

A Brief Note on Redundancy and Fail-over

Introduction

The RAVE family of products support failsafe operation by providing at least two paths for an audio signal to gain access to the network. This is normally accomplished by backing up a primary RAVE product with a secondary unit. During normal operation, the primary RAVE will remain active while the secondary unit is placed in “standby”. Should the primary unit experience a power failure or loss of Ethernet, the secondary RAVE will come out of “standby” and assume the active role.

The redundant features of RAVE were designed when CobraNet support was limited to repeater networks only. As such, the goal was to provide a means for duplicating the role of the transmitter. On a repeater network the audio is multicast to all locations on the LAN. Any receiver wishing to listen to a particular audio channel may do so simply by tuning into the proper bundle address. However, there may be only one active unit transmitting on any given *public* bundle.

Repeater Networks

On repeater networks failsafe operation can be provided by backing up both the transmitter and receiver for any bundle assignment. This is done via the “Sync Output” and “Slave Input” BNC connectors on all RAVE rear panels.

When providing a redundant unit for the transmitter, connect a 50-ohm coaxial cable from the “Sync Output” of the unit that is to be the normally active device to the “Slave Input” of the unit that is to be the normally passive device. Using RAVE’s front-panel rotary switches, duplicate the outgoing bundle number(s) exactly on both transmitting units. From the analog audio source, connect a “Y” cable so that input channel 1 of the primary and secondary units are both receiving the same audio signal. If the source is a mixing console, you may be able to drive the same signal from two separate outputs, one driving channel 1 of the primary RAVE and one driving channel 1 of the secondary RAVE. Repeat the audio connections for each input channel as needed.

** Note that digital audio signals may not be combined. RAVE AES inputs require independent audio feeds from appropriate AES transmitters or a multi-channel distribution device.*

With redundant configuration, the normally passive unit will not have audio access to the network until there is a loss of system clock on the normally active unit's "Sync Output". When the secondary unit senses the loss of clock it will assume the active role. If the primary unit has failed, or its Ethernet connection is faulty, the problem can now be serviced without the loss of audio distribution. In the case of a temporary outage, possibly due to a port failure or delayed link, the return of a valid clock to the secondary RAVE will re-establish the original roles. Note that an intermittent port or flaky cable can cause the two units to toggle between the active and passive roles. This is a rare exception and may require bringing down both units to troubleshoot the fault.

When providing a redundant unit for the receiver, connect a 50-ohm coaxial cable from the "Sync Output" of the normally active unit to the "Slave Input" of the normally passive unit. Using the RAVE's front panel rotary switches, duplicate the incoming bundle number(s) exactly on both receiving units. Note that it is not possible to safely parallel the outputs of the RAVE and RAVE "s" series products. A provision has been included on the "s-24" series RAVEs to accommodate paralleling the audio outputs. On pre "s-24" products, connect each audio output from the primary and secondary unit to an independent console or amplifier input. Optionally, use a splitter box (selectable 2-to-1 multiplexer) prior to driving an amplifier input. To complete the fail-over process, the splitter or console input signal must be switched over from the primary to the secondary path.

** Note that digital outputs may not be combined. RAVE AES outputs must source to independent destination inputs.*

Switched Networks

With switched networks, there are a number of ways to provide for redundant and robust operation. All of the methods employed with repeater networks also apply with switches. However, there are additional methods that make use of CobraNet's advanced version 2 protocol and features that may be available with switch hardware.

The main difference for incorporating redundancy with RAVEs on network switches is that it's now possible to employ unicast addressing of audio bundles. In order for a unicast connection to be established between two RAVEs, there may be only one transmitter and one receiver occupying any given bundle assignment. Therefore, it's possible to assign two transmitters and two receivers the same unicast bundle assignment yet only one pair will be active. If either the transmitter or receiver should fail, a redundant unit will assume the role. Unlike the use of the "Sync" connectors, the unicast addressing method does not rely on the system clock and the redundant unit need not reside next to its Master. The location of the units can be anywhere in the system so long as they have access to the same program material. The toggle effect is also eliminated so

long as both units have the same bundle transmit/receive priority. This is because the conductor grants transmit and receive permission based on a first-come first-served policy. Of course custom configuration allows transmitters and receivers to adopt higher levels of priority access (refer to the CobraNet Technology Datasheet on the Peak Audio website). The significance here is that an intermittent unit or flaky cable can remain wired in the rack until servicing is convenient.

Another benefit of using unicast addressing to handle failsafe operation is that a single RAVE can backup two independent units. It's possible to use a RAVE 161s to backup the transmitters of two RAVE 188s units. Likewise, a RAVE 160s can backup two output bundles of two independent RAVE 188s units. Depending on how a system's bundle configuration is setup, there may be many possibilities in providing redundancy for failsafe operation.

Thus far, the discussion has outlined redundancy of RAVE devices on network switches by implementing unicast addressing in lieu of clock signals. In this application, predictability of the active and standby roles upon power up is not guaranteed. The unicast pair that assume the active role will be those that have come on-line first. A system arbitration or reset event may change these roles, however. When operating the network strictly via the RAVE's front panel switches, it is recommended that the BNC connectors be used. This will assure predictability of the "active" and "passive" pair. However, one unit can no longer effectively serve as a backup to multiple RAVEs due to relying on a single and disabling clock output. A solution to the predictability issue then is to use the management interface. Via SNMP, it's possible to set the priority of each unit on the network to suit the needs of the system. Additionally, this can be done remotely and dynamically after the system has been installed. By configuring the normally active unit's transmit or receive priority to be greater than that of the normally passive unit, the conductor is forced to grant permission based on priority levels rather than the first-come first-served policy.

If redundancy on network switches is needed for receivers implementing multicast addressing, the method incorporating the clock output and BNC connectors should again be used. This is due to the complexity of trying to control audio reception among multiple receivers through prioritization. The procedure for setting up redundancy using RAVE's rear panel connectors is the same as that outlined for repeater networks. Note also that connecting all audio I/O for redundant operation is the same as that outlined for repeater networks above.

Failsafe Network Design

Additional failsafe options may be available depending on the network switches implemented in the LAN system design. Use of the "Spanning Tree Protocol" (STP) can provide a means for connecting a network in a ring topology. This allows multiple paths between Ethernet switches. With the use of STP, only one active path will be enabled. When the switch senses the loss of connection between two devices on the network due

to an inter-switch link, STP will enable a secondary route, if one has been provided. STP can be very useful when there are multiple switch hops between devices.

Other features available with some switches are “trunking” and “meshing”. Trunking can provide multiple paths between switches. This is like having several connections acting as a large switch uplink. The advantage is that there is a larger pipe for data throughput. The maximum bandwidth is the sum of all connections assigned to the “trunk”. Trunks provide redundancy because one or more of the connections may fail and the workload can be absorbed by the remaining, albeit smaller, trunk. Meshing allows two or more switches to share their aggregate bandwidth. This is usually done through dedicated uplink connections. Redundancy may be provided here in that a switch may partially fail or become overloaded while another absorbs the workload.

Other methods for providing network robustness and failsafe operation include multiple power supplies on switch hardware, redundant port cards, overlapping layer 2 filtering and more. There are a number of ways in which to setup the switch hardware with managed devices. Filtering and partitioning may enable the system Administrator to remotely enable selective links. Providing the smallest amount of point-to-point links to isolated VLANs also helps to provide a robust network. Multiple VLANs means multiple “conductors” and less reliance on a single system master for the entire network. If one VLAN dies or must re-arbitrate for the “conductor” role, the rest of the network is not affected. The audio that is affected during the fail over would then be isolated to one local zone.