Ann Robertson

INDRA
Working
Paper IEN - 190

INDRA Note 1114
IEN 190
9th. July 1981

Routing and Access Control in UK to US Services

Robert Cole and Robert Braden

ABSTRACT: The routing and access control problems for US-UK Catenet operations are discussed and an interim solution proposed, based on addressing mechanisms. Given the present generation of internet gateways, it is necessary to curtail the use of the full physical connectivity between the US and the UK to avoid gateway routing problems and undesirable paths.

Department of Computer Science
University College, London

C O N T E N T S

1. Introduction

The removal of the TIP at UCL, and the consequent change to TCP-based service from UCL to the DARPA Catenet, present some interesting problems in routing and access control. These problems are discussed here, and some interim (and ad hoc) solutions are presented.

The changes in connectivity at UCL, and at RSRE, mean that the extension of the DARPA Catenet into the UK must be considered as a whole; hence this document also looks at the position of the RSRE network (PPSN) in the routing schemes.

2. Background

When UCL moves over to an entirely TCP-based access to the Catenet, all of our traffic will be routed via a number of gateways. As most of the traffic will be destined for the ARPANET at least two gateways will be involved. UK users of hosts on the ARPANET can be divided into two groups:

1. Authorised MoD users -- this includes RSRE and UCL experimental traffic.

2. All other users

The MoD authorised users are allowed to use SATNET as a path for their traffic into the ARPANET. The others must use a PTT-supplied path to cross the Atlantic and enter the ARPANET through a special gateway.

The configurations and services UCL are putting forward for all users is presented in [1]. Essentially, MoD users may gain access to ARPANET via:

a. PPSN

b. TAC (dial-in terminals)

c. PAD to PSS and UCL

d. FTP via PSS and UCL

Other users will access ARPANET via:

a. PAD to UCL (via PSS or SRCNET)

b. FTP to UCL (via PSS or SRCNET)

From UCL, access control is applied for each user and the appropriate path is selected.

This note describes how the paths are enforced, both for forward and return traffic, and how a number of routing problems from dual connectivity are avoided.

3. Network Connectivity

The physical connections available between UK systems and the US are shown in figure 1. Note:

i. Gateways are indicated by the letter "G".

ii. That a single gateway has been assumed at PPSN for simplicity.

iii. A single gateway (G') has been shown at UCL. When line 77 is removed, and until the C/70 gateway is installed, this gateway will be implemented by two smaller gateways having only 2 and 3 ports. The exact configuration is yet to be decided.

The "PTT network path" includes the concatenated VAN/PTT networks TELENET and TYMNET in the US, IPSS across the Atlantic, and PSS and SRCNET in the UK. These networks all use X.25, and are joined by X.75 gateways which are not shown here. The gateway which is shown between the ARPANET and the VAN network TELENET is being built by BBN and is therefore know as the "BBN VAN gateway"; it will be sited in the US.

Across the PTT network path, internet datagrams are encapsulated within X.25 packets, over virtual calls which are opened as required by the BBN VAN gateway or by UCL. We use the term "tunnel" for such a path using X.25 encapsulation of IP datagrams. A second tunnel is shown, linking UCL with PPSN within the UK, using the British Telecom network PSS.

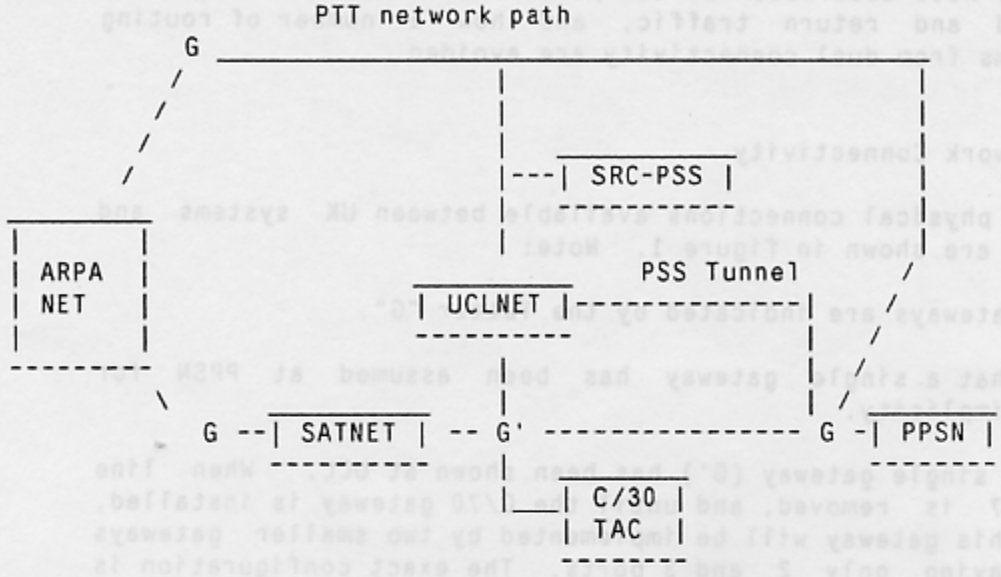

Figure 1. Physical Connections For UK-US Access

The physical connectivity shown in figure 1 presents several serious problems for internet routing algorithms. The present generation of IP gateways [2] assumes all paths between gateways are equivalent in delay, cost, and administrative authorisation. The various paths shown in figure 1 include violations of all three of these conditions. For example, routing MoD or ARPANET traffic over the PTT path, hence over IPSS, between the US and UK will create unacceptable costs.

As an interim solution, we propose to effectively reduce the connectivity, to that shown in Figure 2 [1]. The upper (PTT) path provides ARPANET access for non-MoD users, while the lower (SATNET) path is for MoD use.

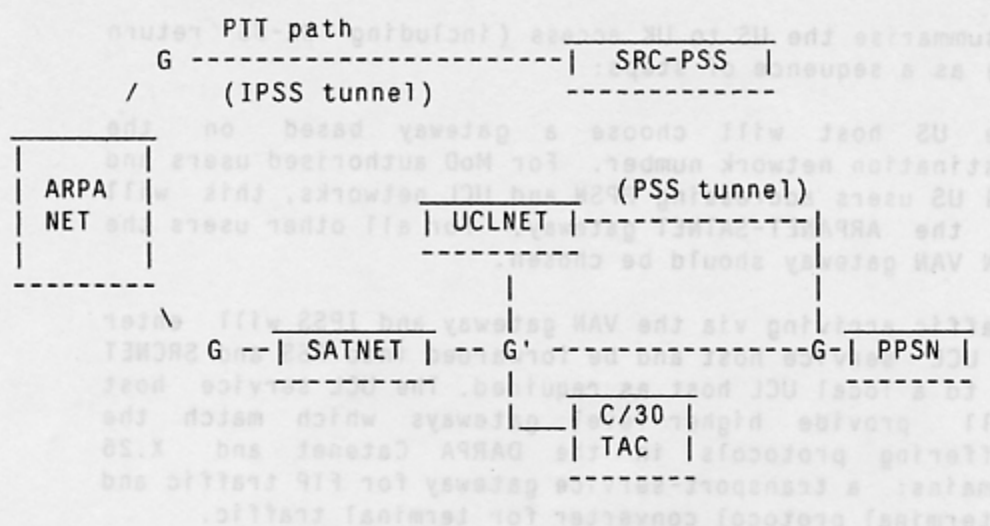

Figure 2. Apparent Connectivity For Service

4. How it Will Work

The reduction in connectivity between the available physical connections and the apparent connections is achieved by using addresses to convince the ARPANET gateways they are connected to logically disjoint networks. We also use addressing to ensure that US hosts use the correct path for return traffic. This is achieved by using different network numbers on each path, enabling the US hosts to select different gateways.

Thus, we will represent all non-MoD users as coming from a network we will call SRC-PSS, and all other users as coming from UCLNET. Of course traffic from PPSN is not affected. In this way the US hosts will select the ARPANET-SATNET gateway (and thus the SATNET path) for MoD authorised users and the BBN VAN gateway for all others.

We can summarise the paths from the UK to the US as:

1. From PPSN:
   this traffic has a direct link to the SATNET gateway, and only allows authorised traffic

2. From PSS:
   this traffic must pass through access control at UCL and thus the correct path is chosen

3. From TAC:
   this traffic has direct access to SATNET, but only authorised users will know the telephone number. Eventually TIP (TAC) login will increase the security.

We can summarise the US to UK access (including UK-US return traffic) as a sequence of steps:

1. The US host will choose a gateway based on the destination network number. For MoD authorised users and all US users addressing PPSN and UCL networks, this will be the ARPANET-SATNET gateway. For all other users the BBN VAN gateway should be chosen.

2. Traffic arriving via the VAN gateway and IPSS will enter a UCL service host and be forwarded into PSS and SRCNET or to a local UCL host as required. The UCL service host will provide higher-level gateways which match the differing protocols in the DARPA Catenet and X.25 domains: a transport-service gateway for FTP traffic and a terminal protocol converter for terminal traffic.

3. Traffic arriving from the UCL-SATNET gateway will pass to either the TAC, PPSN gateway, or into the UCL network via the SATNET Access Machine (SAM) [3].

5. Access Control for the PTT Network Connections

Extensive use is made of the PTT networks for carrying user traffic and encapsulated datagrams. Controls are required to ensure that only authorised users are allowed to connect to the various entry points of the Catenet.

There are two areas to be considered: user level access, and gateway access for encapsulated IP datagrams (tunnels).

1. All user level access to the Catenet from PSS and SRCNET is controlled at UCL by a login procedure, both for terminal access and file transfer. This is discussed in [1]. Similarly, any PSS access to PPSN will be vetted by the MoD's VAN gateway.

2. A more serious problem is the misuse of the tunnels across the PTT networks, for two reasons. First, one or both ends of a tunnel may need to treat the opposite end as a "trusted host" with respect to application of access controls. Since the two ends of the tunnel are connected with switched rather than permanent circuits, one must guard against a third host masquerading as one of the legitimate tunnel hosts. Secondly, the virtual X.25 calls used to carry IP traffic will incur usage charges. Across the international link IPSS, in particular, these charges will be substantial.

Figure 1 shows the two tunnels that are planned.

1. From UCL to BBN VAN gateway via IPSS. (Prior to providing service, we will start accessing IPSS via PSS.)
2. From UCL to PPSN via PSS

The "trusted host" problem can be handled easily, by having each end of the tunnel accept calls only from the expected host at the other end. Since the PTT network provides the remote host address, this tunnel can only be misused by "breaking" the PTT network.

For example, the UCL end of the "PSS tunnel" will accept X.25 virtual calls only from the PPSN gateway, while the MoD VAN gateway to PPSN will accept X.25 virtual calls only from specified calling addresses. Similarly, UCL will only accept calls from the BBN VAN gateway, and the latter will have a list of acceptable calling addresses [4].

Usage charges can pose much greater difficulty.

i. Accounting

The IP tunnels and the VAN gateway [3] will operate only on IP datagrams, and therefore can account for usage only on the basis of addressing that is visible at that level of protocol. Thus, it is possible to account by internet host, but not by virtual call or even individual user. This forces any user-level accounting and access control back onto the hosts at each end of the virtual call, and at present no ARPANET hosts are prepared to accept this responsibility.

Haverty [3] has pointed out a further problem with host-level accounting: the internet gateways do not ensure correct source addresses in the internet datagrams they process. Thus, the VAN gateway cannot rely upon even the alleged internet source host of a datagram destined for the tunnel.

ii. GGP

It is very undesirable to have internet gateways on both ends of a tunnel through a PTT network, since GGP messages interchanged by the gateways will generate excessive usage charges. Thus, in Figure 2

the UCL end of the IP tunnel must be an internet host, not a gateway.

The PSS tunnel for PPSN cannot avoid this GGP traffic, but this path is expected to be used only for experiments in alternate routing, and will not normally exist.

iii. Retransmission

There are serious problems with the use of an end-to-end retransmission protocol like TCP across an IP tunnel. Some host implementations of TCP use retransmission timeout computations that are incapable of adapting to long network delays, causing execessive retransmissions when calling the UK. The result will be to create excessive usage charges which are not under user control.

iv. Call Multiplexing

To minimise usage costs, it is necessary to use the X.25 virtual call as efficiently as possible. In general, there should be at most one virtual call open for a given tunnel. When traffic ceases, this call should be closed after a suitable interval. The minimum call duration charge interval should be considered when choosing this idle-call timeout interval.

If both ends of the tunnel open calls to each other simultaneously, two calls will result, and one must be closed. There can be an agreement between the two ends of the tunnel that one of them will always close a duplicate call. In the absence of such an agreement, a symmetric algorithm must be used to resolve the race. For example, each side can accept the incoming call and process traffic from it, wait for a random interval, and then close the incoming call (unless the other side has meanwhile closed the other outgoing call). In the event that both sides wait the same time and therefore close their calls simultaneously, the algorithm should be repeated.

6. Reconfiguring for Failure

The weak links in the routing seen by the Catenet gateways are the Atlantic connections. When one of these breaks, whole sections of the UK user population will be disconnected.

In fact, the physical connectivity allows us to realign all UK-US traffic onto the single remaining path, and in theory this change is fairly easily made. In practice, the cost of using the PTT path and the politics of using the SATNET path may preclude such operations. However, we wish to understand the technical problems in switching from the normal dual path to either single path.

1. SATNET failure
   The MoD traffic via UCL may be routed via IPSS quite easily. PPSN may open their own virtual call to the BBN VAN gateway, but they would need to indicate themselves as neighbours to obtain traffic from the US side (this requires another null network). If a gateway at UCL made itself known as a neighbour to the BBN VAN gateway, a similar path would exist. In either case, traffic from the TAC would be catered for. See figure 3.

   Notice that the "PTT network path" in Figure 3 really constitutes two or three different internet networks -- SRC-PSS, UCLNET, and perhaps PPSN. For this to work, we require BBN's VAN gateway to support a number of different "local" networks on the VAN side. We believe this to be feasible within the general framework described in [4].

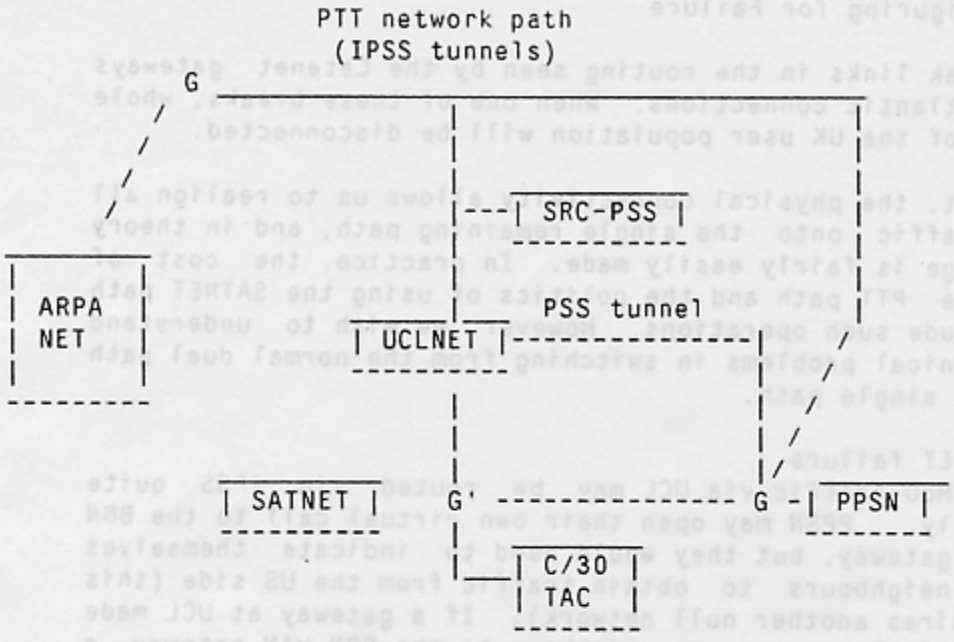

Figure 3. Alternative paths for SATNET route failure

2. IPSS failure.
The non-MoD traffic could be switched by making it appear to come from UCLNET, however all existing connections would be lost as the addresses change. To continue using the PSS-SRC network addresses via SATNET would require UCL to impose a gateway and inform the UCL-SATNET gateway of the new neighbour. The dynamic rearrangement of Atlantic links requires the gateways to issue redirect message and the US hosts to act on them. See figure 4.

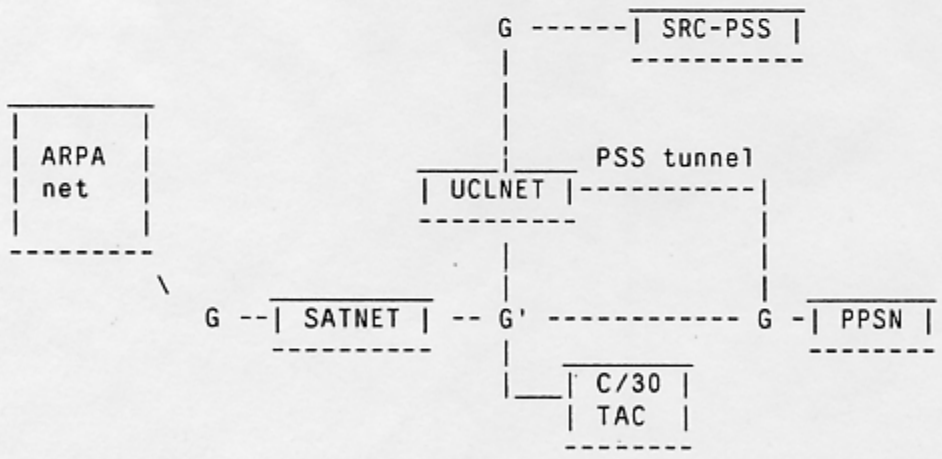

Figure 4. Alternative paths for IPSS route failure

7. Issues in Internetworking

The problems, and solutions, discussed above represent specific cases of more general issues in Catenet design and operation. In [5] [6] [7] Rosen presents a number of problems in the current Catenet design and implementation. He then goes on to propose a new model for catenet operation. This section will look at the UCL problems in this context, as a contribution to the debate.

The addressing solution used at UCL to force a particular route on a packet has 3 disadvantages.

1. It reduces the potentially rich connectivity of the available physical connections to a minimum.

2. It manipulates Catenet routing in an unacceptable way (i.e. a hack).

3. It requires complicated manoevres to restore service via alternative paths when the minimum connectivity is further reduced by failures. It is not clear how easy it will be to automate any switchover.

Ideally, the Catenet routing algorithms, implemented in the gateways, should be capable of knowing the full UK-US connectivity without the UK being used as an expressway for US-US traffic. In the present Catenet design and with the present Catenet protocols, such full knowledge by the gateways would have to be constrained by specifying source and return routes in each IP datagram.

In practice we find that dependence upon source routing is impractical, as it requires every gateway, and every US host that any UK user may access, to support these 'options'. The overhead of this option on IPSS/PSS charges is another concern. Even more seriously, source routing would move the burden of reconfiguration back onto the hosts and, in the end, the users.

In the model put forward by Rosen, gateway routing may be constrained by factors such as cost and suitability of paths. In the UCL case we should also add political factors arising from crossing national boundaries and Telecommunications monopolies. In the Rosen model the full UK-US connectivity would be available to the Switches. Thus we must look at some mechanisms by which we can constrain the actual routes available to particular classes of traffic.

The obvious mechanism is to use a 'class field' in the packet which identifies a potential routing problem to the

switches. However the UK is not alone in its problems, and in a general Catenet we can imagine a large number of paths having a number of user classes giving a very large class-problem space. Of course, if the entire Catenet is to be accessible from everywhere, each class/path combination must have a unique representation to enable the source Switch to plan a route.

The situation is further complicated by allowing a gradation of class. It is not unreasonable to suppose that someone may be prepared to pay for a path (such as IPSS) if the first choice (say SATNET) is unable to provide the required level of service due to traffic loading or malfunction. Thus we also need a mechanism that can say 'if the delay on path X goes above d then use alternative path Y'. Similarly we may have a situation where an administration will say 'allow any users of class m on this path as long as the loading is less than n%'. All of this complexity for each packet! We admit this is an argument about the distant future; however we should consider these general problems now because we already have specific instances of them.

As a final note, it is worth pointing out that the UCL addressing hack works because all UK-US traffic to the DARPA Catenet comes through UCL. This situation may not continue. It is not infeasible that a UK research establishment or a US Defense installation in the UK (or Europe) may obtain independent access to the Catenet. Or, the German SATNET installation might decide to increase their service reliability by adding extra connectivity. Each of these may use existing paths as alternatives. In these cases our addressing solution may fall apart, with disastrous results.

8. Conclusion

The routing and access control problems for US-UK Catenet operations have been discussed and an interim solution proposed, based on addressing mechanisms. Given the present generation of internet gateways, it is necessary to curtail the use of the full physical connectivity between the US and the UK to avoid gateway routing problems and undesirable paths.

Access control between the PTT networks and the DARPA Catenet has also been covered in some detail.

References

1. Higginson, P. and Braden, R., UK-US Services,
Indra Note 1101, IEN 185

2. Strazisar, G, How to Build a Gateway,
IEN 109

3. Cole, R. and Lloyd, P., The SATNET Access Machine,
Indra Note 1113

4. Haverty, J.,
Van Gateway: Some Routing and Performance Issues,
IEN 181

5. Rosen, E,
Issues in Internetting Part 1: Modelling the Internet,
IEN 184

6. Rosen, E,
Issues in Internetting Part 2: Accessing the Internet,
IEN 187

7. Rosen, E, Issues in Internetting Part 3: Addressing,
IEN 188