

DATA AND FILE TRANSFER - SOME MEASUREMENT RESULTS

During the last six months, we have been monitoring (although not continuously) the performance of our FTP-user and FTP-server programs. The purpose of this paper is to 1) discuss measurement criteria, 2) describe the measurement facilities, 3) report the relevant measurement results, 4) discuss the significance of results and compare them with other measurement data, and 5) ask for suggestions on our measurement and summarizing procedures.

I. THE MEASUREMENT CRITERIA

The FTP (Ref. "The File Transfer Protocol", by Abhay Bhushan, NWG/RFC 354, NIC 10596,) may be considered a facility for data transfer between file systems. The relevant measurement parameters for a data transfer facility are:

- 1) Transfer rate (both peak and average, measured in bits per second) which determines the throughput of the data transfer facility.
- 2) Response time or delay (measured in seconds) which determines the "interactibility" of the facility.
- 3) Processing cost (measured in dollars or cpu-seconds per megabit transferred) for transferring the data between the network and the file system. This is only one component of the cost of transferring data, the other component being the communication cost (including IMP processing costs) which we take as given.
- 4) Failure-to-connect rate - average time elapsed between failures to connect to the facility (measured in hours). Failures could be in the Host (processor and file system) hardware or software, or in the IMPs and telephone lines.
- 5) Availability - the percentage of time a given facility is available, or alternately the probability of finding the facility available at a given time.
- 6) Accuracy - measured by the probability of error in transferring bits, bytes, blocks, or files.

II. THE MEASUREMENT FACILITIES

The MIT-CMS survey program (ref. "A Report on the Survey Project" by Abhay Bhushan, NWG/RFC 530, NIC 17375) measures the response-time, failure-to-connect rate, and availability of the Host-logger facility (on socket 1). Our preliminary experiments have indicated that the corresponding measurement results for the FTP are very close to that for the logger (at least they are the same order-of-magnitude). As the use of FTP and the ARPANET is increasing rapidly, most Hosts have their logger and FTP operational whenever their Host and NCP (Network Control Program) are functioning. The response time for obtaining the use of FTP service is very close to that for obtaining the use of the logger service as both involve the use of the ICP (Initial Connection Protocol).

Preliminary results from the Survey Project indicate that the average response time in recent months has been about 2.7 seconds. The average availability has been about 85% with the failure-to-connect rate being about once every 10 hours. Table I shows summary results for the time period August 26 through August 31, 1973, for three Hosts with TENEX operating systems (SRI-ARC (NIC), BBN-TENEXA, and USC-ISI).

The reader is cautioned that the data below reflects the Host performance as seen by the MIT-DMS survey program which surveys the Hosts only once every twenty minutes. Consequently, the actual host performance may be somewhat different. Also, we cannot distinguish between IMP, telephone lines, and Host failures and the response time of a host is affected by its distance (number of IMP hops) from the MIT IMP (IMP 6).

In the data shown in Table II, each success or fail response is considered to have a duration of 20 minutes, so Hosts are given the benefit of the doubt for the time we are not surveying. In addition, the response time has been averaged only for the successful logger available responses. The logger is considered available if the SURVEY program can establish a full-duplex connection within 20 seconds. The Host is considered available when it is not in the "DEAD" state (states in which logger is not up but the Host is available are logger not responding and logger rejecting).

TABLE I

RESPONSE TIME, AVAILABILITY, AND FAILURE RATE FOR SELECTED HOSTS
(based on SURVEY data for 8/25/73 through 8/31/73)

| PARAMETER | NIC | BBN | ISI |
|---|------|-----|------|
| Average Response-time (sec.) | 2.7 | 2.4 | 3.0 |
| Host Availability | 93% | 85% | 87% |
| Logger Availability | 91% | 79% | 83% |
| Failure-to-connect rate for Host (hours) | 18.2 | 9.4 | 18.1 |
| Failure-to-connect rate for logger (hours) | 16.0 | 6.0 | 10.0 |

The details on the above measurements will be reported in a forthcoming paper. This paper will focus on the remaining parameters of transmission rate, processing costs and accuracy, as measured by the MIT-DMS File Transfer Measurement facility.

The FTP measurement facility exists in the MIT-DMS CALICO subsystem. Each time the MIT-DMS FTP-user or FTP-server program in the CALICO subsystem is used to transfer files (and data) via the ARPANET, it records in a local disk file the following transfer parameters: the remote Host involved, the date and time the transfer is initiated, the total number of bits transferred, the real time taken (in seconds) for the transfer, the CPU time (in micro-seconds) used by the program, whether the program is the server or user, and the FTP parameter settings for byte size (BYTE), representation type (TYPE), transfer mode (MODE), and the file structure (STRU). Programs exist in CALICO to display and summarize this data.

It should be noted that no measurements are recorded when the non-CALICO FTP-user and FTP-server programs are used for transferring files. Therefore it should be pointed out that the measurement represents a small subset of our total FTP-usage. The CALICO FTP-server was operated only till May 1973, when we switched to the non-CALICO FTP-server. (The switch was made because CALICO still undergoing development is somewhat less reliable. As CALICO stabilizes we may again operate the CALICO server and continue measuring data transfer.) In addition many users prefer to use the simpler (involving fewer system resources) stand-alone FTP-user

program. The measurement does include the data transferred when FTP is used indirectly by such commands as "copy", "print", "listf", and "mail.file" in the CALICO NETWRK subsystem.

III. THE MEASUREMENT RESULTS

The measurement facility has been operational (though not continuously) since 25 February 1973. It has recorded the transfer of 304 files consisting of 57.6 million bits. Over 90% of the bits transferred (but only 75% of the files) used the more efficient Image-36 stream mode (TYPE I, BYTE 36, MODE S) of transfer. The remainder of the files were transferred using the ASCII-8 stream mode (TYPE A, BYTE 8, MODE S). It should be noted that even though block mode was available, it was never used by our users (primarily because many FTP-servers do not implement it, and it is less efficient to use). All the files had a sequential non-record file structure (STRU F). A summary of the measurement results is shown in Table II.

TABLE II

SUMMARY OF FTP MEASUREMENT RESULTS

| Subset of data | # Files | # bits Mbits | Av. File Kbits | Speed Kbps | CPU-use sec/Mb |
|------------------|---------|-----------------|-------------------|---------------|-------------------|
| Total | 304 | 57.6 | 189 | 7.56 | 4 |
| Image 36 mode | 223 | 53.6 | 240 | 9.35 | 3 |
| ASCII-8 mode | 81 | 4.0 | 49 | 2.09 | 19 |
| Server sending | 62 | 3.8 | 61 | 7.50 | 2 |
| Server receiving | 110 | 19.8 | 180 | 7.44 | 1 |
| User receiving | 83 | 22.8 | 276 | 7.92 | 6 |
| User sending | 49 | 11.1 | 225 | 7.09 | 4 |

The entire display of the measurement data and the summaries shown in Table II are generated by the "PFTPST" (Print FTP Statistics) program in the CALICO subsystem. A sample of the data displayed is shown in Table III. The BPS (bits per second) and the M/B (CPU microseconds per bit or CPU seconds per Megabit) information is calculated by the displaying program. The largest file transferred was 5.03 Mbits, a "STOR" by the FTP-user to MIT-AI. The transfer took 10 minutes of real time for a transfer rate of a little over 10 Kbps. The highest data transfer rate recorded was 27.8 Kbps, a

"RETR" from BBN-TENEXA to MIT-DMS FTP-server. The length of the file in the above case was 28 Kbits. Needless to say that both of the above transfers used the more efficient Image-36 mode for transfer. The smallest file and the smallest transmission rate recorded was an 80 bit "MLFL" to MIT-ML (using ASCII-8) which took 7 seconds real time for 11 bps transfer rate.

TABLE III

SAMPLE DISPLAY OF FTP MEASUREMENT DATA

| -#- | ---HOST--- | COMM | --DATE-- | --TIME-- | --BITS-- | -BPS- | M/B | T | BY | PRG |
|-----|------------|------|----------|----------|----------|-------|-----|---|----|-----|
| 2 | sri-arc | STOR | 73/08/09 | 18:19:49 | 121392 | 1395 | 21 | I | 36 | U |
| 198 | mit-ml | STOR | 73/08/15 | 15:00:30 | 50688 | 5336 | 8 | I | 36 | U |
| 198 | mit-ml | RETR | 73/08/15 | 15:01:14 | 50688 | 10137 | 12 | I | 36 | U |
| 198 | mit-ml | STOR | 73/08/15 | 15:02:33 | 255456 | 8808 | 7 | I | 36 | U |
| 198 | mit-ml | RETR | 73/08/15 | 15:03:58 | 258048 | 8601 | 12 | I | 36 | U |
| 134 | mit-ai | STOR | 73/08/15 | 15:13:17 | 286720 | 1898 | 29 | A | 8 | U |
| 134 | mit-ai | RETR | 73/08/15 | 15:18:39 | 258048 | 9557 | 14 | I | 36 | U |
| 134 | mit-ai | STOR | 73/08/15 | 15:19:42 | 258048 | 6974 | 7 | I | 36 | U |
| 2 | sri-arc | RETR | 73/08/15 | 15:31:20 | 7236 | 3618 | 22 | I | 36 | U |
| 2 | sri-arc | STOR | 73/08/15 | 15:32:55 | 49428 | 8238 | 31 | I | 36 | U |
| 2 | sri-arc | RETR | 73/08/15 | 15:34:56 | 49428 | 3530 | 15 | I | 36 | U |
| 2 | sri-arc | STOR | 73/08/15 | 15:38:09 | 49428 | 7061 | 8 | I | 36 | U |
| 2 | sri-arc | STOR | 73/08/20 | 15:18:26 | 35460 | 2364 | 9 | I | 36 | U |
| 2 | sri-arc | RETR | 73/08/20 | 16:08:09 | 58832 | 426 | 153 | A | 8 | U |
| 2 | sri-arc | RETR | 73/08/22 | 12:46:10 | 10512 | 166 | 247 | A | 8 | U |
| 2 | sri-arc | RETR | 73/08/23 | 16:29:37 | 320 | 64 | 369 | A | 8 | U |
| 2 | sri-arc | RETR | 73/08/24 | 12:25:38 | 9992 | 262 | 254 | A | 8 | U |
| 2 | sri-arc | RETR | 73/08/24 | 12:27:26 | 9992 | 454 | 250 | A | 8 | U |
| 198 | mit-ml | STOR | 73/08/29 | 10:40:58 | 768924 | 7538 | 7 | I | 36 | U |
| 198 | mit-ml | STOR | 73/08/29 | 10:44:09 | 166572 | 5552 | 7 | I | 36 | U |
| 198 | mit-ml | STOR | 73/08/29 | 10:54:32 | 166572 | 7932 | 7 | I | 36 | U |
| 198 | mit-ml | STOR | 73/08/29 | 13:48:18 | 158040 | 12156 | 7 | I | 36 | U |
| 69 | bbn-tenexa | MLFL | 73/08/29 | 22:30:55 | 5600 | 1866 | 51 | A | 8 | U |
| 69 | bbn-tenexa | MLFL | 73/08/29 | 22:31:42 | 5600 | 2800 | 50 | A | 8 | U |
| 86 | usc-isi | MLFL | 73/08/29 | 22:33:55 | 5600 | 1400 | 54 | A | 8 | U |
| 69 | bbn-tenexa | MLFL | 73/08/29 | 22:36:15 | 5600 | 2800 | 48 | A | 8 | U |
| 69 | bbn-tenexa | MLFL | 73/08/29 | 22:36:54 | 5600 | 2800 | 49 | A | 8 | U |

It should be pointed out that recent measurement data for ASCII-8 transfer includes retrieval of "NIC Journal" documents ("

In transferring files we found the ARPANET and the FTP to be quite reliable. On numerous occasions we transferred complete listing of our operating system (about 6 million bits), reassembled it and ran it with no problem. No data lossage problems have been reported to us as yet.

IV. THE SIGNIFICANCE OF MEASUREMENT RESULTS

First of all let me state my complete agreement with Barry Wessler (Ref. "Revelations in Network Host Measurements" NWG/RFC 557, NIC 18457) that the measurement results should be taken in the spirit: "Here is a place to make the Network better" rather than: "Look, isn't the Network terrible." We take these measurements in the same spirit and have found the measurement effort to be quite fruitful. In several instances, with the aid of our measurement facilities, we have been able to improve the performance of our Network programs by an order-of-magnitude (just as Don Allen at BBN improved Greg Hicks' RJS program).

Our measurement results are in close agreement with the BBN FTP measurements (8.2 cpu seconds/Mb for 8-bit byte and 2 CPU seconds/Mb for 36-bit byte transfers). We also find the 36-bit byte transfer to be an order-of-magnitude more efficient than 8-bit byte transfer. The processing cost (assuming \$6.00 per CPU minute) for transferring a Megabit of information comes to about \$1.90 for ASCII-8 mode as compared to only \$0.30 for Image-36 mode. The difference in transfer rate is equally astounding being 9.4 Kbps for Image-36 as compared to only 2 Kbps for ASCII-8.

It is therefore recommended that Image-36 mode be used as much as possible to transfer data between PDP-10s (of which there are many on the ARPANET). It is strongly urged that protocols and programs allow (and use) the Image-36 mode for all data transfers including mailing files (MLFL), listing directories (LIST, NLST), and sending/retrieving NIC Journal documents. Many of the MID-DMS user programs such as "COPY" and "FTP" take advantage of the fact that the remote Host is a PDP-10 (there is a table of PDP-10's in "COPY") and use the more efficient Image-36 mode. Such a procedure is highly recommended.

The effective IMP-IMP data transfer rate is about 37.5 Kbps over the 50 Kbps telephone line (Ref. McQuillan John M., "Throughput in the ARPA Network--Analysis and Measurement," BBN Report 2491, NIC 14188, January 1971). The Host-to-Host data transfer measurement performed by BBN (above reference, p. 28) have indicated a transfer rate of 30-35 kbps BBN-to-BBN (0 IMP hops) and 12-16 Kbps BBN-to-SRI (5 hops) using single link. As FTP transfers data via a single link, a maximum transfer rate between 12 and 35 Kbps (depending on number of

IMP hops) can be expected if that file transfer is the only activity going on. In this light our maximum transfer rate of 27 Kbps to BBN (2 hops) is probably the most one can expect out of any program. The average transfer rate of 9.4 Kbps (for Image-36) transfer also appears reasonable in view of the fact that during many of the transfers other network activity is also going on, and that many of the transfers are performed when the respective computer systems are quite heavily loaded. Our measurement data does reveal that transfer rate is appreciably higher during the times a computer is likely to be lightly loaded.

The above does not mean that improvements are not possible or not required in the state of the ARPANET data transfer. Our measurement data has revealed areas in which improvements can be and should be made. For example, the transfer of data to other MIT Hosts (0 IMP hops) and back to ourselves should be faster than what we currently achieve (transfer to BBN is faster!). The probable reason for the above discrepancy is that our allocation (Host-Host protocol) is very small (2944 bits) as compared to that provided by BBN (17724 bits). This means that to transfer data our Network Control Program (NCP) has to wait for an allocation many more times while communicating to an ITS system than to a TENEX system. Large allocations are always desirable but even more so while transferring files. NCP designers can (and should) modify NCP's to allow large allocates (larger NCP buffers) for file transfer even at the expense of smaller allocates for other types of connections (such as a terminal connected to a computer system) which do not require or use the larger allocation. In addition, a new allocate should be sent as soon as data is read by the receiving program (the NCP should not wait for the allocation to become zero before sending the new allocate).

We also observed that small files are transferred at a significantly lower transfer rate than large files but beyond a file size of 40 Kbits, the file size makes little difference in transfer rate or processing cost per bit transferred. The figure of 40 Kbits is probably related to the size of sending and receiving buffers used by the programs. In general, for most practical values of buffer size, the larger the buffer size and allocations, the faster and more efficient will be the transfer. Unfortunately, large NCP buffers are not easily available in many systems and come at a premium. The information on average file size (240 Kbits for Image and 40 Kbits for ASCII files) may be helpful in optimum allocation of buffer space.

V. REQUEST FOR COMMENTS AND SUGGESTIONS

It is hoped that the above measurement results and our FTP and SURVEY measurement facilities will help ARPANET users plan their modes of Network usage and help Network programmers in making the Network better. This RFC is indeed a Request For Comments and your suggestions on the way we collect, store, and display measurement data will be greatly appreciated. We can break the measurement data by Hosts and will be happy to provide the information if it is considered desirable. Please let me know what other parameters we should record or display. You may communicate with me via the ARPANET (AKB at MIT-DMS (Host 70), NIC Ident AKB), via telephone (617-253-1428 or 1449), or US mail (Rm. 208, 545 Tech Square, Cambridge, Mass 02139).

[This RFC was put into machine readable form for entry]
[into the online RFC archives by Robert Baskerville 9/98]

