

Backup Access to the European Side of SATNET

Robert Braden

DISCUSSION

The purpose of this RFC is to focus discussion on a particular Internet problem: a backup path for software maintenance of the European sector of the Internet, for use when SATNET is partitioned. We propose a mechanism, based upon the Source Routing option of IP, to reach European Internet sites via the VAN Gateway and UCL.

This proposal is not intended as a standard at this time.

1. Introduction

During several previous SATNET meetings, it has been observed that it would be useful for BBN to be able to access the European side of SATNET indirectly via the VAN Gateway, when direct SATNET connectivity has been lost. This short paper proposes a possible approach to such "backup" access, using the source routing option of IP.

Figure 1 illustrates the problem we wish to solve. The US host H is used for diagnosis and control of the SATNET SIMP's S1 and S2 as well as the gateways B and G and the UCL TAC (not shown, but connected to G).

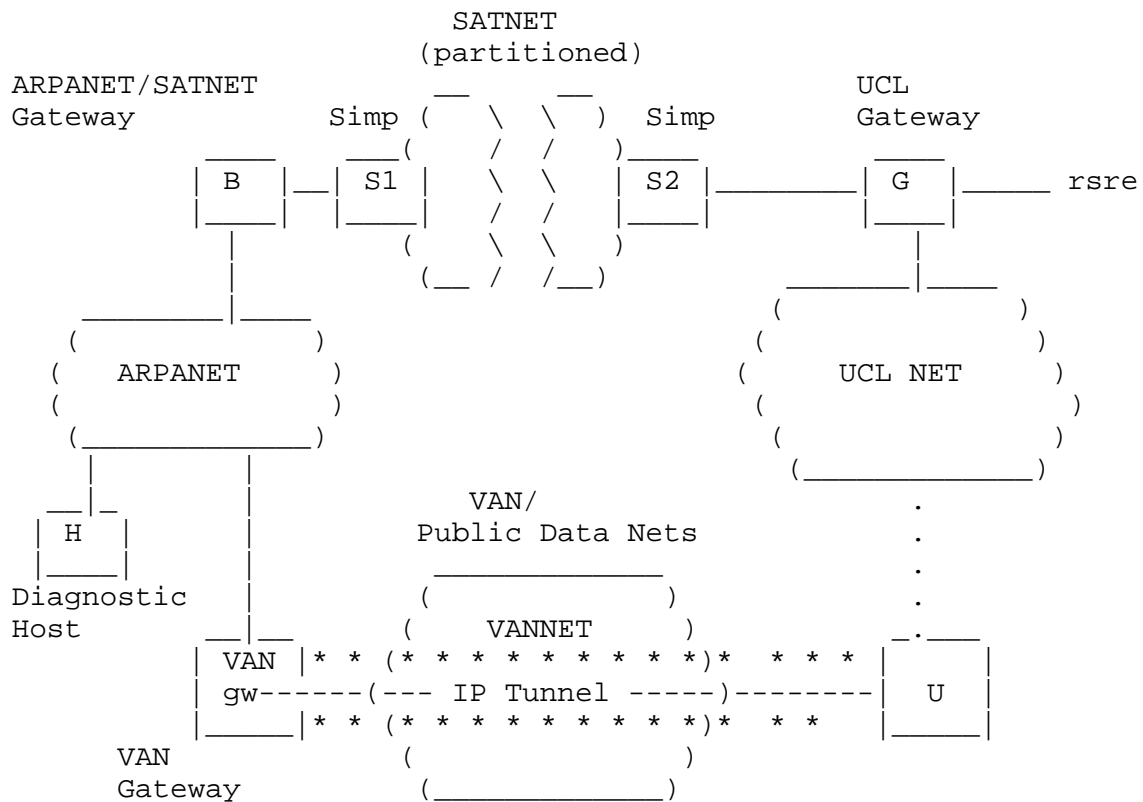


Figure 1. US/UK Connectivity with Partitioned SATNET

VANgw is the VAN Gateway which encapsulates IP datagrams in X25 packets for transmission over VAN/PTT virtual circuits. The collection of these paths, called "IP tunnels" by UCL, is addressed from the Internet as a distinct network, VANNET.

U is a UCL host, the Terminal Protocol Converter, which provides a path to UK X25 networks. However, to the Internet world U looks like a host on VANNET, so the path from U to UCLNET (shown dotted) does not appear to exist.

Now suppose SATNET is partitioned between S1 and S2. Then we wish host H to be able to exchange IP datagrams with S2 via the "back door" path:

H - Internet - VANgw - VANNET - U - UCLNET - G - S2

There are some important rules in this game, however.

- (1) U may only be a host, not a gateway.

This is because we do not want the Internet to route ALL its traffic (e.g. rsre traffic and UCL traffic that is required to use SATNET) via the IP Tunnel. So the VAN Gateway (VANgw) must not discover it can get to UCLNET through U.

- (2) To implement the back door path to S2, we are willing to have some special code in H and/or in U, but not in G, S2, or VANgw.

Note: Jack Haverty is allowed to violate this assumption, though we doubt that he will want to. But we must stick to it.

Given these constraints, we claim that the only possible solution is to "mung" the headers of IP datagrams at UCL. Thus, when SATNET is partitioned:

- (1) The IP addresses of S2, G, and the UCL TAC are unreachable from all US gateways. Therefore, if H sends a packet addressed to one of these destinations, it will be discarded and an ICMP unreachable message returned.

- (2) Similarly, the IP address of H is unreachable from the UK side. Hence, if the XNET debugger in a UK host emits a return packet addressed to H, that packet will be dropped.

Therefore, the destination address of each packet from H must be changed in order to reach the UCL side of SATNET (S2 or G), and the source address of each of these packets must be changed so that return packets can reach H. For this purpose, we introduce the Munger host M (see Figure 2).

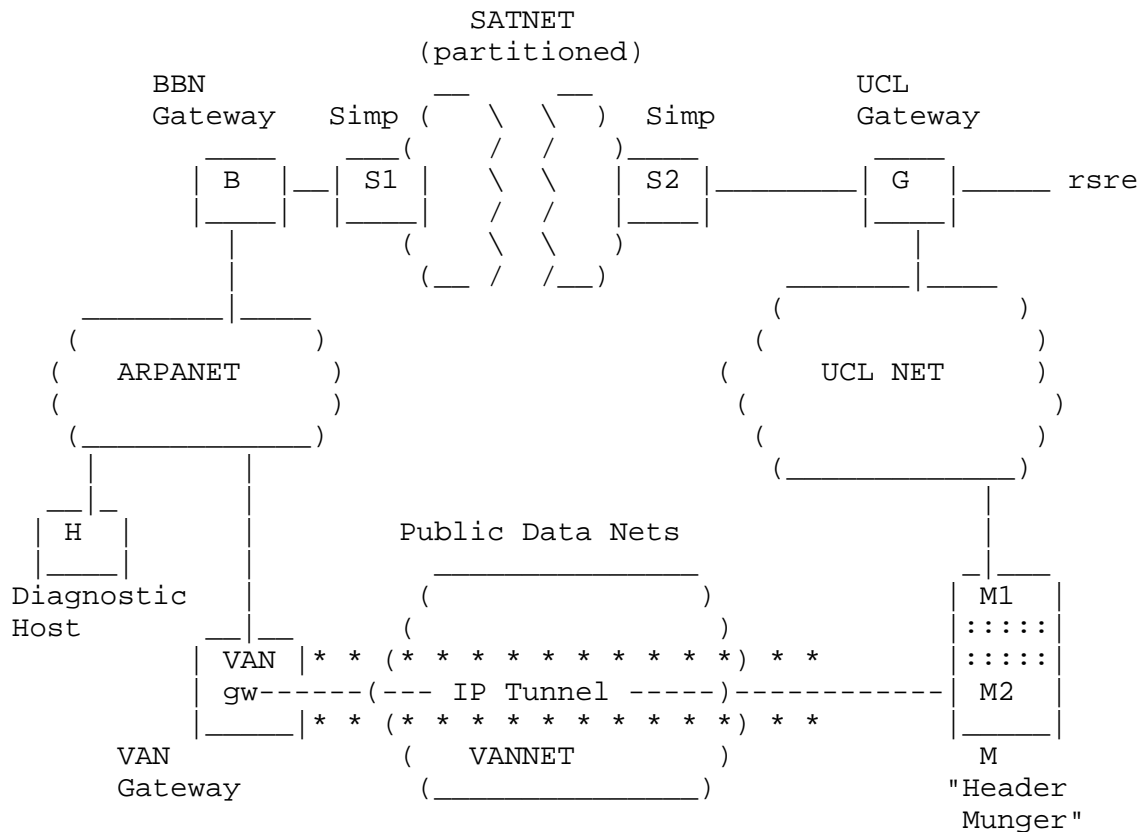


Figure 2. Introduction of Header Munger at UCL

Host "M" (M1/M2) is multi-homed, appearing as host M2 on VANNET and as host M1 on UCLNET. Like host U (shown in Figure 1), host M2 is the end of an IP Tunnel which communicates with VANgw over an X25 virtual call.

Suppose for example that host H desires to reach the XNET

debugger in the SIMP S2. H must send its packets with
destination address M1; these will be routed to M1 via VANGw
and the IP Tunnel. Host M will change the headers of these
datagrams to contain source address M1 and destination S2.
S2 will return packets to M1, and M1 will change them back
to M2->H packets and launch them back through the VANNET to
H.

How does M know how to change the headers?

- (1) M could respond to a range of M1 and M2 addresses,
and have a fixed table of correspondence.
- (2) We propose instead to use the SOURCE ROUTING option
in the datagrams. This assumes that H is able to
build source-routed datagrams, and is not upset that
the intermediate host in the route is not a gateway.

If we further assume that the IP layers in G and S2
can handle source and return routes, then the task
is simple. M must contain the source routing
algorithm of a gateway, but otherwise act as two
hosts (no routing updates, etc).

- (3) Although G supports source routing, S2 and the TAC
may not. In that case, S2 and the TAC will not be
able to recognise the return route in a received
packet and use it as a source route in packets sent
in reply.

This possibility calls for additional complexity in
M, a combination of (1) and (2):

- * In the US -> UK direction, the Source
Routing option would be used.
- * In the reverse direction (UK -> US),
mapping of datagram addresses would be
controlled by a table in M.

We suggest that M use source routing to get packets from H to S2, and meanwhile build a "soft state" table showing this mapping. When a packet comes from S2 without source routing, M would consult this soft state table to discover how to alter the addresses to reach H again. This would allow only one US host at a time to access a given SATNET host, but surely this is no restriction.

In practice, M2 and U should have different IP tunnels and hence different DTE addresses. Since the caller pays the X25 charges, the IP Tunnel for U will normally be opened only by UCL. On the other hand, the IP Tunnel to M2 will be opened from the US end. Since UCL has only one PSS line, this requires the use of separate X25 subaddresses. The VAN gateway must handle 14 digit X121 addresses, as well as 12 digit addresses.

2. Acknowledgment

Robert Cole of UCL has made major contributions to the contents of this paper. In particular, he suggested the use of the Source Routing option.