

## Guidelines for Configuring a RAVE Switched Network

The information presented in this paper should be used as general guidelines only. The complexity of a large network with multiple forms of traffic can be overwhelming. We wish to address here only the concerns for setting up a layer-2 Ethernet switch and RAVE™ products for use on a switched network. Additional documentation available from QSC Audio and Peak Audio provide information on the CobraNet™ Discovery application as well as getting started with simple network management protocol (SNMP).

All of the general guidelines regarding switched network design, as outlined on the Peak Audio website, still hold true. The Peak Audio website and CobraNet specifications should always be considered the authoritative source for CobraNet network design. The following is a quick checklist and reference of configuration rules for setting up a RAVE network for operation on Ethernet switches.

- All RAVEs must connect directly to their own switch port. No repeater hubs may be placed between a RAVE device and a switch port. If an in-line media converter is required for use with fiber optic cable, the media converter must be an intelligent switch such that the configuration is consistent with the no repeater rule.
- All ports on the Ethernet switch that support RAVE devices must be set to advertise themselves as all-capable via auto-negotiation. It is through the auto-negotiation process that RAVE products determine what type of network device they are linking to. Forcing a switch port to operate at 100 Mbps will defeat this auto-sensing feature and trick the RAVE device into linking in the repeater mode. This will result in an unreliable connection to the switch.
- Although multicast traffic is tolerated on a CobraNet network while operating on Ethernet switches, we recommend that unicast traffic be used wherever possible. Currently, multicast bundles should be limited to a total of 4 per LAN or virtual local area network (VLAN). If the system requirements call for a large amount of multicast traffic in specific areas, it may be possible to split the network into smaller segments thereby isolating traffic into smaller concentration zones. This can be done by configuring multiple VLANs on a managed switch or by segmenting the network with multiple discrete non-managed network switches. Note that all units can still

reference a common conductor clock by slaving multiple LANs or VLANs to the primary conductor clock using the “external synchronization” feature with RAVE.

CobraNet bundle assignments dictate which delivery method is implemented. When using the RAVE’s front-panel rotary encoding switches, positions **10** through **FF** are reserved for unicast bundles. Positions **01** through **0F** are reserved for multicast bundles. Refer to the reference document on RAVE bundle assignments for further information on setting the front-panel rotary encoding switches.

- RAVE devices that are to communicate with one another must be assigned to switch ports on the same LAN or VLAN. This refers to a common broadcast domain. Devices must belong to the same broadcast domain in order to communicate with each other. Most network switches ship out-of-the-box with a default VLAN for the entire switch. If the system requirements call for splitting the switch into multiple VLANs, care must be taken to ensure that the appropriate devices retain communication. Generally, each VLAN will have a unique Internet Protocol (IP) address associated with it. If software control and monitoring of the RAVE is required, each RAVE must be assigned an IP address within the *same subnet* as that of the VLAN to which they belong.
- Best results are achieved through isolation of the RAVE audio network. Although not a requirement, isolation can eliminate the headaches associated with data communications broadcasts and unregulated traffic. Segmenting the network into VLANs, each associated with a common data type, provides isolation. For example, the system may be segmented into a RAVE audio VLAN, a video VLAN, a QSControl VLAN and a computing VLAN.

The question of whether to split a network into multiple VLAN segments may depend on the size of the overall network, whether other forms of traffic exist on the same network hardware with CobraNet audio and what type of system management and fault tolerance is employed. In addition to elimination of problems associated with data communications on the network, robustness can be built into the system by operating RAVEs on an isolated VLAN. This is because a managed switch provides the ability to configure the audio network for the highest priority of packet delivery in the event that bandwidth is exceeded or an inter-switch link becomes severed. With managed switches it is also possible to assign a dedicated block of bandwidth to support the RAVEs. This can virtually guarantee that saturation will not occur on the audio segment.

If the decision is made to split the network into multiple VLANs, it may be beneficial to split the RAVE audio network further into multiple segments. When using a point-to-point (PTP) architecture with unicast delivery between RAVE devices, it’s possible to assign a separate VLAN for each PTP segment. The advantage to this is that it is possible to service equipment on one segment without bringing down the network or affecting any other units on neighboring segments. Additionally, each segment will have its own CobraNet conductor. So bringing new units on-line or

taking units off-line have less affect on the network in regards to audio dropout. There are some disadvantages to segmenting however. Control and monitoring become difficult due to the isolation of units from a central control point. Additionally, network management increases with each additional partition and the toll on network resources increase if dedicating a defined amount of bandwidth to each RAVE VLAN. A compromising alternative is to segment the audio network into multicast groups. Segmenting 4 multicast transmitters and a handful of unicast transmitters into a VLAN group may prove to offer the best in efficiency. A central management console can then belong to multiple VLANs for the purposes of control and monitoring.

- RAVE devices have the ability to support two CobraNet bundles coming into the unit and two leaving the unit. This effectively allows for control of 16 audio channels in and 16 audio channels out via CobraNet. However, the RAVE hardware only supports up to 16 channels of physical audio I/O. The hardware configuration of the RAVE product line is as follows:

*RAVE 81s and 161s/s-24:* 16 physical audio channels into the box, which are split into two bundles transmitting via CobraNet.

*RAVE 80s and 160s/s-24:* receives two CobraNet bundles, which are routed to 16 physical audio outputs.

*RAVE 88s and 188s/s-24:* 8 physical audio channels into the box, which may be transmitted on one or two CobraNet bundles. The 88s and 188s/s-24 also receives one or two CobraNet bundles and selects up to 8 audio channels to route to its 8 physical outputs.

Obviously the RAVE 88s and 188s/s-24 are the most versatile of the product line.

It's important to keep in mind that a RAVE operates internally in half-duplex mode. In all default configurations the half-duplex issue poses no problems. However, there are some conditions to be aware of when configuring the RAVEs for operation on network switches.

The first condition to be aware of occurs with the RAVE 81s and 161s/s-24. When these units are transmitting two bundles onto the network the level of bandwidth utilization is considerably more than other RAVE devices. We therefore recommend that units in this configuration not hold the conductor role, as the conductor is responsible for servicing bundle requests from all transmitters and receivers as well as delivering the system sync. As a comparison, a RAVE 161s-24 acting as the system conductor *and* delivering 2 bundles *and* responding to management requests can be consuming nearly 30 Mbps bandwidth. In contrast, a RAVE 160s-24 can be consuming as little as a few hundred bytes per second. It should be noted that the conductor priority structure on RAVE products ensures that the RAVE 81s and 161s/s-24 devices are the least likely to become conductor. But these units are not

prevented from becoming the system conductor. However, the conductor role can be managed appropriately (forced) via the management interface using SNMP.

The second condition to be aware of occurs when a RAVE 88s or 188s/s-24 is configured via SNMP to transmit on two separate CobraNet bundles. This is done by copying audio channels from the default bundle onto a second out-going bundle. In this case, the 88s or 188s/s-24 unit should also avoid holding the conductor role.

Most of the small to medium sized networks still employ hardware control with use of an SNMP supporting software application as a monitoring tool. Issues with audio dropouts due to multiple transmitters and over utilization become more of a concern with larger network diameters, large amounts of multicast bundles and invoking additional audio processing in RAVE through software configuration. In these larger systems it is a simple process to manage conductor variables in order to prevent a unit(s) from assuming the conductor role while leaving a large number of conductor-capable devices available for a reliable system implementation.

- RAVE products must contain the correct CobraNet firmware version for the application intended on network switches. QSC Audio has shipped CobraNet versions 2.5.16, 2.6.9 and 2.8.5 firmware in the RAVE product line. Version 2.5.16 supports operation on network switches and provides an added measure of safety when upgrading new code. However, CobraNet versions 2.5.16 and prior do not support SNMP. CobraNet versions 2.6.9 and above do support access to management interface (MI) variables via SNMP as well as support for increased network diameters. However, version 2.6.9 does not provide a means for permanently saving user-configurable parameters in the event of a power loss or reset cycle. The “persistence” feature, which allows saving user-configurable parameters, was first shipped with CobraNet version 2.8.5 in RAVE products. Refer to documents on the QSC Audio and Peak Audio websites regarding firmware revision history and feature set.