that group. If the CPU sees no packets being received for the group in question, then it will permanently remove the hardware forwarding entry, as it is now serving no purpose.
    - This act of periodically removing the hardware forwarding entries for these unregistered multicast groups (for which the switch has not recently received any IGMP reports) causes the corresponding streams to be briefly flooded (if they are still being transmitted to the switch).
3. **Flood multicast by default**: Unless an IGMP snooping switch's switch chip contains a hardware entry to suppress the forwarding of a given group, or to forward that group to a certain set of ports on which reports for that group have been received, then it will flood data that is received for that group.

Implications for the operation of the IGMP snooping devices in the network

These three characteristic behaviors of IGMP snooping switches are not ideal in video surveillance networks because:

- Flooding IGMP queries throughout a VLAN means sending the queries to all cameras in a VLAN. This means that cameras receive unnecessary IGMP queries, the processing of which can put extra load in their CPUs, and disrupt their video transmission. Given that cameras will never be requesting to listen to video streams, there is no need to send queries to the cameras. Hence, in a video surveillance network, sending IGMP queries in any direction—except towards the recording and monitoring equipment— is unnecessary and potentially disruptive.

Freeing up hardware forwarding entries in edge switches is unnecessary.

- The edge switches will have a given set of cameras connected to them. These cameras will continue sending their video streams to the same IP multicast group address day in and day out. The edge switch in a video surveillance network has just a small set of multicast groups to forward, and these groups very rarely change.
    - This is quite in contrast to the situation for an edge switch in a video transmission network, where the listeners attached to the switch may request any number of different multicast groups, and will change between groups reasonably often.
    - So, edge switches in a video surveillance network can save themselves the bother of periodically removing the hardware forwarding entries for unregistered groups, to check if they are still receiving data destined to those groups. If the cameras transmitting to those groups are still attached and running, then the switch will still be receiving data for those groups.
    - By not periodically removing these forwarding entries, the switches avoid flooding data to other cameras.
- Flooding unregistered multicast is inherently not desirable in a video surveillance network for the very reason that it causes multicast data to be sent towards the cameras, thereby potentially disrupting the cameras' video transmission as their CPUs process these unnecessary packets.

Useful AlliedWare Plus commands

This section provides information on the commands addressing the multicast issues specific to an IP video surveillance network.

1. Allied Telesis AlliedWare Plus has a feature that globally disables the flooding of IGMP Group Membership Specific Queries to all ports in the VLAN that the group membership is registered in.

   When the flooding of Specific Queries is disabled, the Specific Queries will only be sent to the group member ports in the VLAN, i.e. those ports on which joins have previously been received. This means that the cameras will not receive the Specific Queries, which they do not actually need, and so avoid unnecessary processing load on their CPUs.

   - The global command is: **(no) ip igmp flood specific-query**

2. If the video streams are not being continuously recorded, and if the monitor is cycling through the different multicast groups at a relatively slow rate, so that the period between when it looks at each given group is longer than the timeout for the multicast forwarding entries, then the forwarding entries will time out. When a forwarding entry times out, the UDP data coming from the associated camera will be briefly flooded, as it will be unregistered multicast traffic at that moment. The flooding will be stopped quickly thereafter because IGMP snooping sees the flooded traffic (because the flooding also sends the traffic to the CPU), and will put a "*do not forward this group anywhere*" entry into the hardware table. However, even a brief period of flooding can disrupt the operation of the cameras.

   - The global command to prevent this flooding is: **platform stop-unreg-mc-flooding**

   This command will also stop the multicast flooding when a new group is learnt (as the hardware entry is just about to be added). This applies for x210, x230, x310, x510, and x610 Series switches.

3. The switch can also be configured to never timeout these entries.

   - The global command is: **ip igmp snooping source-timeout 0**

   This command will stop the group timing out once it is learnt, until it is cleared manually.

   Then, if data arrives again, the **platform stop-unreg-mc-flooding** command will come into action again (if also configured) to stop the new group from being flooded. This command can be applied globally, as well as per VLAN.

Other considerations

Reducing STP topology change events

In a network where spanning tree is being used, port state changes can have a surprisingly disruptive effect on the network. If a port is an active spanning tree port, and it changes state, then the spanning tree network can see this event as a **Topology Change**. As a result, Topology Change Notifications(TCNs) are sent through the network. When a switch receives a TCN, it needs to prepare itself for the possibility that traffic it was receiving on some ports may now arrive on other ports (as the active paths in the network may have changed). Therefore, the switch must clean out its MAC forwarding database (FDB), and its ARP table, ready to relearn entries potentially via different ports.

This means that some, or often all, switches in the network suddenly have empty FDB and ARP tables, resulting in a lot of data being flooded, and delayed. If the network is busy, this will result in a brief period of congested links, and dropped packets. In a video surveillance network, dropped packets are highly undesirable, as a dropped packet in a video stream is dropped forever, it is not retransmitted.

Therefore, it is very important to take measures to reduce the incidence of STP topology change events in the network. By far the most effective such measure is to ensure that all *edge* ports in the network (ports that are not attached to other switches, but are attached to cameras, recorders, workstations, etc.) are configured with portfast. This configuration gives the ports a special STP property whereby state changes on those ports are not treated as STP topology change events. Thereby, cameras can be connected, disconnected, rebooted, etc. without causing TCN-initiated disruption in the network.

The relevant commands are:

```
awplus# configure terminal
awplus(config)# interface portx.y.z
awplus(config-if)# spanning-tree portfast
```

Core switch resiliency and link resiliency

Frequently, the design of a video surveillance network will involve core or distribution layers, where connections from access switches are concentrated together. These concentration points are key elements in the network. Failures at these nodes have a high impact on the operation, as such failures will cut off the video streams from large numbers of cameras.

Therefore, it is highly recommended to build in resiliency at these points by populating them with stacked pairs of switches, rather than single stand-alone switches. The Allied Telesis x-series switches provide a highly reliable stacking feature known as Virtual Chassis Stacking (VCS), that enables two (or more) switches to operate as a single entity, in an active-active formation. If one of the switches in the stack fails, or loses power, the other stack members will continue to operate, with barely noticeable traffic disruption.

To take full advantage of the VCS resiliency, it is necessary to use aggregated links from the stacked switches to the access switches. By aggregating together links that terminate on different stack members, connectivity between the stack and the access switch is maintained even if one stack member fails. Do note that using aggregation for this link resiliency is far more desirable than using spanning tree to provide the link resiliency because:

- Link aggregation is an active-active process. All links in the aggregation carry traffic. When using spanning tree, only one link in the bundle is active, and all others are blocked.
- Failover is very rapid when a link in an aggregation goes down. Spanning tree can take some seconds to fail over.
- No topology change event is generated when a link in an aggregation fails.

C613-16181-00 REV A

Allied Telesis™ the solution : the network

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com