Distribution of software in academic and research network: overview of recent trends

An unsolicited note

prepared by:
L.Chiappetti, IFCTR

23 May 1991

For distribution to:

SAX GSWG: all members
ASI SAX GS Project Manager
Telespazio (via ASI)
SAX DAWG coordinator

Table of content

1. Purpose of this note

Situation in academic and research network is subject to rapid evolution. A previous technical note (*Academic and research networks. An overview.* June 1989. *SAX/GS/IFCTR/TN-001*; hereafter *TN-001*) gave a comprehensive description valid at the time it was written.

However since then, things have changed in the world and in Italy. More than re-issuing the original note, it has been felt useful to provide an overview of the major changes (essentially due to the large diffusion of the Internet), at the same time focussing on the tools commonly used for distribution of information (public domain software, documentation, data etc.). In particular the latter point is felt relevant in the SAX context in view of the Mission Support functions attributed to the SDC (distribution of documentation, software, calibration data etc. to the general observer community).

2. Internet is spreading

One of the more striking facts of the evolution of academic networks in the recent period is the fact that the Internet, and the suite of Internet protocols, is rapidly spreading in the research environment.

Internet was described in section 2.2 of *TN-001* and we would not recall here any technical detail. The main Internet network services (sometimes called *ARPA services*) are *telnet* (remote login), *ftp* (file transfer), *smtp* (mail). *DNS* (the Domain Name Service) shall not be forgotten. These high level layers rest upon a suite of low level protocols, broadly known as *TCP-IP*.

TCP-IP and ARPA services are native on Unix systems; however they were originally designed for interoperability* among machines of different vendors, and this is actually their major strength, and the reason for their widespread and rapid diffusion.

One may incidentally note that also Unix systems (fostered by the RISC hardware platforms, and by their relatively low cost) are becoming more and more widespread: this is a further boost for the diffusion of the Internet.

Among the other reasons which may be at the ground of more and more people moving to the Internet, one has facts like:

° Internet is providing a wider suite of services w.r.t. a network like Bitnet.

° Internet connections are generally faster (in terms of bit rate) than Decnet connections.

2.1 in Italy

The number of Internet connections and sites in Italy is increasing. Major computer centres (CNUCE, CINECA, CILEA) have now a number of Internet nodes. Internet is one of the protocols running on the high speed national backbone funded by the GARR (Gruppo Armonizzazione Reti Ricerca). It is also strongly supported by CNR (for instance telnet is now replacing the IBM PVM "passthru" protocol as the principal mean of remote login; CNR.IT is a noationwide domain, with local subdomains like . MI.CNR.IT, BO.CNR.IT, CNUCE.CNR.IT etc.).

A top domain .IT is officially defined; all sites having network connections in Italy (including those in Bitnet or Decnet) should have a valid domain address by which they can be reached through a suitable gateway. (In fact for Bitnet nodes it is largely transparent to the user whether mail travels over Bitnet or Internet stretches of the network).

It is true however that, while Internet is faster

* Note that interoperability does not mean portability, nor does it imply transparent access to services under a common interface.

then Decnet and offering more services than Bitnet, it requires a larger effort to be managed (particularly in the setting up of name servers), and that there is a certain lack of manpower in this area.

For more details, a good contact is A.Blasco Bonito of CNUCE (*blasco@cnuce.cnr.it*), responsible of the CNR department for Network Infrastructures for Research. A mailing list (TCP-ITA@ICNUCEVM) is set up to discuss Internet problems in Italy at the Pisa LISTSERV.

2.2 in USA

Internet is now definitely the largest network in the USA even in the academic environment (note that US Internet supports also non-research government, military and commercial organizations).

More and more sites and persons, even in the astronomical community, are moving from Decnet to Internet. In fact NASA itself has now officially disbanded SPAN (the NASA-supported Decnet), and is strongly supporting NSN (the *NASA Science Internet*).

It has to be noted that the exact details of what this "disbandment" means are not known. Previous Decnet nodes are still operating, and the overall Decnet named ESNET (merge of SPAN and HEPNET) is still active. It is expected that more details be known at the Joint SPAN/ESIS Users' Meeting (which will take place in October 91 in ESRIN).

It is also notable the large diffusion of *ftp anonymous servers* as a way to distribute public domain software. Major US astronomical sites (like NOAO, the National Optical Astronomical Observatories, and STScI, the Space Telescope Science Institute) have adopted this mean for the dissemination of software, documentation and data.

2.3 in Europe

The situation in Europe is less clear. Some areas (like e.g. Scandinavia and Germany) have clearly moved to the Internet. Almost all countries (including Holland and France) have officially defined a top national domain (.DE, .NL, .FR), and a majority of Bitnet and Decnet nodes can also be reached using Internet-style domain addresses (even if they may not be full-fledged Internet nodes).

Concerning ESA, it is still supporting two main channels (Bitnet for internal mail exchange using IBM PROFS, and SPAN for communication with the scientific community and support of projects like ESIS). There are however rumours of movement to Internet. Here too it is expected that more details be known at the announced Joint SPAN/ESIS Users' Meeting.

2.4 in the SCS

We briefly summarize here the current networking connections (with emphasis on, but not limited to, Internet links) of SAX sites, as they are known to the author.

Milanoarea

The VAX 8250 of IFCTR is connected to a LAN extending (via Trans-LAN bridge) to various CNR sites (Area della Ricerca), and including a number of Unix workstations in the Institute. This LAN offers the possibility to exit to Decnet (via a router at SIAM connected to CILEA) and to Internet (using the SIAM IBM3090, also Bitnet node IMISIAM, as gateway; again the connection is thru CILEA to the national high speed backbone).

The VAX is the only machine publicly known to Decnet (node number 38.698), while one Sun workstation and some PCs have Decnet "hidden" numbers. All workstations, the VAX, and some PCs have valid Internet numbers (192.65.131.nn, with a reserve of 20 numbers for IFCTR), and a name in . MI.CNR.IT.

Internet protocols are widely used for local connections (among workstations and with VAX), and for remote connections (more than 40 Mbyte of software has recently been downloaded in a very satisfactory way from the USA).

Internet usage among Unix workstations (including usage of the NIS, or "yellow pages", to handle the necessary databases in a joint way) is very satisfactory. Some problems in the Internet connectivity of the VAX are due to the fact that the Digital TCP-IP emulator (UCX) implements only a subset of the functions. For this reason for instance the VAX needs two domain names, one used for telnet and ftp, and the other one for routing mail thru a gateway. The gateway is currenlly located at CNUCE (using PMDF), but installation of PMDF is due on a VMS machine at SIAM (a non-Digital TCP-IP emulator, TGV Multinet, officially recommended by CNR, will be used). This machine is also planned to run DNS. When this will be implemented, a number of subdomains will be set up on Institute basis (this way all nodes at IFCTR will have names like: host.ifctr.mi.cnr.it).

Bolognaarea

The situation at ITESRE is similar to the one described for Milano. Machines are connected into an extended LAN managed by the Area della Ricerca jointly among CNR Institutes.

The number of publicly known Decnet nodes is however larger, including more clusters with VAX 8250, MicroVax and Vaxstations. A number of Unix (including Ultrix) workstations are also connected to the LAN. The wide-area Decnet and Internet connections share the same line, using a multi-protocol CISCO router.

Most machines have Internet numbers and names assigned to them in the .bo.cnr.it subdomain (it has to be noted that the mail gateway for this subdomain is currently at CNUCE too).

As far as the author knows, the situation of connectivity is also quite satisfactory, with similar problems concerning the TCP-IP emulation on VAX.

RomeFrascatiarea

IAS Frascati has a number of Decnet nodes (including the VAX 8250) and an IBM mainframe connected to Bitnet. The wide area Decnet connections are ensured by a router to INFN, but are quite unstable. The situation of Internet connectivity is unknown to the author.

IOA Rome is Decnet node ASTROM. The situation of Internet connectivity is unknown to the author.

Palermoarea

IFCAI Palermo has a number of Decnet nodes (including the VAX 8250), connected via a router to CERE. The situation of Internet connectivity is unknown to the author (it however known that other Internet nodes exist in Palermo).

Utrecht

An Internet domain sron.ruu.nl is defined, and this form of addressing has now replaced the previous Bitnet connection. It has to be noted that the main machine at SRU is a Data General, which provides TCP-IP emulation in a slightly unusual way.

The situation of Decnet connectivity is unknown to the author.

SSDESTEC

SSD has access to all network services on the ESTEC LAN (including an IBM mainframe). The main protocol used is Decnet, particularly on the EXOSAT VAX cluster.

The situation of Internet connectivity is unknown to the author.

3. Tools for software distribution over the network

We give here first a logical overview of the ways to distribute software packages (but also related documentation, or calibration datasets), followed by the description of some working examples.

3.1 logical overview

One may consider the following requirements for the distribution of software (programs, documentation, data):

a) occasional distribution of one or a few (small) files

b) distribution of major releases of a large software package. This will be distributed in bulk, possibly in compressed form, and will need some installation procedure at the receiving site. This distribution will occur either once (to give the receiving site the whole package) or seldom (once/twice per year typically).

c) distribution of updates. This may occur as a major release (see above) in bulk, or otherwise the user may want to retrieve updates to individual modules only (possibly after having been notified of a change).

One may consider the following procedures to implement the above requirements:

a) using e-mail. An human agent, or a software server will send the files encapsulating them in an e-mail message, after a request by the user. Suitable for small text (source programs, ASCII documentation) files. Works easily across network (with size limitations). Does not need any authorization (except for sufficient quota or spool space at the receiving end) to reach the end user.

b) sendfile model. Logically similar to above, but suited to generic files (binary, any size). The idea is that the user issues a request (e.g. by mail), and this is serviced asynchronously (again by a human or a program) by sending files to a spool or scratch area at the user site. This does not require any authorization, nor that the user be logged on (this is extremely comfortable). The server could take care of keeping a log of the files distributed.

This is natively implemented by IBM SENDFILE, and could be implemented in Decnet (and perhaps uucp ?) provided a publicly writable area exists at the receiving end (e.g. the FAL$SERVER directory in VMS); this may not be desirable for security reasons, and will require special s/w at the receiving site. This cannot unfortunately be set up to work with ftp (wants a password at the receiving end).

c) remote copy. Naturally implemented in Decnet. It is bilateral and synchronous. Either the user at the receiving end retrieves the file from a remote node in his own area (no authorization needed provided the origin area is world-readable), or a server at the other end does a copy, provided it knows a password, or there is a publicly writable area (the relevant security problems are described in b above).

Cannot be implemented with plain ftp (wants a password for the connection). Logging of the transactions cannot be controlled.

d) anonymous ftp model. Implemented at some Internet sites. The user at the receiving end connects without need of a password (this could be implemented in VMS by a captive account, and is provided in Unix by "anonymous ftp") and retrieves the files he wants. In anonymous ftp, one generally provides an identification (his name, or the Institute name) instead of the password. I ignore how this information is used for logging, and I also ignore whether anonymous ftp can be setup natively under Unix, or whether it requires additional software.

3.2 practical implementations

LISTSERV

The Revised LISTSERV is described in section 5.2 of TN-001. It is an asynchronous server, with operations initiated by e-mail or interactive requests. Files are then sent back via sendfile, or e-mail. The e-mail modality is the only one possible across different networks, and shares the usual limitations (only ASCII, max 80-byte records).

Among the useful features the Automatic File Distribution or File Update Information in case of updates.

remotecopyof singlefiles

This way can be used (it is one of the possibilities offered by NOAO and STScI) to retrieve over Decnet a small number of files. It requires some previous directory browsing to known what to retrieve.

It also requires typing often long and awkward file access strings (node, disk, directory, file name), prone to errors. Works only for Decnet nodes (there may be problems with nodes using Decnet emulators in case of binary files).

It is also in common experience that Decnet connections are quite slow.

remotecopyof backupsavesets

In the case of distribution of a large software package, sometimes the distributor prepares a VMS BACKUP save set on disk, and the remote user can this way retrieve the entire package, issuing a single COPY command (e.g. in a batch running in off-peak hours).

In case of large datasets the transfer may take a long time, and also require lot of scratch space to store the save set before being unpacked. Obviously this methods works only with VMS Decnet nodes at both ends.

anonymousftp

Anonymous ftp is widely used for retrieval of public domain software (also in the astronomical community: e.g. IRAF and STSDAS are offered this way).

In order to have easy access, it is necessary to set-up a well organized directory tree on the server (e.g. by cathegory or sub-package). Possibly each subdirectory shall contain a README file with a description of the content and instructions for retrieval. A good example of such organization is the STEIS service at STScI.

Ftp requires TCP/IP to be running at the receiving site, and is generally able to handle also binary files. Normal access is interactive. Transfer of more than one file can be handled via the mget command (if available) or using local batch facilities.

anonymous ftp of (compressed) tar files

A way used (at least by NOAO and STScI) to distribute a large package in bulk, is to write it to a tar file (using Unix tar utility). This is similar to the BACKUP solution described above and shares similar problems (one need large scratch space, and access to tar).

In order to alleviate the load on the network, compression techniques can be effectively used: the following examples are based on practical experience with NOAO and STScI.

Documents can be kept in PostScript files (ready to be printed), and such files can be compressed with the Unix *compress* utility. At the receiving end they can be uncompressed and printed "on the fly" (without extra space), using Unix pipes.

The following technique is used instead for tar files. They are compressed with the Unix *compress* utility, and split into chunks (say of 0.5 Mbyte each) to reduce the probability of files corrupted during transmission. One can then retrieve each chunk separately, verify correct reception, concatenate the chunks, uncompress and de-tar them (the last steps can be joined by Unix pipes). In order to check errors, a small file with checksums is provided, and the user should verify the checksum of the received files (using Unix *sum* command).

Compression is very effective. This way I succesfully loaded the equivalent of 21 Mbytes just by transferring 11 files of 0.5 Mbytes each. With data files compression is even more efficient (up to factor 8).

application to non-Unix systems

The procedure described above appears to be limited to the case of Unix systems at *both* ends. This is not exactly true, as tools to emulate standard Unix commands on non-Unix systems are available. Here are some examples:

° an emulator for tar exists for instance under IBM VM

° a "portable" tar (called *rtar* and *wtar*, particularly suitable for exchange between Unix and VMS systems) exists in the framework of IRAF.

° NOAO and STScI offer utilities (in executable and source form) to emulate under VMS the Unix *compress/uncompress* and *sum* commands.

4. Conclusive recommendations

It is therefore suggested that, in the framework of the SAX SDC, attention is paid to the need of Internet connectivity, considered the strong trend towards Unix and Internet existing in the research community. This might also include a selection of TCP-IP emulators under VMS consistent with choices officially supported at national level.

In particular for what concerns the distribution of software programs, documentation and calibration data, one of the possibilities offered should include the use of ftp anonymous servers. It is suggested to gain experience with this sort of tool, and with the working examples existing in the astronomical community.