

INDRA
Working
Paper

INDRA Note 959
TSIG 4.0
IEN 153
29th July 1980

Realization of the Yellow Book Transport Service Above TCP

C. J. Bennett

ABSTRACT: This note defines an enhancement of the service provided by the US DoD Standard Transmission Control Protocol (TCP) sufficient to meet the requirements of the UK Network Independent Transport Service (the Yellow Book).

1. Introduction

This document defines a means of providing the Yellow Book Transport Service [1] above the DARPA Internet facilities, in particular TCP [2], so that this can then be used to support other services such as endpoint file transfer without requiring UK hosts to implement the Internet family of protocols. It assumes familiarity with both TCP and the Yellow Book.

The basic approach taken is to enhance the TCP service along the lines suggested for enhancing X25 in Annex I of the Yellow Book, taking into account the different services provided by TCP. In addition, the note discusses how to integrate Yellow Book TCP so that it can run alongside ordinary TCP - an issue the Yellow Book ignores for Yellow Book X25.

2. Deficiencies of TCP

A comparison of the services provided by TCP and those provided by the Yellow Book reveals that TCP is unable to support directly, either in whole or in part, the following Yellow Book features:

- (i) The RESET and ADDRESS primitives.
- (ii) The Yellow Book multiple-domain addressing structure. The TCP address space constitutes a single naming domain in Yellow Book terms. Consequently, features involving addressing - notably ACCEPT - are inadequately supported by TCP.
- (iii) Much of the subsidiary information provided with Yellow Book primitives. The fact that the source address provided with certain actions such as DISCONNECT is not provided is again a limitation of the global TCP naming domain. The Yellow Book 'Explanatory Text' parameters have no corresponding feature in TCP.
- (iv) The closest equivalent to Yellow Book EXPEDITED data - theoretically requiring a priority data channel - is TCP URGENT data. However, TCP URGENT data remains in sequence, and the URGENT pointer only marks the end of the data. Its beginning is not delimited.
- (v) The Yellow Book DISCONNECT is a full-duplex close, whereas the TCP CLOSE is only half-duplex. The TCP RESET is a unilateral close, used in error conditions. Connection closure

provides particularly subtle problems.

Hence in order to provide a Yellow Book service above TCP an enhancement of TCP is necessary. The remainder of this document discusses such an enhancement.

3. Principles of the Enhancement

The basic principles of the enhancement are as follows:

- (i) Where a TCP function corresponds directly to a Yellow Book function that TCP function is used directly.
- (ii) Where the Yellow Book function requires more information or action, the TCP function is associated with a TCP Control Message in a defined way. This message is a TCP letter of defined format containing the information required.
- (iii) Where there is no TCP function even remotely corresponding to the required Yellow Book function, a control message is defined which may be used by the source and destination processes if possible, and may be forwarded into other transport domains more capable of taking the appropriate action.

4. The Yellow Book TCP Enhancement

4.1 Distinguishing the Yellow Book TCP

There are several ways a TCP connection providing a Yellow Book service could be distinguished from a normal TCP connection. These are as follows:

- (i) A single TCP port could be reserved for Yellow Book service. Additional multiplexing could then be provided using the lines proposed in Annex III of the Yellow Book.
- (ii) A number of TCP ports could be reserved for the different higher level services required, each one of which would be defined as including the enhancement defined within this document.

- (iii) A TCP option could be defined which when associated with a SYN would denote the 'enhanced TCP service' option - i.e. the data stream would be governed by this document.

The first of these possibilities leads to unnecessary duplication of multiplexing facilities - perfectly adequate multiplexing is already done within the TCP. The second is potentially restrictive in that it limits the growth of future services, and would probably eventually lead to the informal adoption of the first solution. This notes supports the third course, subject to approval by the DARPA Internet Group. Proposed text defining the TCP option, using the format of the TCP specification, is as follows:

Enhanced TCP Service

```
+-----+
|00000010|
+-----+
Kind = 2
```

This option specifies that the TCP connection is providing the data formats and command messages necessary to support the enhanced services offered by the UK standard Network Independent Transport Service. This option may only be specified in the header of a packet which has the SYN bit set to 1. If the receiving process is unable to support this option, the connection should be aborted via a RESET.

The adoption of this option does not imply that a TCP itself is required to understand the structures of this document. It does allow such TCPs to be built. For TCPs which do not support these structures directly, the option is effectively a null option, whose presence should be indicated to the receiving process.

Failing the adoption of this facility, the first choice (of a single reserved port) is supported here.

4.2 Identification of TCP Command Messages

Once the connection is identified as providing a Yellow Book TCP service, the next problem is how to identify TCP Command Messages within the data stream. The X25 enhancement made use of the X25 Q-bit for this

purpose, but there is no corresponding function within TCP. The choices are as follows:

- (i) Provision of a TCP Q-bit facility. This is unlikely to occur, not least because the example of XXX shows that it is liable to misuse.
- (ii) Further definition of the 'Enhanced Service' option, so that the occurrence of this option with a letter specifies that the letter is a Command Message.
- (iii) Structuring the data stream itself.

Despite the marginally greater data efficiency of the second choice, the last choice is supported here, as it requires no modification of the TCP user interface. It does, however, require the transmission of a data octet which will require a fairly clever user interface if extensive buffer copying is to be avoided.

The structure proposed is as illustrated in Figure 1. Each TCP letter is prefaced by a single octet, known as the TYPE octet. This takes a value of 0 for data letters and a value of 1 for command messages. All other values are undefined.




Figure 1: Letter Structure in Yellow Book TCP

4.3 Structure of TCP Command Messages

A TCP Command Message is a Yellow Book TCP Letter of TYPE 1.

The first octet of the message defines the COMMAND CODE of the message. These codes are defined in subsequent sections, and have been chosen to correspond to the command codes of the X25 Command Messages.

Following the command code is a number of PARAMETERS. The significance of these parameters is defined by their position in the parameter sequence for each command; thus no intermediate parameter may be omitted, though it may have a null value. Parameters, however, may be omitted if they and all succeeding parameters have null values.

Most parameters have a free field format. For this reason each parameter is constructed of a number of FRAGMENTS. These fragments consist of a header byte and the body of the fragment, which may have a maximum length of 127 octets.

The fragment header is a single octet. The high-order bit of this octet, if set to 1, declares that the current fragment is the last fragment of the current parameter. The remaining seven bits define the length in octets of the current fragment.




Figure 1: TCP Command Message Structure

A parameter with a null value is represented by a fragment header whose EOP bit is set to 1 and whose count field is set to 0. The rules governing the syntax of free-field parameters are the same as those defined in section 5.7 of the Yellow Book, based on the use of the IAS character set.

The sole difference between this structure and the X25 Command Message structure is that the count field in the fragment header is extended by one bit - this bit is used for a specific purpose in X25 CONNECT messages which does not arise with the TCP. This doubles the maximum fragment size. Because of the similarity of structure the same terminology has been used, though the term 'fragment' is somewhat

This structure is summarised in Figure 2.

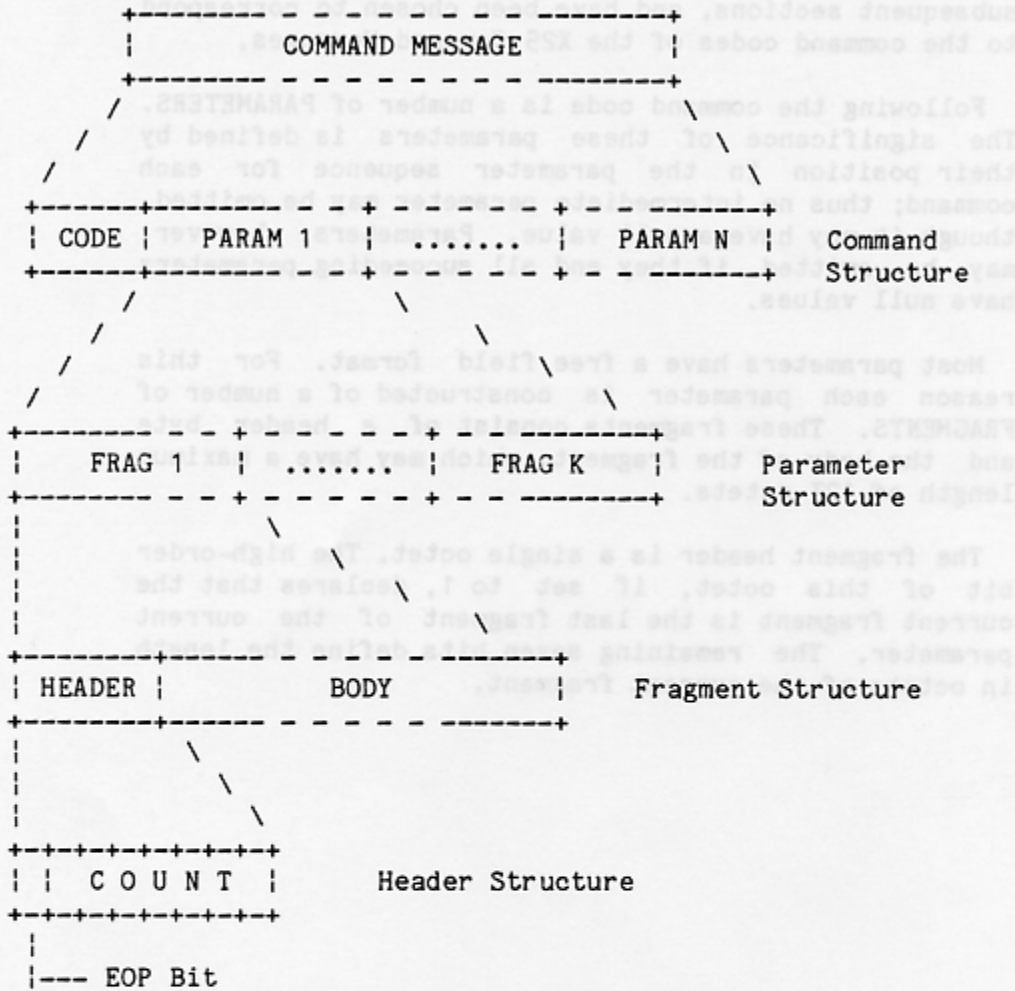


Figure 2: TCP Command Message Structure

A parameter with a null value is represented by a fragment header whose EOP bit is set to 1 and whose count field is set to 0. The rules governing the syntax of free-field parameters are the same as those defined in section 2.7 of the Yellow Book, based on the use of the IA5 character set.

The sole difference between this structure and the X25 Command Message structure is that the count field in the fragment header is extended by one bit - this bit is used for a specific purpose in X25 CONNECT messages which does not arise with the TCP. This doubles the maximum fragment size. Because of the similarity of structure the same terminology has been used, though the term 'fragment' is somewhat

unfortunate in the DARPA context through its specific association with the Internet Datagram Protocol.

4.4 Yellow Book Commands and TCP

The following general remarks apply to the following specification:

- (i) All codes are specified as decimal integers.
- (ii) All address fields include the appropriate TCP address components, specified as /<NET ID>+<INTERNET HOST ID>+<PORT NO>/, where the bracketed fields are the character strings of the octal representations of those TCP fields, except where otherwise noted. Thus, the NIFTP port (octal 10651) at ISIE (net 12, host 1, logical host 0, IMP 64) will be denoted by the string:

/12+200064+10651/

4.4.1 CONNECT

The CONNECT command message is defined by the following code and parameters:

- Code = 16
- Parameter 1: Called Address
- Parameter 2: Calling Address
- Parameter 3: Quality of Service
- Parameter 4: Explanatory Text

This message will be preceded by the usual TCP three-way handshake. Where possible or appropriate, the quality of service parameter will be used to select TCP quality of service from the options defined in the TCP specification.

The CONNECT message will be the first message sent by the calling party. It will be possible for the calling party to initiate the transfer of data before the arrival of an ACCEPT message.

4.4.2 ACCEPT

The ACCEPT command message is defined by the

following code and parameters:

Code = 17
Parameter 1: Recall Address
Parameter 2: Quality of Service
Parameter 3: Explanatory Text

This message will be the first message sent by the called party after the three-way handshake, unless the call request was rejected (see DISCONNECT, below). If the recall address and the quality of service do not differ from those specified in the CONNECT, the ACCEPT will normally consist of the code octet only.

4.4.3 DISCONNECT

The DISCONNECT command message is defined by the following code and parameters:

Code = 18
Parameter 1: Reason
Parameter 2: Address of DISCONNECT Initiator
Parameter 3: Explanatory Text

The reason parameter is a single octet giving a machine-oriented encoding of the reason the DISCONNECT was initiated. The defined reasons are listed in Appendix B of the body of the Yellow Book. Parameter 2 is included to cover the case where the DISCONNECT was initiated by some intermediate gateway (where 'gateway' is used in the Yellow Book sense).

DISCONNECT will always cause the TCP to issue an URGENT call. On receipt of a DISCONNECT message, no further data may be sent and all data currently queued for transmission should be flushed if possible. No data will be passed across to the user after a DISCONNECT has been issued.

Beyond this, the exact DISCONNECT sequence used varies depending on the state of the connection, as follows:

- (i) If the DISCONNECT is being used to reject a CONNECT request, the DISCONNECT message will be followed by a TCP RESET. This will abort the TCP connection, flushing all outstanding data. No response is expected. The URGENT pointer points to the TCP RESET.

- (ii) In the normal case of closing an open connection, the DISCONNECT issued by the initiator will be followed by a TCP FIN. The remote party will respond with an optional DISCONNECT message accompanied by a FINACK and a FIN. The URGENT pointer points to the TCP FIN.
- (iii) For error terminations, a DISCONNECT message should be answered with a TCP RESET. The issuer of the DISCONNECT will also issue a TCP RESET after a timeout period, if a RESET has not already been received. The URGENT pointer points to the end of the DISCONNECT message.

4.4.4 DATA

DATA is sent as a sequence of Yellow Book TCP data letters, as defined above.

4.4.5 PUSH

The PUSH function is conveyed by use of the TCP EOL function, pointing to the data octet at which the PUSH was issued.

4.4.6 EXPEDITED

The EXPEDITED command message is defined by the following code and parameter:

Code = 21
Parameter 1: EXPEDITED data

EXPEDITED data is accompanied by a TCP URGENT pointer pointing to the end of the letter.

It should be noted that this will cause the receiver of the URGENT to process all data up to the URGENT pointer in 'fast' mode, whether EXPEDITED or not. As noted above, TCP has no direct equivalent of a priority data channel, but the mechanism described allows the preservation of EXPEDITED data so that it may be passed as such in subsequent networks.

4.4.7 RESET

The RESET command message is defined by the following code and parameters:

Code = 19
Parameter 1: Reason
Parameter 2: Address of RESET Initiator
Parameter 3: Explanatory Text

TCP has no equivalent of the RESET function (a TCP RESET is something else entirely). Thus the only TCP action taken with a RESET message is to accompany it with an URGENT pointer pointing to the end of the RESET message.

As with DISCONNECT, the defined RESET reasons are those listed in Appendix B of the main portion of the Yellow Book. The address parameter is again included to cater for the case where a RESET was initiated in some intermediate network.

A RESET may only be issued if the connection is fully open and there are no RESETs already outstanding. A RESET message must always be replied to with another RESET message, leaving the connection open, or with an error DISCONNECT message followed by a TCP RESET, which will abort the connection. All data and control messages, with the exception of DISCONNECT, received after a RESET has been issued and before a RESET reply has been received, will be discarded without informing the user. In the case of DISCONNECT, the connection will be considered as having closed in an abnormal state. If a DISCONNECT has been issued, a received RESET will be ignored.

Where possible the issue of a RESET should cause the sender to flush its transmission buffers.

4.4.8 ADDRESS

The ADDRESS command message is defined by the following code and parameters:

Code = 20
Parameter 1: Address
Parameter 2: Address Qualifier

The ADDRESS qualifier is a single octet parameter taking one of the values defined in the Yellow Book:

- 0: The message is passing towards the addressed object.
- 1: The message is passing away from the addressed object.
- 2: An addressing error has been detected.

There is no associated TCP action taken with an ADDRESS message.

The receiver of an ADDRESS message will perform the appropriate ADDRESS transformation as defined in the Yellow Book.

It is recommended that the ADDRESS function should not be used.

5. Conclusions

One of the difficulties of writing a note such as this is that it is addressed to several audiences with different interests and not necessarily a great deal of overlap either in aim or in background.

The immediate audience is the team at University College London who are involved in implementing a 'Protocol Converter' to make possible direct access between hosts in Britain using the CCITT and UK national standards and the hosts in the US based on the DARPA Internet standards. For this audience, the document hopefully defines an answer to what will soon be a practical need, though it is a matter for continuing debate to what extent the full enhancement defined here will be implemented.

Within Britain, the wider audience aimed at is centred on the Transport Service Implementors' Group. For this group, the aim of the document will be well understood - it is defining a service enhancement similar to the one that is already defined for X25 and the ones they are defining for their local campus networks. The aim is to provide a common transport service for all these systems. They will be unfamiliar with the detailed nature of the TCP itself, but this is not particularly important. The major interest of the document lies in the fact that the system being

enhanced is not an ordinary local network, but an entire family of networks, and the resultant enhancement will make possible direct authorised access between UK and US hosts. I would also like to point out the issue of separating Yellow Book enhancements from ordinary uses of a network service. This issue is not addressed by the X25 enhancement specification.

The US Internetwork Group is likewise unfamiliar with the concepts and details of the Yellow Book Transport Service. For them, the document will be of interest in that it shows how to coordinate two very different approaches to internetworking. The Catenet, based on TCP, can be described as a strongly connected internetwork system, with common transport protocols and a common address space. The UK Transport Service takes an approach based on providing a common service interface, which leads to a weakly connected system with common service protocols and no common address space. Within this approach, the entire Internet system appears as a single component network, as delimited by the TCP and its address space.

Beyond this the document proposes a new TCP option. For most existing TCPs this option will effectively be a null option which must be signalled to the user at call setup time. However, there is no reason why a TCP could not be built which contained the features of this note as optional enhancements selectable by choosing the option.

6. References

- [1] - PSS User Forum Study Group Three: "A Network Independent Transport Service" SG3/CP(80)2. February 1980.
- [2] - Information Sciences Institute: "Transmission Control Protocol" IEN 112. August 1979.