

Network Working Group
Request for Comments #176
NIC #7100
Categories: C and D

A. Bhushan, MIT
R. Kanodia, MIT
R. Metcalfe, MIT
J. Postel, UCLA
14 June 1971

Comments on Byte Size for Connections -----

There are at least the following three views on the use of byte size for network connections*:

- 1) Byte size should not be used at all.
- 2) Byte size is solely for the convenience of NCP's.
- 3) Byte size choice is a user-level prerogative.

According to the first view, network connections are bit streams, and messages should contain bit counts (i.e., a byte size of 1). This view existed before the "Glitch Cleaning" of RFC 107, and was discarded in favour of byte stream because of stated reasons of efficiency in storage management and message concatenation.

The second view represents a special interpretation of RFC 107. According to this interpretation, byte size is entirely a 2nd level (i.e., NCP) issue. There is no requirement that 3rd level user processes be able to specify byte size. This view is indicated in RFC 151 by Shoshani.

* Byte size for connection is the byte size selected by sending NCP, as explained in RFC 107 (Output of Host-Host Protocol Glitch Cleaning Committee).

According to the third view user processes are always allowed to choose byte size for connection, either explicitly (specify a specific byte size parameter) or implicitly (byte size depends on I/O mode). An NCP is allowed to use a default byte size, if the user does not specify it.

The Correct View

The third view should be considered the correct interpretation of RFC 107. In fact, RFC 107 states on page 2, "the choice of the byte size for a connection is a 3rd level protocol issue." To be consistent with TELNET, ICP, and other 3rd level protocols which require that a specific byte size be used for connection, it is imperative that corresponding 3rd level processes be able to specify (and_impose) a particular byte size to the NCP. NCP implementors should take note of it.

On Specifying Fixed Byte Sizes in 3rd Level Protocols

Holding the view that byte size choice is a 3rd level issue, we are still faced with the following two questions. First, is it appropriate for 3rd level protocols to legislate a specific byte size for all connections using that protocol? Second, if it is appropriate to specify byte size, then what should this choice be?

There are two arguments in favour of using specific byte size in 3rd level protocols. First is that a potential mismatch problem exists because RFC 107 does not require that NCPs be capable of handling all byte sizes 1 through 255. Using a fixed byte size of 8-bits or 8-bit multiples resolves the problem as this is acceptable to all hosts (including terminal IMPs).

The second argument is one of efficiency. If it is agreed before hand that only a specific byte size would be used, it is possible to make programs more efficient (i.e., reduce program space, and possibly run time). The efficiency argument assumes that the byte size for connection represents the natural byte structure of data being transferred over the connection.

For TELNET and ICP, the byte size choice is straight forward as data is naturally in 8-bit multiples (8-bit ASCII characters in TELNET, and 32-bit socket numbers in ICP). But for data transfer protocols, the byte size choice is more complex, as data may be structured in a variety of byte sizes. Specifying a byte size for a data transfer connection reduces efficiency in instances where connection byte size does not correspond to data byte size. Further, filler fields will be required for data blocks which are not a multiple of the fixed byte size. This imposes an additional overhead.

Even if all hosts were to accept arbitrary byte sizes, and the 3rd level protocol does not legislate a specific byte size, the inefficiency problem will not be solved entirely. Under current specifications "the byte size is fixed for the life of a connection".* This means that byte size cannot be varied during the life of a connection even if structure of data varies. The problem of inefficiency is solved only for instances in which data has a constant byte structure.

Given the current state of the network, it appears that specifying fixed byte size in 3rd level protocols is a good idea. This eliminates the potential byte size mismatch problem, and improves efficiency at least for TELNET and ICP. In data transfer, the efficiency issue is more complex, as discussed earlier. It is not clear that overall efficiency would be degraded if a fixed byte size was required.

On Reopening the Byte Size Issue

The above discussion exposes certain weaknesses in the efficiency arguments for having byte streams on network connections. In moving from bit stream to byte stream, we may have lost generality, and it is not clear how much overall efficiency is gained. Sometimes, the gain in NCP efficiency may be at the expense of user process efficiencies.

* RFC 107, page 2

It is also clear that for efficiency arguments to hold, the byte size choice should not be an NCP prerogative. It is the combined efficiency, rather than NCP efficiency which should be our primary concern. Restricting byte size choice to NCPs has the further disadvantage of potential byte size mismatch not only between communicating NCPs but also at the user-NCP interface. Therefore, allowing a user process to specify byte size is a step in the right direction, given that we have adopted byte streams.

It is our opinion that the issue of bit stream and byte stream be set aside until serious consideration can be given to a major Host-Host Protocol overhaul. At a later stage we will have a better idea of the relative efficiency merits.

[This RFC was put into machine readable form for entry]
[into the online RFC archives by BBN Corp. under the]
[direction of Alex McKenzie. 12/96]

