

Network Frames Addition

for datagram activity

Sat, Jun 14, 2003

The data stream called `NETFRAME` is used to log network activities, including both datagrams transmitted as well as received. This note develops a new way to mark whether the datagram was received or transmitted.

The determination of whether a given datagram was received or transmitted has been based upon the buffer pointer included in the data stream record. If the buffer address has bit 17 set, which corresponds to mask `0x20000`, as in `0x160000–0x17FFFF`, it was interpreted as a datagram that was received; otherwise, it is assumed that it was transmitted. The corresponding transmit buffer area was `0x180000–0x19BFFF`. This empirical scheme works, but it is not flexible enough. In the IRM, a reassembled datagram is stored in allocated dynamic memory, so its buffer address may not follow the empirically chosen convention described above. A better method is needed. It would be good if a new scheme for this does not cause problems with any existing code that interprets these records.

The fields in the 16-byte record are the 4-byte buffer address, the 2-byte datagram size, the 2-byte node number, and the 8-byte time-of-day in the usual BCD format, with the last byte a binary value of half milliseconds within the indicated 15 Hz cycle of the calendar second.

Consider that a new way to indicate whether a datagram was sent or received can be done by usurping the year byte in the time-of-day for the purpose. Suppose that new code that writes records into this data stream sets the year byte to `0xC0` for reception and `0xC1` for transmission. The new client code that interprets these records can determine whether the new format is used, allowing the new interpretation, or whether the old format is still being used, in which case it would have to use the old interpretation. (Note that valid year byte values do not exceed `0x99`.)

Loss of the year byte value is not deemed significant. When time filtering is used on Page F, it can only match the current hours/minutes value. The listing of network activity produced by Page F only announces time to hours and minutes. On a 32-character “little console” line, it appears as follows:

```
01234567890123456789012345678901  
C509 003E R 167504 1132:33-03+18
```

The new scheme allows more flexible assignment of network buffer areas, and it eliminates misinterpreting a large reassembled datagram that is deposited into allocated memory.

It has been learned that a very large number of datagrams arriving nearly simultaneously at an IRM suffer from misinterpretation due to wrapping of the receive buffer. The receive buffer occupies 128K bytes, allowing for 85 buffers of 1500 bytes each. If 142 datagrams are received in response to a multicast UDP Echo request, they can all be received, all right, but the subsequent processing is late, resulting in apparent duplicate responses, which actually came from different nodes. With the new scheme, we can increase the receive buffer area to 512K bytes, say, allowing for 340 full-frame buffers. This should eliminate the problem due to receive buffer wrapping. A side benefit of having a larger receive buffer area is that capturing the contents of a received datagram is easier, since it will take longer for the receive buffer to wrap. For the IRM, memory in the range `0x280000–0x2FFFFFF` might be used for a larger receive buffer area. Since the larger area might not be deemed necessary for every IRM, it may be useful to make the use of this memory for this purpose optional, perhaps by installing something in the TRING table.