

Self-Reception via Ethernet

Acnet request handling

Wed, May 12, 1999

Ethernet hardware does not permit reception of frames by the transmitting node. Since the server request logic was developed with token ring, which does permit this, the system fails to fulfill a request for data from multiple node, when one of the nodes is the node that received the server request. This note explores how to support this.

As an example to make things more concrete, suppose that the Booster data server node `node06CF` receives a request for a word of data from all Booster HLRF IRMs, including `node06CF`. The request will be forwarded to the multicast IP address that targets all Booster nodes. Because the network is ethernet, however, `node06CF` does not receive this forwarded request. (There is no logic to separate the local data from the entire request, as there is in the Classic protocol case, because it was thought that this would very much complicate the implementation of `RETDAT` support.) All the rest of the Booster nodes do receive the request, and each arranges to deliver its own contribution of replies to the requester, in this case, `node06CF`. As each contribution is received by `node06CF`, it copies it into the appropriate places in the reply buffer that will ultimately be delivered to the Acnet client node. But since the forwarded request was not received locally, `node06CF` does not transmit a reply to itself of its own contribution to the request. This results in an error code being returned for the data from `node06CF`.

Note that from the example, not only must a node be able to receive a multicast message that it transmits, assuming it is enabled to receive such, but it must also be able to receive a message that it sends to itself.

There is a mechanism for faking reception of network messages internally. It is used by the logic that supports IP fragment reassembly to fake reception of the reassembled datagram. The network frame is housed inside a "completed datagram" block and queued to the message queue that SNAP (IP reception support) awaits. This same mechanism may be used to accomplish reception of multicast datagrams and datagrams sent to the local node.

In `NetXmit`, a frame to be transmitted is assembled for passing to network hardware. In the case that it is an Acnet IP-based message, then check whether it is being sent either to a multicast destination or to itself. If it is, then we have the special case. Build a copy of the frame inside a completed datagram block, then queue it to the message queue. If it is to be sent to a multicast destination, send the frame. If it is only to be sent to itself, skip the actual transmission, but set all the appropriate used bits in the `OUTPO`.