Introduction to IP Multicasting

VBrick Video Network Appliances support IP multicast for the transmission and reception of MPEG audio/video. This paper provides a brief overview of the concepts and techniques used to deploy video in a multicast-enabled network.

Unicast

In conventional networking, one computer connects to another computer and that one-to-one relationship continues for the duration of a session. There are well known mechanisms for computer "A" to discover where computer "B" is located, and once that address is known, computer "A" simply sends its traffic to computer "B's" address and the network delivers it.   In this conventional Unicast world then, computers talk to computers using end point network addresses. Example: A VBrick with address 172.16.2.101 sends traffic to a VBrick with address 172.16.2.102. The network delivers the traffic between these two points.

Multicast

In Multicast networking, a computer does not send its traffic to another computer. Rather, it sends it to a special address whose physical location is actually inside a router and/or switch. A multicast client can then inform the router that it wants to receive a multicast stream. When so informed, the router replicates the traffic to that client (and to all other clients who inform the router of their desire to "join the session").

The illustration shows that a VBrick with an IP address of 172.16.2.101 is sending TO multicast address 225.1.1.1. VBrick 172.16.2.102, VBrick 172.16.2.103, and a PC with an address of 172.16.2.104 have each sent a message to the router to join the multicast session. Upon receipt of this message, the router replicates the traffic and sends it to each receiving device.

Leaving the Multicast

The network should stop sending the traffic to any device that no longer desires it, or who is no longer connected to the network. But the router must continue to send the multicast traffic it to all other members of that multicast group. It is important that the multicast stream ceases to be delivered to a device that no longer wants it, as it would otherwise adversely affect that device's ability to receive other

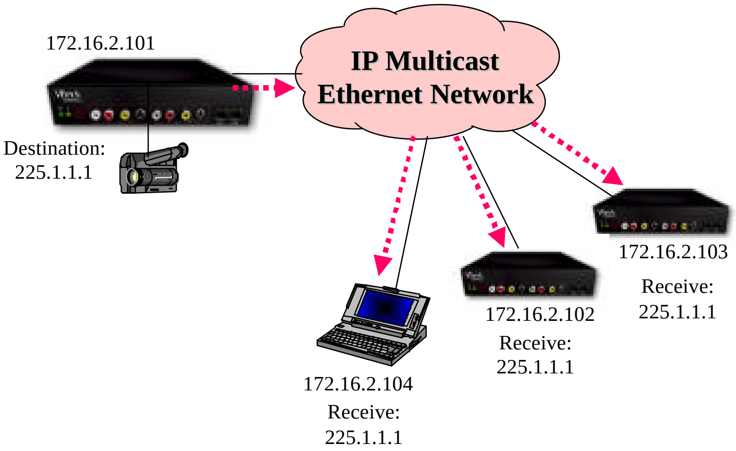

streams or to engage other data traffic. Routers periodically ask the members of a multicast group if they are still active, and will stop sending multicast traffic to a client when either of two things occur:

1. The client sends a message to the router telling it to stop sending a particular multicast stream (“leave the session”), or

2. The client no longer responds to the router’s periodic poll, in which case the router stops forwarding the multicast traffic.

Multicast Standards

While the basics of multicasting are easily understood, a new set of routing protocols are needed for a multicast network to scale to global size. Several routing standards are available, and they are summarized here.

Internet Group-Membership Protocol (IGMP) is the multicast protocol. IGMP relies on Class D IP addresses for the creation of multicast groups (defined in RFC 1112). IGMP is used to dynamically register individual hosts in a multicast group with a Class D address. Hosts identify group memberships by sending IGMP messages, and traffic is sent to all members of that multicast group. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on particular LANs. Routers communicate with each other by using one or more protocols to build multicast routes for each group.

Protocol Independent Multicast (PIM) sparse mode is optimized for internetworks with many data streams but a relatively few number of LANs. It defines a rendezvous point that is then used as a registration point to facilitate the proper routing of packets.
When multicast traffic is sent, the router nearest the source sends data to the rendezvous point. When a receiver wants to receive data, the last-hop router (with respect to the receiver) registers with the rendezvous point. A data stream flows from the sender to the rendezvous point and to the receiver. Routers in the path optimize the path and automatically remove any unnecessary hops, even at the rendezvous point.

The Multicast Open Shortest Path First (MOSPF) is an extension of OSPF. MOSPF employs a unicast routing protocol so that every router is aware of all available links, and each calculates routes from the source to all possible group members. MOSPF calculates the routes for each source/multicast group pair when the router receives traffic for that pair, and routes are cached until a topology change occurs. MOSPF then recalculates the topology. MOSPF works only in networks that use OSPF, and where there are a small number of multicast groups. Also, MOSPF takes up a lot of router CPU bandwidth when there are many multicast groups, or where those groups change often.

Beyond the standards, many Ethernet switch and router vendors have developed features to assist in the roll out of multicast services. For example, Cisco provides broadcast/multicast suppression that prevents switched ports on a LAN from being disrupted by a broadcast storm. A broadcast storm occurs when broadcast or multicast packets flood the LAN, creating excessive traffic and degrading network performance. Errors in the protocol-stack implementation or in the network configuration can cause a broadcast storm. Cisco, and other vendors provide a means to control the multicast bandwidth on any switch port.

ReCasting & Tunneling

In certain cases, a source or destination may not be attached to a multicast-enabled network. For example, a branch office with multiple receive clients may be connected via a WAN connection to headquarters. If more than one client at the branch office were to request a stream, the network might attempt to multicast the stream at the headquarters (sending multiple copies of identical streams), rather than sending one stream and multicasting it at the branch office. An effective solution is to unicast from the headquarters to the branch office, and use IP ReCaster at the branch office. ReCaster accepts unicast streams and converts them into multicast. This is similar to tunneling, where a multicast stream is carried within a unicast session.