How to configure resilient Layer 2 multicast networks using IGMP

Introduction

This document describes the recommended configuration for designing a resilient network for multicast distribution across a Layer 2 infrastructure.

Many TV streaming networks in environments such as apartments, hotels, and cruise liners, are required to operate with very high reliability. A failed core switch in the network can cause failure of a large number of edge connections and therefore a resulting loss of revenue. The scenarios below show a typical basic multicast distribution network, with a suggested configuration, then a resilient multicast network using IGMP, with a suggested configuration.

Switch configurations

The scenarios in this document use AT-9812T switches as core switches, and Rapier 24*i* switches as edge switches. You could also use an AT-8700 series switch as an edge switch, and a SwitchBlade as a core switch.

- Rapier or AT-8700 series switches should be running software release 2.6.1 with the latest patch.
- AT-9800 series or SwitchBlade core switches should be running software release 2.5.1 patch 13 or above.

Basic multicast distribution network

In the basic multicast distribution network shown in Figure 1, multicasts are launched in by the servers. Clients on the edge switches will join and leave the groups via Internet Group Management Protocol (IGMP) v2.

Figure 1: A typical multicast distribution network

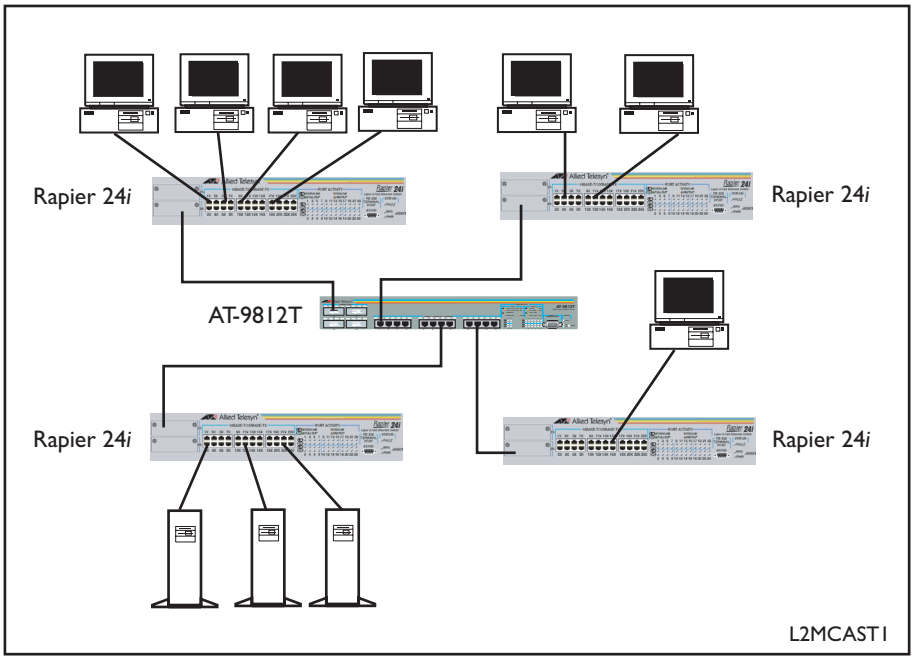

The examples in this document use AT-9812T switches as core switches, and Rapier 24i switches as edge switches.

To control where the multicasts go, the core switches are running IP IGMP and the edge switches are running IGMP snooping. The core switch periodically polls all the devices on the network with an IGMP general query. Any hosts that are joined to a group respond with a report. By 'snooping' these IGMP exchanges, the edge switches are able to maintain an up-to-date table of which port requires which multicast. If a multicast is not required by a host then it will not be sent to that switch.

Configuration of the switches in this case is quite simple:

Configuring the core switch (AT-9812T switch in the scenario)

```
ena ip
add ip int=vlan1 ip=192.168.1.1
ena ip igmp
ena ip igmp int=vlan1
```

Configuring the edge switches (Rapier 24*i* switches in the scenario)

```
ena ip
add ip int=vlan1 ip=192.168.1.13
```

The edge switches in this configuration example are running IGMP snooping by default.

No other configuration should be necessary in this scenario. All multicasts launched into the network are forwarded back to the querying switch. This happens whether the server joins its own group or not. The operation is cleaner if servers do join the group they are sending on but some do not. Diagnostics will show if this is happening, as you will see unregistered groups when looking at the output of the `show igmpsnoop` command.

This system works well and is the most cost effective solution. However, if there is a high up-time requirement, then the event of a core switch failing will cause all the edge switches and clients to lose service until this is repaired.

Setting up a resilient multicast network using IGMP

Further to the previous scenario, you can add a second core switch to the system and use Rapid Spanning Tree Protocol (RSTP) to provide a backup link should one of the core switches fail, as shown in Figure 2.

Figure 2: A Resilient Multicast Network Design

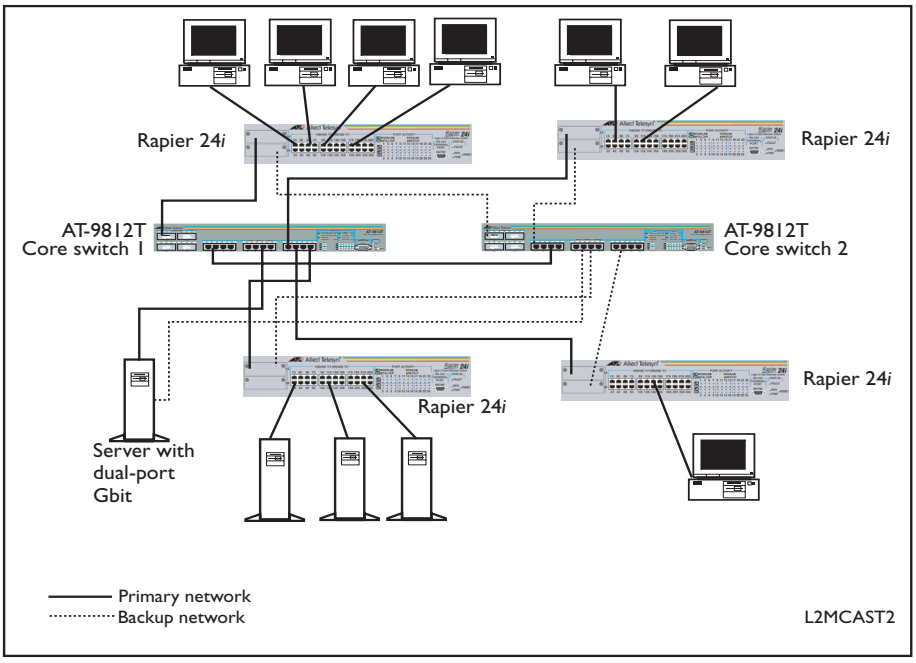

In the configuration shown in Figure 2, there are two core AT-9812T switches and each Rapier 24*i* edge switch is dual-homed to each of these. The servers could be moved to connect directly to the Gbit core switch and use a dual port Gbit network adapter with link fault tolerant drivers (e.g. AT-2970T2). Rapid Spanning Tree Protocol (RSTP, 802.1W) is used to control which of the links between the switches is passing traffic.

In the event of a failure, RSTP quickly recovers network connectivity, normally not losing more than one ping on a test. It can do this more quickly than conventional STP (802.1D) which has to go through listening and learning phases before ports will start forwarding traffic. This can take up to 50 seconds with default settings. RSTP sets the blocking ports to a standby mode of 'alternate' and starts using them as soon as a failure is detected on the main 'designated' port.

The next step to recovering multicasting is that the IGMP has to relearn where to send the streams. This requires a general query from the IGMP switch. Depending on which switch fails, it may be the existing querier or a new querier that takes over this function. In the configuration for IGMP resilience, set the query timer much lower than the defaults, normally around 15 seconds, with a 40 second timeout. This allows all the multicast throughput to be recovered within 15-30 seconds, depending on how soon after the failure the next scheduled IGMP general query occurs. It is possible to set the IGMP timers lower, but this recovery time is usually acceptable.

Configuring core switch 1 (AT-9812T switch in the scenario)

```
Ena ip
Add ip int=vlan1 ip=192.168.1.1
Ena ip igmp
Ena ip igmp int=vlan1
Set ip igmp query=15 timeo=40
Set stp=default mode=rapid prio=9000
Ena stp=default
```

Configuring core switch 2 (AT-9812T switch in the scenario)

```
Ena ip
Add ip int=vlan1 ip=192.168.1.2
Ena ip igmp
Ena ip igmp int=vlan1
Set ip igmp query=15 timeo=40
Set stp=default mode=rapid prio=18000
Ena stp=default
```

Configuring edge switches (Rapier 24*i* switches in the scenario)

```
Ena ip
Add ip int=vlan1 ip=192.168.1.x
Set ip igmp query=15 timeo=40
Set stp=default mode=rapid
Set stp port=1-24 edgepo=yes
Ena stp=default
```