Ethernet 1:1 and 1:N Redundancy Strategy for Simple Bridged Modems

R. Montgomery

Revision 1 – May 2006

Background

In satellite communications, it is desirable to have standby modems available to switch into an active configuration when an online modem experiences a failure.

In redundant configurations, it is customary to have an external switch between the modem or modem bank and RF equipment to properly switch over the satellite connections from a failed online modem to the modem that is coming online to replace the failed modem. All modem settings are also transferred. Two types of modem switching are currently supported by CEFD: 1:1 and 1:N redundancy. In 1:1 redundancy, a single modem is provided as a backup for a single online modem. In the 1:N configuration, multiple online modems coexist with a single modem that is ready to switch into any position, in case one of the online modems experiences a fault. This solution has existed for years and works nicely with the traditional terrestrial data interfaces supported by CEFD modems. This whitepaper assumes general understanding of existing CEFD redundancy designs.

With the introduction of an Ethernet data interface into modems, a new problem arises with regard to redundancy. This problem exists on the Ethernet side of the modem—multiple modems will need the Ethernet stream feed. In the event of a redundant switchover, the secondary modem can immediately begin forwarding the stream. Given this requirement, an aggregation point feeding all the modems needs to exist. The simplest method for providing this aggregation point is with a hub or switch. Of course, this doesn't completely solve the problem. With the introduction of an aggregation point, care must also be taken not to introduce circuit loops, multiple IP frame copies, and multiple responses hitting the local Ethernet segment.

In the Ethernet packetized world, all packets carry an address that identifies the next piece of hardware for which the packet is destined. This address is known as the Media Access Control (MAC) Address. An Ethernet enabled device must know the MAC of the next piece of equipment ready to receive the packet (next HOP). Most devices maintain an Address Resolution Protocol (ARP) table that maps the destination IP Address to the next hardware MAC address. Once a sending device populates an entry, the mapping is maintained until no traffic is destined to that IP address for a given amount of time (timeouts are typically 4 to 5 minutes, but vary across devices). Once the time expires, the entry is flushed from the table. If a new packet arrives that needs to be sent to the IP Address again, the system will re-ARP for the MAC address and place the mapping back into the ARP table, and the cycle continues.

Let the piece of equipment sending traffic to the current online modem be defined as the *upstream* piece of equipment (computer, router, firewall, etc). As pictured in Figure 1, the Ethernet switch is needed between the upstream equipment and the modem bank for aggregation purposes, as described previously. The upstream piece of equipment, following standard ARP requirements, maintains an ARP entry that identifies the MAC address of the next device or target device.

In this very specific case, we are considering the link to be working in bridging mode. Picture the end-to-end modems, including the satellite link, to be the equivalent of a very long wire or Ethernet cable. Therefore, as depicted in Figure 1, PC1 (upstream device) has determined the proper MAC address for the target PC (PC2), and entered the mapping into its local ARP table. Conversely, PC2 has identified PC1 in the same manner and has recorded a similar entry as well.

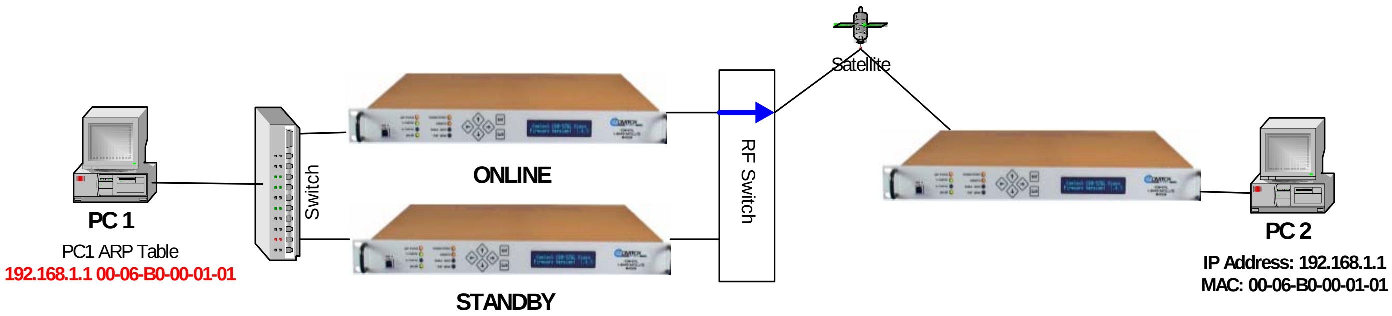

Note 1: The switch has mapped PC2's MAC (00-06-B0-00-01-01) to port 1
Note 2: The switch has not mapped anything regarding the standby unit, which is plugged into Port 2.

Figure 1 Normal Operation

As noted in Figure 1, the Ethernet switch's standard operation must be considered. The Ethernet switch maintains a table mapping MAC addresses to switch ports, known as the dynamic Content-Addressable Memory (CAM) table. When a packet arrives, and an entry for a specific destination MAC is unknown to the switch, it sends the packet out all ports and waits for response packets. This is an attempt to identify the port on the Ethernet switch where a particular MAC is externally connected. Once the device responds, the switch is able to map the MAC to the port in the CAM table. Any new packets arriving destined for that MAC address will only be forwarded out the port identified in the CAM table. Most Ethernet switch CAM tables do have a flush or timeout value, but are normally set to a very high number to limit the amount of times the switch has to broadcast a packet out all ports. It is also important to note that CAM entries for a specific port are cleared when a port link goes down.

Problem Statement

Finally, given these key concepts, we identify the problem with redundancy on the Ethernet side of modems. When a redundancy switchover occurs, the satellite link is broken and switched to the standby modem, as expected. However, the Ethernet switch will continue to send packets to the original port because its CAM entry is still mapped with the destination MAC residing out of port 1. This is the basic problem. The Ethernet switch has no way of knowing to now send packets out of port 2. This is illustrated in Figure 2. It is accurate to say that the switch would *eventually* figure out the change and update the CAM table because of packets coming from PC2 destined to PC1, but the goal is not to wait for this to occur before beginning to pass traffic. The redundant switchover needs to occur as seamlessly as possible.

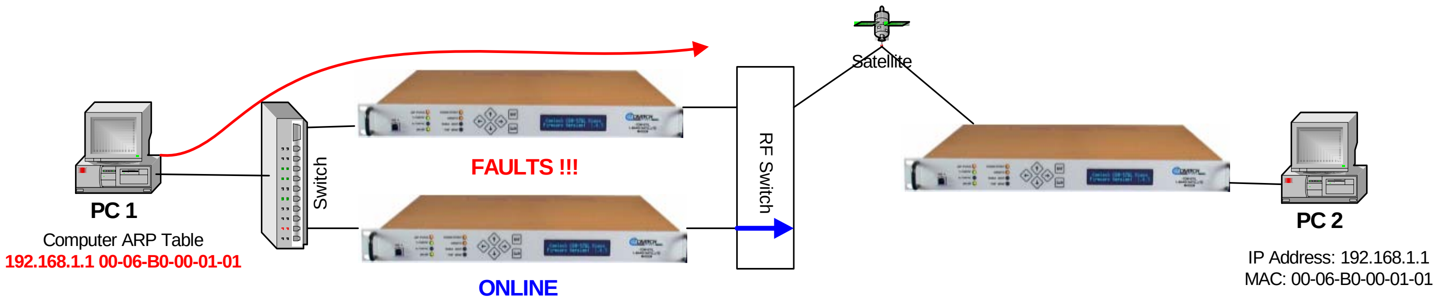

Note 1: The switch still has mapped PC2's MAC (00-06-B0-00-01-01) to port 1
Note 2: The RF has been switched to the second modem, but IP traffic won't go out Port 2 on switch.

Figure 2 Redundant Switch Has Occurred

Solution

To solve the problem on the Ethernet side, the strategy is still to rely on the Ethernet switch to quickly and seamlessly remap the MAC and port in its CAM table. Please note, at this time we are not discussing a hub – a true switching device is required.

As described above, the switch will properly clear CAM entries for a given port when the link status of a port is taken down. When a modem is taken offline, the modem will turn off its Ethernet port electrically. This will force a dropping of the link on the Ethernet switch and will force the Ethernet switch to relearn the associations between the port and far side MAC(s). At this time, any packets will be sent out all ports and responses will now be seen on the new port (Port 2 in this example), and the Ethernet switch will update the CAM table.

In short, the mechanism for upstream Ethernet redundancy is to use the learning and port management characteristics of a typical off-the-shelf switch.

Notes

1. This method of redundancy is intended to be equivalent to pulling the Ethernet cable from one port and putting it into a different port on the same switch.
2. While the Data Ethernet port will be unavailable when offline, the M&C Ethernet port will still be functional and available for command and control.
3. Installation must use a true switching device, not a hub device.
4. This analysis and solution only pertains to modems in a bridging mode.