JANET - the United Kingdom Joint Academic Network

Mike Wells

Paper presented at the UK Serials Group Annual Conference, Leeds, March 28-31, 1988.

History

Since the early 1970s the United Kingdom academic community has been moving towards the use of a single data network linking all Universities, Research Council sites, and other institutes of tertiary education. The first proposals arose from a small working group which I chaired, and which in 1973 pointed out the benefits possible by simplifying access to computer systems, which were then seen as being expensive to install and to operate. Over the next ten years a number of networks were developed and brought into service, but until 1982 it proved impossible to unify these into a single system. Finally, in the winter of 1982, the Computer Board, which is responsible for providing funds for University Computing Services, took the bold step of agreeing to underwrite the costs of running a single network, and sought a Director of Networking to manage the transition from a number of separate networks serving different sectors of the academic community to a single network, serving the whole community. In April 1983 I accepted an appointment as a part-time Director of Networking.

My brief was straightforward. I was to formulate objectives for a single network, and to devise suitable management structures to ensure that the then separate networks, which were owned and operated by separate agencies, were brought into a single structure. All this was to be done without disrupting existing services, and the resulting network was to cost no more than the several separate networks. I would like here to pay tribute to those who constructed these networks, for their skill and perseverance in creating the networks, for their forbearance in seeing them taken over by an alien group and for their very positive attitudes towards the unified network which has resulted.

Mike Wells *is Director of Networking, Computer Board for Universities and Research Councils*

Objectives

The objectives established were to provide an integrated network offering the following facilities:

terminal access ;
file transfer ;
electronic mail ;
job transfer ;

To understand these it is essential to realise that the network is a wide-area network which serves only to interconnect the local-area networks at each connected site; there are NO terminals and NO host computers attached directly to the wide-area network as such. Thus terminal access is to be understood as from a terminal on the local-area network on one site to a host computer on the local-area network on some other site. Similarly, file transfer is between a host connected to the local-area network on one site to a second host connected to the local-area network on some other site. The wide-area network is thus simply a 'bearer' network between local-area networks and is in no way involved in the details of the type of work being transported. I shall return repeatedly to this point, which is central to the way in which the network is operated and managed.

The wide-area network is usually referred to as the 'Joint Academic Network', or JANET, and from now on I shall use that name. In practice, the term JANET is also often used to refer either to the physical network, or to the community who actually make use of the network; the context usually makes it clear whether the reference is to the network proper, or to the user community. With one exception, the capital and running costs of JANET are covered by a major grant from the Computer Board, with supplementary grants from the Research Councils, and there is no accounting or charging for use of the facilities. The single exception concerns the provision of connections to some Polytechnics, where a charge is made for the connection. This decision is quite deliberate, and is based on two motives. Firstly, the inclusion of accounting and charging represents a substantial extra complication, and the actual collection of revenue is very expensive of effort. Secondly, the whole purpose of creating JANET was to encourage the use of networking within the academic community and the introduction of charging would have been counter-productive.

JANET does provide access to certain chargeable services, of two fundamentally distinct types. The first of these is for access to other network services, such as British Telecom's PSS (Packet Switch Stream) service, where the supplier, in this case BT, bills the user of the network by virtue of his address on the BT network. One of the services offered via JANET is access to PSS, and it is therefore necessary that such access is made via an 'accounting gateway', which can ensure that the JANET user has an account on the gateway, and can generate suitable accounting and charging records for the use of PSS. The second type of access to chargeable services is to host-related services, such as access to various commercial databases. In these cases, the supplier bills the user by virtue of his identity on the host system, and there is no requirement on JANET to apply controls or to generate accounting or charging records. We shall see that this distinction between the identity of an entity on the network, and the address of an entity on the network is valuable in a number of other situations.

The access to other chargeable networks such as BT's PSS, and via this to the International Packet Switching System, IPSS, is greatly simplified because as a design objective JANET uses standards which are compatible with PSS and IPSS. The standards are not identical with PSS/IPSS, but full inter- working is possible. There are, of course, other networks which do not use these same standards, and some of these are of great interest to the JANET community. In particular there are gateways between JANET and the European Academic and Research Network EARN, between JANET and USENET, and between JANET and ARPANET. These gateways handle predominantly file transfer and mail, and in some cases do not support terminal working or job transfer; indeed in some cases terminal working is not possible due to the nature of the second network.

Network Structure

JANET is a packet-switching network. It uses eight separate switching centres, each housing a GEC 4000 series computer which performs the switching function. Each switching centre is connected to several other such centres by 'trunk' connections; each site is connected to one of the switching centres by a 'link' connection. Both the links and the trunks use leased lines, and there is no physical distinction between the links and the trunks. The trunks operate at either 48 kilobits/second or 9.6 kilobits/second, whereas the links mostly operate at 9.6 kilobits/second. The switching centres are all located at University or Research Council sites, and each switching centre is operated under a contract with a small group, the Network Executive, which is responsible for the running of JANET. The Executive deals with GEC who are the supplier of both the switch hardware and the associated software; the Executive actually generates the software for each switch and oversees the installation of new versions of the software. Detailed changes to the hardware configuration, for example the addition of new sites to a switch are overseen by the staff at the switching centres. The local staff are also responsible for the day to day liaison with GEC's hardware maintenance staff, and with British Telecom's staff who maintain the leased lines.

The JANET bearer network uses the 1980 version of the X-25 protocol to interconnect the campus networks at linked sites. There are some minor deviations from the version of X-25 implemented on British Telecom's PSS service, among which is the ability to allocate a range of X-25 addresses, rather than a single address, to a gateway. This is designed to allow each site to be allocated a range of addresses which it can, in turn, allocate to entities connected to its own local-area network. In this way, the management of addresses on each site is delegated totally to that site, and the overall JANET management is freed from any involvement with detailed address management on a site. An X-25 address is presented as twelve decimal digits, possibly followed by a two decimal digit sub-address, and this form is used within the bearer network. However, much of the software associated with the systems connected to JANET uses a 'transport service'. The transport service is presented as an additional service running over the X-25 service, and one of its major functions is to provide a much more flexible form of addressing, which takes the form of a highly structured character string with a number of fields. Each field can be used as a mnemonic, and the whole address can be readily understood by a human user, the whole being capable of mechanical processing to yield an X-25 address.

Transport service addresses are very important in the management of JANET; in particular, the use of transport service addresses allows a decoupling of the allocation of physical addresses from the logical addresses of entities. This has two benefits. Firstly, addresses can be assigned by local management, without reference to any form of centralised authority, whose role is restricted to allocating a range of addresses to each site management for its own use. Provided a site management does not allocate addresses outside its proper range, there can be no possibility of assigning identical physical addresses to different logical entities. Secondly, a site management can announce the transport service form of an address once and for all, but can, at its discretion, reassign the physical (ie X-25) value of the address as it chooses. To cater for this, a 'Name Registration Scheme' (NRS) operates, which provides a service to the whole network. The scheme allows a duly authorised, NRS administrator, at each site, to assign the X-25 address which is referenced by a transport service address. All users, and all software authors, are encouraged to refer to an entity by its transport service address, and to rely on the transport service at each site to carry out the process of conversion to an X-25 address. The advantages in flexibility, and in allowing a site to manage its own addressing structures, are considerable. We shall see later that the mechanism also provides a powerful tool in managing the introduction of new protocol implementations.

Types of Heterogeneity

The JANET network, and the community it serves are both heterogeneous. This applies to practically every aspect of the situation:

Supplier
Technology
Application
End user
Management
Funding

'**Suppliers**' covers at least 15 different suppliers of host computer and at least 25 different operating systems running on these hosts; it also covers six different suppliers of networking equipment.

'**Technology**' refers to the different network technologies in use for the wide-area network and for the campus local-area networks which include X-25, Cambridge Ring and Ethernet as well as character switches such as the Gandalf.

'**Application**' covers an enormous spectrum. Many of the applications use networking in a purely incidental way, in the sense that the user happens to be at one site, while the serving host is at another. The applications working in this way range from scientific and engineering calculations, through to work in arts and social studies. Other applications, such as file transfer or electronic mail arise only in the context of a fully connected network. Similarly the '**end user**' category spans a range from undergraduates, via research assistants and senior academic staff across to technical staff and administrators. Naturally this catholic mixture of users further widens the range of applications.

The JANET bearer network requires a quite specific style of **management**, and this must co-exist with the very different management structures which are found in different Universities, Polytechnics, and Research Council establishments. Finally the **funding** arrangements for the JANET activity, and for the sites connected to it constitute a rich cross-section of the ways in which UK academic establishments are funded. A university receives most of its funding from the University Grants Committee (UGC), a Polytechnic is funded by its local education authority (LEA), and a research council's site is funded by the respective council, who each in turn receive their funds from the Advisory Board for Research Councils (ABRC). In the final analysis the UGC, the LEAs and the ABRC are all funded by the Department of Education and Science (DES), but the different funding lines inevitably generate different priorities. In fact, this convergence of funding lines represents one of the two facets of the JANET community which are in any sense homogeneous, and it has had important bearing on our ability to create JANET in the first place. In the United Kingdom, just as in other countries, it is necessary to have a licence to operate a network which transmits data for entities other than the 'owners' of the network. The fact that all sites are funded whether directly or indirectly by the DES has raised the possibility that the JANET community is a single legal unit. At the same time the complexities of the funding arrangements for those who exploit JANET has meant that no licence has yet been issued!

With all this heterogeneity one is tempted to ask what, if anything,is not heterogeneous. The answer is, in a single word PROTOCOLS and much of what follows will be greatly concerned with this topic. It is a highly technical subject, and contains a rich vein of jargon. I hope to stay at a safe distance from both technicality and jargon, without in the process losing accuracy or becoming verbose.

Software Structures

In discussing network software structures there are two jargon words in widespread use, and the distinction between the two is central to the production of usable networks. The two words are INTERFACE and PROTOCOL.

As in any other environment, a large task is broken down into many separate sub-tasks, until the sub-tasks are sufficiently simple that they can be reasonably understood by the person implementing that sub-task. Normally this process of sub- division is left to the whim of those in charge of the products. However, in a networked environment it has been found that there is enormous value in carrying out the sub- division once and for all, with each group implementing the product using the same sub-division. (This is a gross over simplification; in particular the actual implementation need not, and in general does not, reflect the sub-division I shall outline, but the implementation behaves as if it DID reflect it)

The sub-division is initially based on the extent to which each task needs a detailed knowledge of the details of the actual mechanism of the data transmission medium. The separate tasks are referred to as 'layers', with the 'lowest' layer being most aware of the physical reality of the transmission medium. The lowest layer at the two points which will intercommunicate is connected via the physical medium, and only this layer can actually transmit or receive data. Each successive layer offers a service to the layer immediately above it in such a way that the higher layers gradually approach more and more the task which the entire product is to achieve. The ultimate purpose of this structure is to achieve communication between physically separated points, and the structure is such that a layer at one point communicates only with the layers immediately above and below it at that same point, with the local system in which it is embedded, and with the same layer at the remote system. This last communication is achieved by transmitting down through the layers, across the medium, and back up to its corresponding layer at the remote site.

The exchange of data at a given site, up and down the layers, is governed by interfaces; the exchange of data between corresponding layers at separate sites, is governed by protocols. To a first approximation, a protocol defines every possible valid message between the corresponding layers at separate sites, and also defines the actions to be taken when invalid messages are received, or when there is a failure in one of the lower layers. Each layer's protocol assumes that the layers below it will provide a set of network services, and the protocol in turn will enhance these and offer to the layer above a service which more closely approaches the needs of the user. These days, protocols are usually defined by international standards bodies such as ISO and CCITT, in some cases by adopting existing proprietary protocols, but more usually by amending existing protocols in the light of experience. The standards bodies do not, in general, define the interfaces to be used in the up-down transfers between layers at a given site, or the means by which a layer communicates with its local environment. Indeed the standards bodies are careful to emphasise that the behaviour at a site is entirely for that site to determine, provided that it is consistent with the behaviour expected of the set of protocols when observed from outside.

At the time when the first networks using non-proprietary protocols were being developed in the UK the academic community was obliged to define its own standards. This work was supported by the Computer Board and the Research Councils, who jointly funded a small unit to assist in the standards work, which was in part based on the activities of working groups associated with the GPO's 'Experimental Packet Switched Service', EPSS. As the protocols were specified they were published in booklet form, each bound in a distinctive colour; in time the protocol became known by the colour of its binding, and the set of protocols became known as the 'Coloured Books'. The major protocols were:

YELLOW BOOK	Transport Service	transport service over X-25
GREEN BOOK	Terminal Access	terminal access to a network
BLUE BOOK	File Transfer	transfer a file between hosts
GREY BOOK	Electronic Mail	mail message to a remote host
RED BOOK	Job Submission	submit a job to a remote host

This list is not exhaustive, and there are other protocols defined to support the use of transport service. The fact that there were full definitions of these protcols, and that in many cases working implementations were available, made it possible to proceed with the integration of the separate networks which preceeded JANET in a way that would otherwise have been quite impossible.

Network and Application Relays

The structure of protocols and interfaces in the previous section make it possible to have two quite distinct types of linkage within a network.

One of these is concerned solely with the transmission of data, while the second is concerned solely with the transmission of an application. They are referred to as a 'network relay' and an 'application relay'.

A network relay operates on the lower layers, typically up to layer 3 or 4, and serves to interconnect two different implementations of the lower layers, or possibly two similar implementations which operate under different management. The network relay will take no notice of the contents of the data it transmits, provided of course that it represents valid packets at the highest layer it handles.

An application relay operates on the upper layers, usually from layer 4 upwards, and serves to interconnect two different implementations of the upper layers, or possibly two similar implementations of the upper layers, which operate under different management. The application relay is unaware of the presence of the lower layers, provided of course that they function correctly in moving the data into and out of the relay.

The present structure of JANET, as a set of local-area networks using a wide variety of different technologies, and with each local-area network operating under separate management, relies very heavily on network relays to manage the traffic, and to provide a mechanism for separation of the management of the many local-area networks from the single wide-area network. However, to a first approximation there are no application relays within the present JANET structure, since all the connected sites use the same higher level protocols. We shall see later that a requirement for the introduction of application relays has been identified in managing the introduction of new protocols.

Acquiring Protocol Implementations

The early success of the academic wide-area networks in the UK was achieved by the community's success in introducing a common set of non-proprietary protocols. This required two steps, neither of which is cost-effective when viewed in isolation. These steps are the definition of protocols, and the acquisition of working implementations of the agreed protocols.

Agreement on a protocol is never easy. It requires compromises as to what is desirable, and what is practicable; however, agreement is easier when the parties to the agreement have minimal commercial involvement in the outcome of the debate.

In the early stages of the move towards UK academic community protocol standards it was felt that the standards were sufficiently fluid that it would be unrealistic to expect computer manufacturers or software houses to provide implementations, and the first products were written within the community. It was soon realised that although this made it possible to modify the protocol definition in the light of experience gained in the field, it carried the penalty of a continuing support load which the community was (and is) ill equipped to bear. It quickly became clear that it was preferable to acquire protocol implementations on the same basis as other software products, as a fully supported part of the software configuration on a host computer.

The acquisition of implementations of agreed protocols is even less easy than reaching agreement on the protocol. Remember that a protocol implementation is generally closely linked to the operating system of the host computer on which it will run, and indeed is often implemented as an extension to the operating system. Thus, on any given occasion, a manufacturer will be able to maintain that, either his own existing realisation of an application is cheaper or more efficient, or that the cost of producing an implementation to the new standards is so high that the customer would not want it anyway. Frequently he will maintain both points simultaneously. The answers to these points are as follows. The manufacturers' own application product will, of course, interwork very well with his own system; ask him how well it will interwork with other suppliers. As regards price, the answer is to make it clear that the implementation is sought not as a special product on one occasion, but that all future orders from that manufacturer will include provision of the product, and further that the fact that he has the implementation as part of his product line will count in his favour.

This rather robust approach can only be made to work if there is what amounts to a buyers' cartel operating. This was realised nearly ten years ago by the Computer Board and the Research Councils in the UK, who established a 'Joint Network Team' whose range of duties included a vigorous programme to encourage the development of protocol standards, and their subsequent implementation. They were helped in this programme by the Computer Board's willingness to fund work on protocol definitions, and on their subsequent implementation within the academic community. However, much more important was the Board's willingness to insist that major procurements of new host computer systems could only proceed if the supplier was prepared to undertake to provide and support implementations of these same protocols. The Board's success in this is a tribute to the pressure which a well-organised customer can apply to even the largest supplier.

The situation was that, by 1982 there was a significant proportion of the UK academic community which had access to working implementations of protocols for transport service, terminal access, file transfer, electronic mail and job submission. Since that time the number of suppliers who have completed implementations has risen still further as shown by the table in Appendix 1. However, this very success has created a new set of problems. The JANET community now has working protocol products for the commonly used services, but the protocols are those of the JANET community, and are not in use elsewhere. The standards bodies are now beginning to deliver definitions for new protocols to replace those in the JANET community, and some suppliers are beginning in turn to offer these as product line items. This whole process will take at least five more years, and in the meantime the JANET community must operate both the old and the new protocols. Before discussing how this will be achieved, I want to look briefly at one other aspect of the way the JANET community runs its affairs.

User Groups

One of my early actions was to institute a user group for the JANET service. There are two difficulties in creating a user group for such a widely spread service, used by such a large community. One is to identify the users, and the second is to create a user group structure which does not involve all those involved in large amounts of travel. To put things in perspective, I estimate that JANET occupies an area about 1000 kms by 600 kms, with about 100 sites connected, about 800 host computers, about 10,000 terminals, and has about 250,000 potential users. Clearly a single user group is out of the question. It was therefore decided that the user group would be regionally organised, into a total of seven geographical regions and each site was invited to join in the region which would be most convenient in terms of travel. Each regional user group is free to organise its own affairs, with some exceptions which will be mentioned shortly. These regional user groups each send two members to a national user group. The chairman of the national user group is automatically a member of the Network Advisory Committee, which is the governing body for the network.

There are, in fact, two quite distinct types of person 'using' JANET. There are users who are ordinary members of academic departments, or who work in a Research Council laboratory as physicists or biologists. These are 'end-users'. Notice that because of the way JANET works - by interconnecting campus networks, and because it is to the campus network that these end-users connect - they do not use JANET directly. There is also an intermediate class of 'users', namely those responsible for operating the campus networks to which end-users connect, and who 'use' JANET in order to achieve interconnection of the campus networks. 'Users' in this second group are likely to be very interested in networking, since that is how they earn their living, and therefore very keen to come to user group meetings. At the same time, an end-user's perception of the quality of service offered by JANET is likely to be coloured by the quality of the campus network, and by the level of advice and support which he receives from his site's network staff. It is thus essential to involve these intermediate users of JANET in user group meetings, without letting them totally dominate the meetings. It is in this respect that the terms of reference for the regional user groups were carefully regulated. I had anticipated that this regulation could not be replicated at the national level; in the event I am happy to say that I was wrong, that regulation was in any case not necessary, and that a number of regional user groups have ensured that both end- users and intermediate users go forward from the regional to the national level. The first Chairman of the national user group is in fact an end-user, being a Professor of Geography at Glasgow University.

Current Management Issues

Currently there are two major management issues. The first of these is an upgrade of the basic trunk network which interconnects the JANET switches, and the links which connect each site to its switch. The trunk network at present uses either 48 kilobits/second or 9.6 kilobits/second leased lines. Part of the trunk network is to be replaced by a set of 256 kilobits/second circuits implemented as virtual circuits on 2 megabits/second carriers. This requires some enhancement to the hardware of the switches, and the installation of suitable multiplexers at the switching centres. The requirements for testing and phased introduction of the new systems, while maintaining continuity of the existing services, present a number of problems in the technical management of the network, but overall the way forward is reasonably straightforward. At the same time, the links between sites and the switches, which for the most part run on 9.6 kilobits/second circuits will be progressively upgraded as the loads expand to run at 48 kilobits/second. Again, this is a straightforward technical programme, requiring only careful supervision.

A much more challenging problem is the transition from the present, community generated protocols, to the newly emerging international protocol standards. These new protocols will affect the JANET network in terms of the lower levels of the protocol stack, and will also affect the JANET community's host systems (connected via their associated local-area networks) in terms of the upper levels of the protocol stack.

The changes to the wide-area bearer network have been public knowledge for some time, and in essence involve a move to the 1984 version of X-25. There is a close correspondence between the present version of X-25 and the new version, and the transition on the wide-area network is again well understood. However, this change must be made in such a way as to accommodate the situation in which the gateways between the wide-area network (JANET) and the sites' local-area networks do not all simultaneously change to the new protocol, and in which the wide-area network and the local-area networks may, for a period, be using different protocols. Clearly, this is exactly the situation which is handled by a network relay, where the same high-level protocol is carried by two different low-level protocols offering a similar upper interface.

The transition from the present application protocols to the new protocols is far more complex, and is generating a substantial technical as well as managerial problem. A working party has studied these, and its proposals are now being circulated within the JANET community for discussion. The transition raises two inter-related issues:

(1) what is the mapping between the features of the old and the new protocol, for example between the present protocol for file transfer (the 'blue book') and the new protocol for file transfer and file access (FTAM) ;

(2) how can we manage the situation in which a user, with no knowledge of which protocol is in use on a remote host, can successfully communicate with that host.

The main means of managing the transition will be via the name registration scheme. As explained earlier the NRS provides a single centrally regulated index of the network addresses of hosts on the various local-area networks. In fact the NRS goes further than this. The names registered are not simply those of a host, but of a particular service on that host; thus the name for the mail server on a host will indicate this fact and will yield a different network address from, say, the file transfer server on that same host. The contents of the NRS are to be expanded, so that when the NRS is searched for the network address of a particular service on a particular host, the result yielded will depend on the context in which the search takes place.

Suppose a host, using 'blue book' file transfer, wishes to find the network address of a second host; if the second host also uses blue book, the NRS will return the network address of that host's blue book file transfer server. However, if the second host uses FTAM, the NRS will return the address of an application relay, which is capable of carrying out the mapping between blue book and FTAM. The first host will now send its file transfer to the application relay, together with the NRS name of the host to which it actually wishes to make the file transfer, and the application relay will translate the blue book file transfer into an FTAM transfer. The details are rather more complex, since during the transfer process there may well be negotiation about details of the transfer, and so the application relay must be prepared to act as an intermediary for messages in either direction between the two hosts.

A much more detailed account of the proposals for transition is given in the March, 1986 'Report of the Academic Community OSI Transition Group', which is available from the Joint Network Team. This gives details of the proposed mappings between the present JANET community protocols and those ISO protocols which have been defined, and an account of the use of the Name Registration Scheme.

Summary

I would like to emphasise four main conclusions from my work with JANET.

(1) Heterogeneous networks can be built and can be made to work. Further, the advantages to the end-user, of being able to retain existing systems, and to have freedom of choice when acquiring new systems, will substantially outweigh the advantages to the service providers, of dealing only with a single supplier. The problems of maintenance in a

mixed supplier system tend to be overstated.

(2) A heterogeneous system typically involves a number of intermediary service providers, who are involved in converting a raw, lower-level service into an applicable, higher-level service. These intermediate providers need to have clearly established relationships with other providers, and with end-users. In particular, they must be included in, but must not be allowed to dominate, any form of user organisation.

(3) Users, and intermediate providers, need to have access to information about network availability and loading. If they do not, they will invent such information, and will project onto the wide-area network (for which they have least information) all their problems of poor response or inability to connect. By contrast, if information is made available, users may even vote for further expenditure on network provision.

(4) As in any other activity, the successful realisation of a project requires that those affected are whole heartedly committed to seeing it succeed. In this respect I count myself as being very lucky, having always enjoyed the unstinting support of every sector of the JANET community. I would like to take this opportunity of putting on record my gratitude to all my colleagues and friends.

Appendix 1
Protocol Implementations available as at Summer 1986

	Protocol										
	X-25	TS	CR82	X-29		TS29		FTP	JTMP		MAIL
Machine range/ Operating System				PAD	HOST	PAD	HOST		RJE	HOST	
ICL/GEORGE 3	a	a		a	a			a			
ICL2900/VME	A	A			A			A	A	A	(f)
ICL2900/EMAS	(a)	(a)	(a)		(a)		(a)	(a)			(a)
ICL PERQ/PNX	a	a		a	a	a	a	a	d		a
GEC/OS4000	A	A		A	A	A	A	A	F	F	A
DEC10/TOPS10	a	a		a	a	a	a	a			a
DEC20/TOPS20	A	(a)			A			(a)			(d)
VAX/VMS	A	A	f(a)	A	A	A	A	A	A	A	A
VAX/UNIX	a	a	f	a	a	a	a	a	a		a
PDP11/UNIX	a	a		a	a	a	a	a	f		a
PDP11/RT11	a	a		a		a		a			
PDP11/RSX11M	A	a		A				a			
CYBER/NOS	A	F			A			F	D	D	(d)
HONEYWELL/MULTICS	A	A		A	A			A	A	D	A
HONEYWELL/GCOS	A(a)	a			A(a)			a	a		a
PRIME/PRIMOS	A	A		A	A	A	A	A	A	D	A
HARRIS/VOS	A	D		A	A			(d)	(d)	(d)	
IBM/VM/CMS	Aa	Aa	a	Aa	Aa	Aa	Aa	Aa	A	A	Aa
IBM/MVS	A	A(a)			A		D	(a)		(a)	
NORD/SINTRAN	A	A			A			A	A		
ITL/DIAMOND 7			A					A			
GEAC	A	F			A			F			
DG/MV ECLIPSE	A	D		D	A	D	D	D	D	D	
GOULD/MPX	A	D		D	D	D	D	D			D
GOULD/UNIX	A	D		D	D	D	D	D			D

The column headings are:

- X-25 PSS-compatible X-25
- TS Yellow Book Transport Service over X-25
- CR 82 YBTS over Cambridge Ring as defined in CR 82
- X-29 Green Book terminal support into an X-25 call
  - PAD Function as a Packet Assembler Disassembler into X-25
  - HOST Accept incoming X-29 terminal calls
- TS-29 Green Book terminal support over YBTS
  - PAD Function as a Packet Assembler Disassembler into YBTS
  - HOST Accept incoming TS-29 terminal calls
- FTP Blue Book File Transfer Protocol
- JTMP Red Book Job Transfer and Manipulation Protocol
  - RJE Remote Job Entry via Red Book
  - HOST Host Job Execution via Red Book
- MAIL Grey Book Mail Protocol

The product status is indicated by:

- A Manufacturer supported product available
- a JNT sponsored product available to community sites
- F Manufacturer product in field trial
- f JNT sponsored product in field trial
- D Manufacturer product under development
- d JNT sponsored product under development

If a product is shown as (a), (d) or (f) then it is available on a 'goodwill' basis from elsewhere in the community. Where several products exist they are separately listed.