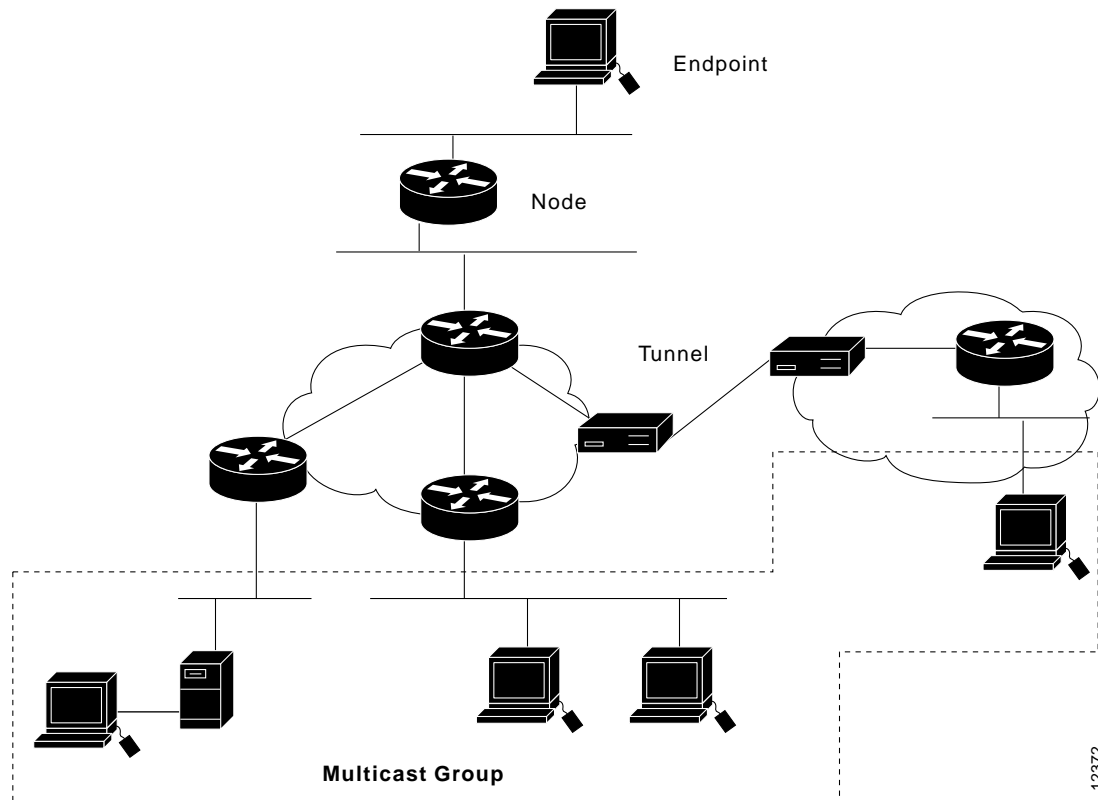


Simple Multicast Routing Protocol (SMRP)

Background

The *Simple Multicast Routing Protocol (SMRP)* is a transport-layer protocol developed to route multimedia data streams over AppleTalk networks. It supports Apple Computer's QuickTime Conferencing (QTC) technology. SMRP provides connectionless, best-effort delivery of multicast datagrams and relies on underlying network-layer protocols for services. In particular, SMRP facilitates the transmission of data from a single source to multiple destinations. This chapter focuses on the functional elements and protocol operations of SMRP. Figure 45-1 illustrates a generalized SMRP environment.

Figure 45-1 A generalized SMRP environment runs from a multicast group to an endpoint.



12372

In creating SMRP, Apple borrowed a number of strategies and concepts from other protocols and technologies. In doing so, many terms were adapted to have specific meanings in Apple's SMRP environment. Table 45-1 provides a summary of SMRP-specific terms and definitions. These terms are used throughout this chapter.

Table 45-1 SMRP-Specific Terms and Definitions

Term	Definition
Adjacent endpoint	In relation to a node or endpoint, an endpoint on the same local network or a node connected through a tunnel.
Adjacent node	In relation to a node or endpoint, a node on the same local network.
Child endpoint	An adjacent endpoint to which a node sends multicast data.
Child node	In relation to a member node of a group, a neighbor node farther from the creator endpoint.
Child port	In relation to a group, a port that is the interface to one or more child nodes.
Creator endpoint	The endpoint that requested creation of the group, and the source of data forwarded to the group.
Creator node	The primary node that created the group.
Designated node	An SMRP router that has been designated as a primary or secondary node.
Destination tree	The spanning tree rooted on a local network with paths directed toward the local network.
Distribution tree	The active subset of a source tree that is used to forward multicast data for a group.
endpoint	A nonrouting source or destination of multicast packets.
Group	A set of recipient endpoints or a multicast address.
Join	Term for the process of becoming a member of a group.
Joining path	A path on the destination tree for a local network used to reach a creator node and constructed using the SMRP distance-vector algorithm.
Leave	Term for the process of relinquishing membership in a group.
Local net	A shared-access data link and its associated network-layer protocol. A LAN might support more than one local net.
Member endpoint	An endpoint that is a member of the group.
Member node	A node that is on the distribution tree of a group.
Neighbor node	In relation to a member node of a group, an adjacent node that is on the distribution tree for the group.
Node	A router implementing SMRP.
Parent port	In relation to a group, the port that is the interface to the parent node.
Parent node	In relation to a member node of a group, the neighbor node closer to the creator endpoint.
Port	A local network or tunnel interface on an SMRP router.
Port parent	The address of the node that is responsible for handling group requests.
Primary node	The node on a local network responsible for creating groups.
Reverse path	The reverse of a joining path, a path on the source tree for a local net used to forward multicast data.
Secondary node	The node ready to take over for a disappearing primary node.
Source tree	The spanning tree rooted on a local network with paths directed away from the local network.

Term	Definition
Spanning tree	A connected set of paths using local networks between all nodes on an internetwork with only one path between any two nodes.
Tunnel	A point-to-point connection between nodes on nonadjacent networks through routers not implementing SMRP.

SMRP Multicast Transport Services

SMRP is designed to enable routers and end stations to exchange multicast packets over network-layer protocols. SMRP provides the capability to manage multicast address assignment and enables a single source to send data addressed to a unique multicast group address. Receivers join this group if they are interested in receiving data for this group. In support of these functions, SMRP involves a number of services. The discussions that follow focus on the key processes and technologies that enable SMRP services, including address management, the Multicast Transaction Protocol (MTP), node management, multicast route management, data forwarding, and topology management.

SMRP Multicast Address Management

SMRP addressing is based on the local network of a creator endpoint. An SMRP address consists of two parts: a 3-byte network number and a 1-byte socket number. Each local network is configured with a range of unique network numbers.

In network number mapping, network numbers must be assigned to local nets for SMRP and must be unique throughout an entire internetwork. Each local net can be assigned any contiguous range of 3-byte network numbers. The number of multicast groups available for a local net is the number of network numbers assigned times 254. Network numbers can be configured or can be mapped from the network number of underlying network-layer protocols. Unique network number ranges can be reserved for supported network protocols.

In the case of multicast address mapping, SMRP multicast addresses must be mapped to network-layer multicast addresses, and these in turn are mapped to data link layer multicast addresses. For each network-layer type, a block of multicast addresses must be obtained for SMRP. In the best case, these addresses will map directly. In most cases, a direct mapping is not possible, and more than one SMRP multicast address is mapped to a single network-layer multicast address.

The manner in which multicast addresses are mapped to network-layer addresses is network-layer dependent. When SMRP transport-layer multicast addresses do not map directly to network-layer multicast addresses, filtering of the SMRP multicast addresses is required. When network-layer multicast addresses do not map directly to data link layer multicast addresses, the network layer is expected to filter out multicast addresses that have not been subscribed.

Network-layer multicast addresses are preset for *AllEndpoints*, *AllNodes*, and *AllEntities* addresses. *AllEndpoints* messages sent to this multicast address are relayed to all endpoints on a network. *AllNodes* messages sent to this multicast address are relayed to all SMRP routing nodes on a network, and *AllEntities* messages sent to this multicast address are relayed to all endpoints and all SMRP routing nodes on a network.

SMRP Multicast Transaction Protocol (MTP)

SMRP involves a multicast transaction protocol (MTP) that provides for three transaction types: node, endpoint, and simultaneously node/endpoint. Communications between adjacent nodes and between nodes and endpoints occurs through request/response transactions.

Responses always are unicast. MTP provides for the retransmission of requests and/or responses in case of network errors. Only hello and designated node-request packets are sent as multicast messages; all others are unicast. Endpoint-to-node requests are sent as multicasts, while node-to-endpoint requests are sent either as unicasts or multicasts.

The basic MTP design as implemented in SMRP routers uses two queues for all transactions: a request queue and a response queue. The request-queue entries are deleted after the router processes the response that it received. The response is processed, when matched with a request, using a callback specified in the entry.

After response processing, the request is discarded. If the request is unanswered, an internally generated reject response, with the error MCNoResponse, is sent to the callback. Requests can be sent to a unicast address or the AllNodes or AllEndpoints multicast address, depending on the context. Unless explicitly redirected, requests are sent to the AllNodes multicast.

The response-queue entries are created upon receipt of a request packet. The entry is referenced during all processing of the request, and the processed entry remains in the queue until it expires and is deleted from the queue. If a duplicate request is received, it is ignored if the SMRP router is still processing the original request, or if a duplicate response is generated if processing is complete. Responses always are unicast to the requestor. Some received requests require an SMRP routing node to generate additional requests. In this case, the original request(s) will be processed by the callback handler of the routing node's request entry.

SMRP Node Management

SMRP relies on a number of node relationships, including *designated nodes*, *adjacent nodes*, and *tunnel nodes*, to permit transport of multicast datagrams.

Designated nodes are SMRP routers that have been specified as primary or secondary nodes. A designated primary node is responsible for allocating group addresses. A primary node is required for each local network with SMRP nodes. A designated secondary node is required if a local network has more than one node. The secondary is used to maintain a copy of the Group Creation table, and it becomes the primary node if the primary node for a network fails.

The basic process of primary- and secondary-node determination begins at startup, when a node first tries to become the designated secondary node on each local net. If successful, the node then tries to become the designated primary node. Transactions are initiated by either a primary-node request or a secondary-node request. No response to the request indicates that the negotiation succeeded, while a positive response indicates that the negotiation failed. If two nodes try to become the designated primary node or the designated secondary node at the same time, the node with the lower network-layer unicast address becomes the designated node. A primary node then sends add-group entry packets and remove-group entry packets to the secondary node for a local network to maintain an identical Group Creation table.

In relation to a specific node or endpoint, an adjacent node exists on the same local network. Nodes periodically send out hello packets on each port. If a hello packet is not received from an adjacent node within a certain interval of time, the nodes's adjacency state is changed to Not Operational, and associated routes are marked unreachable. Notify packets are sent to each adjacent node whenever the state of a port in the node changes to a different operational state. Each node maintains an entry in the Node table for each an adjacent node. The table entry is allocated the first time it receives a packet from adjacent node. Table entries include the time of the most recent hello packet and its state.

Tunnel nodes are point-to-point connections between nodes on nonadjacent networks through routers not implementing SMRP. Two distinct tunnel nodes are defined: tunnels between nodes and tunnels between a node and an endpoint.

Tunnel nodes are maintained as entries in the Adjacent Node table in every node in the same way as for other adjacent nodes with respect to the use of hello packets and notify packets. Similarly, SMRP enables tunnel nodes to join and leave groups in the same manner as any other adjacent node.

Note Cisco does not support tunnel nodes. SMRP, however, can be enabled to run network-layer tunnels between nonadjacent nodes.

SMRP Multicast Routes

SMRP relies on a spanning tree-based forwarding scheme to determine routing paths for multicast traffic. This route-determination process relies on the use of a distance-vector algorithm. A node sends distance-vector request packets to adjacent nodes at startup time and when routes change. The distance specified in the vector is the number of hops needed to reach a particular network number range. Nodes contain a vector for each entry in the Network Route table and send as many packets as necessary to send all the vectors. When routes change, each node sends distance-vector request packets to every adjacent node.

When a route is received on a port, the port-parent address must be set for the route for all ports. Because the group address is bound to the network address, the port-parent address also is used if a node is to handle a request for specified groups. When the port-parent address is the node's own address, the node is responsible for the request. Equal path nodes decide which node is responsible for a request by determining which node has the highest network address.

When a distance-vector request with entries for unknown local networks is received by a node, network ranges for associated local networks are added to Network Route table for the node, with a received distance incremented by one. The adjacent node that sent the distance-vector packet then becomes the parent node for the local network. The table entry is updated if a distance-vector packet is received for known local networks, and the distance-vector packet plus one is less than the entry in the Node Route table. A tie breaker is used if a distance-vector packet is received from an adjacent node with the same distance to a local network. The tie breaker is determined to be the adjacent node with a higher network-layer unicast address. That node is identified as the parent node for the local network.

SMRP Multicast Group Management

In SMRP, multicast group participation is managed via a process involving negotiations among endpoints and nodes. An endpoint attempts to join a group by contacting a node on a local network. Any contacted node is responsible for joining the distribution tree for the group by activating paths to an existing distribution tree. Nodes leave a distribution tree for a group by deactivating paths whenever no more member endpoints for the group exist on those paths. Four basic processes are required to manage SMRP groups: *creating*, *joining*, *leaving*, and *deleting*.

An endpoint sends a create-group request to the designated primary node when it wants to start sending data to a group. The primary node then assigns an unused group address and allocates an entry in the Group Creation table. The primary node finally returns the group address to the creator endpoint and sends an add-group request to the secondary node, if it exists.

Endpoints send requests to initiate joining a multicast group. The parent node for a group on a local network responds to endpoint-join group request packets. (A node determines whether it is the parent node by examining the network number in the group address.) When the parent node for a group gets a join-group request packet and that node is not yet a member of the group, the node forwards the join request toward the creator node of the group. Eventually the join-group request packet reaches a member node or the creator node for the group, and a join-group confirm packet is

sent back along the reverse path. The member or creator node adds a child port to the Group Forwarding table if the join was received on that port. When data arrives in the reverse path, it is forwarded to all child ports. When the creator node receives the first join request for a group, it forwards the request to the creator endpoint to enable it to start sending data.

To leave a multicast group, endpoints send leave-group request packets on their local net. The parent node for the group on a local net returns a leave-group confirm packet to the endpoint and sends out a group-member request packet on the child port. If the parent node does not get a group-member confirm packet on the child port from a member node or endpoint, the parent node removes that port from the entry. If the parent node has no child ports left in the entry, it sets the state of the entry to Leaving and sends a leave-group request packet up the distribution tree to its parent node. Each respective parent node removes the entry from its Group Forwarding table when it receives the leave-group confirm packet.

The endpoint sends a delete-group request when it wants to stop sending data to the group. Only the designated primary node responds to this request.

Forwarding Multicast Datagrams

SMRP data forwarding involves nodes forwarding multicast datagrams on active paths of the source tree for a particular group. An active path has member endpoints on it for the group, or is a path needed as a transit path to reach other active paths. The subset of active paths for the source tree is the distribution tree for the group. Data forwarding under SMRP involves a series of negotiations among endpoints and nodes. In general, nodes receive multicast datagrams when endpoints send data to a group. The creator endpoint can send data packets with a network-layer multicast address to its local network after it receives a join request from the creator node. Parent nodes on the local network receive this multicast and forward the packet to all child ports in the forwarding table for the group. A node multicasts a packet on a local network only if it is the parent node for the group on that local network and if the data was received on the parent port for the group. Nodes also forward data to adjacent tunnel nodes that are members of the group. In the case of an SMRP tunnel, multicast datagrams are encapsulated in a unicast network-layer packet.

Handling SMRP Topology Changes

Topology maps are maintained by SMRP entities to manage path or membership changes within an SMRP environment. SMRP anticipates a number of typical topology changes and defines specific techniques for handling them.

Disappearing Member Endpoints

To detect disappearing member endpoints, nodes periodically send a group-member request packet to each active child port. Each member node and endpoint returns a group-member confirmation packet to the parent node. If no group-member confirmation packets are received by the parent node, the node sends a leave-group request packet to its parent node and then deletes the group entry.

Stranded Groups

To detect stranded groups, creator nodes periodically send a group-creator request packet to the creator endpoint. If after a number of retries no group-creator confirm packets are received by the creator node, the group is deleted. Network route tables are kept up to date by nodes sending distance-vector packets to their adjacent nodes when routes change. This allows nodes to change multicast group routing based on changes in topology.

SMRP Transaction Example

A typical SMRP-based transaction session involves a Macintosh workstation creating a group, other Macintosh workstations joining the group, and data being sent to the group members.

In a typical SMRP transaction session, a Macintosh (call this system Creator-Mac) first sends a create-group request to all nodes on a particular network. The primary router (Primary) for the local network assigns an unused group address and returns that address to the Creator-Mac. A Macintosh on a distant network (called Member-Mac) finds the Creator-Mac via the Name-Binding Protocol (NBP).

Creator-Mac then responds with the group address via an NBP response. The Member-Mac sends a join-group request to all nodes. A remote router (say Router M) with a valid route to the group and a correct port parent sends a join-group request toward the Primary.

The Primary finally receives the join-group request and sends it to the Creator-Mac. It also adds the incoming port to the group in the forwarding table. The Creator-Mac confirms the join-group request and sends data to the group. The Primary receives the data and forwards it to the group's child ports.

Finally, the data is received by Router M, which looks up the group in the forwarding table and forwards the multicast data. The Member-Mac then receives data for the group.

SMRP Packet Format

Figure 45-2 illustrates the general SMRP packet format.

Figure 45-2 A general SMRP packet consists of five fields.

1	1	2	4	Variable
Protocol Version	Type	Sequence Number	Group Address	Data

it4302

The following descriptions summarize the SMRP packet fields illustrated in Figure 45-2:

- *Protocol Version*—Indicates the version of SMRP.
- *Type*—Consists of 2 subfields. The first 2 bits modify the packet type specified by the bottom 6 bits in order to identify whether a packet is a transaction packet, and if so, what type of transaction.
- *Sequence Number*—Matches responses to requests in transactions in order to avoid duplicate requests and responses. All packet types are transaction packets and will have a nonzero sequence number (with the exception of multicast data packets and hello packets whose sequence numbers are set to zero).
- *Group Address*—Serves as the designated primary node and assigns group addresses for all multicast sources on the local network. A particular local network can be assigned more than one network number, but multiple network numbers must be in a contiguous range. Nodes must configure network numbers that are unique to each local network and each primary node to prevent collisions of multicast addresses. When a primary node assigns a new group address, it arbitrarily assigns any unused group address for its network number.

- *Data*—Varies depending on SMRP packet type. Table 45-2 summarizes data characteristics based on packet type.

Table 45-2 Data Characteristics Based on Packet Type

Packet Type	Data Carried	Size
Multicast Data	Data	Variable, depending on network-layer datagram size
Hello	Port state	2 bytes
Notify	Port state	1 byte
Designated Node	None	0 bytes
Distance Vector	Multicast vector	8 bytes
Create Group	None	0 bytes
Delete Group	None	0 bytes
Join Group	None	0 bytes
Add Group Entry	Network-layer unicast address	Variable, depending on network-layer address format
Remove Group Entry	None	0 bytes
Leave Group	None	0 bytes
Creator Request	None	0 bytes
Member Request	None	0 bytes
Reject	Error indication	Short integer ranging from -7700 to -7710, depending on errors