

A Comment on Current Source Routing Techniques

Michael Witt
UNISYS Defense Systems
5151 Camino Ruiz
Camarillo, CA 93011

Introduction

The purpose of this note is to air the author's belief that there is a certain deficiency in the way the "source route" option is defined in both the ISO and DoD Internet Protocols. The two protocols handle source routing in essentially the same manner, and do not provide the full capability that one might like to have from this option. It is proposed that a very slight addition would rectify this situation in both ISO and DoD IP.

Definition of Source Routing

Source routing is a method whereby the host computer that originates an IP datagram is able to specify all or part of the path which that datagram will take through the Internet. This is in contrast to the more typical situation where an originating host simply "injects" the datagram into the Internet and assumes that the intermediate systems (gateways) will route it on to its final destination.

The Usefulness of Source Routing

There are probably many situations in which a source routing capability would be useful. Some possible uses are:

- (1) Allowing the use of connectivity information which is available to the originating host but not to the intermediate systems. In other words, the sending host knows how to route the IP datagram to its destination, but the gateways do not. An example of this is in "down-line loading" a gateway with its initial operating software. Since the gateway is not yet active, it has not broadcast news of its existence to its neighbors. Hence, the host responsible for initializing the gateway must specify the route explicitly.
- (2) *Restricting* the route that a datagram may take. For instance, it may be necessary to explicitly route control information around some "trouble spot" in the internet. In some cases a malfunctioning gateway could become a sink for all traffic in the area unless there is some way of forcing datagrams to take a different route. (An analogous problem was once seen the Arpanet subnet when a particular IMP malfunctioned and decided that it had zero delay to every other IMP in the network.)
- (3) Causing a datagram to be routed to a specific set of intermediate systems. One might wish to propagate control information to a group of gateways by sending a single datagram which would be routed from one to the other. This could be much more efficient than sending a separate control message to each individual gateway in some cases.

How Source Routing can be Accomplished

In looking at what techniques would be appropriate to implement a source routing capability we need to consider what the requirements actually are. The three cases that have been described above seem to provide a good starting point. Note that the requirements are somewhat different for each case.

In case (1) we basically need a way to give routing information to the gateways. This might come in the form of complete routes or simply in the form of "hints". If our main objective is to get the datagram to its intended destination, then we would like to be able to take advantage of any information known to the originating host as well as any information that the gateways possess. Therefore we want the gateways to be able to use the information that the host provides, but we do not want them to be constrained by this information if it should prove insufficient.

In case (2) we do want to constrain the route to the one specified by the originating host. We specifically do not want any gateway to do routing based on its own perception of the current Internet topology.

In case (3) the desired effect is to cause the datagram to visit more than one location, specified by the host. It would be possible to accomplish this by *completely* specifying the route, as in case (2). It would be more desirable, however, to allow the gateways to use any routing information available to them, as long as the route eventually passes through all of the specified destinations. That way, the originating host only has to know the addresses of the destinations it wishes to communicate to and does not have to construct the entire route.

Discussion of ISO and DoD Source Route Options

The ISO [1] and DoD [2] Internet Protocols handle source routing in essentially the same way. Both have two types of source route options, called "strict source routing" and "loose source routing".

In strict source routing, the originating host specifies a list of addresses. During each hop the datagram makes, it must be forwarded to the next address on the list. If, at any given hop, the gateway possessing the datagram cannot forward it directly to the next address on the list, it must be discarded.

Loose source routing is like strict source routing except that, during any given hop, the gateway may forward the datagram in the general direction of the next address on the list. This means that the originating host does not have to specify the route completely. The gateways will do their best to route the packet to all of the specified addresses. Still, the packet will only be directed toward the next address on the list. If the next destination is ever determined to be unreachable, the datagram will be discarded.

It would appear that strict source routing is ideal in case (2) above. If a host needed to ensure that a particular message would not deviate from the specified route, the strict source route option would accomplish this very nicely.

Loose source routing, on the other hand, might be useful in case three. The loose source route option could be used to cause a given message to be circulated among a specified set of intermediate nodes. This could be done even if the originating host only knew the addresses of the intended recipients. Loose source routing is still, however, a bit too "strict" to be ideal for this purpose. We would like the message to continue to circulate to the remaining addresses on the list even if one of them should happen to be unreachable, but the way the option is specified the datagram would have to be discarded in this case.

Neither strict nor loose source routing supports case (1) perfectly. If we use strict source routing it will be necessary for the originating host to know the entire route to the destination. Even if the host does possess this knowledge, it will not be possible for the gateways to make use of alternate routes. If we use loose source routing, then the host will not have to know the whole route. But, if the intended destination cannot be reached through a path that includes *all* of the specified intermediate systems, then the message will have to be discarded, even if an alternate route was available. The gateways are able to make use of the information provided by the host, but they are also, unfortunately, *constrained* by this information.

Suggestion

It is not clear that there is an application of the loose source route option that would require that a packet be discarded if, at any point along the way, it was not possible to reach the next specified intermediate system. It might be worth considering a re-interpretation of the loose source route option that eliminates this feature.

If it is desirable to maintain the current loose source routing functionality then perhaps it might be appropriate to consider a third type of source routing option that would allow an originating host to specify a partial route without forcing the gateways to discard the datagram if a known, alternate route is the only possible means of reaching the destination.

It might be argued that any host engaging in source routing could be expected to be able to specify the route well enough that it should never need to make use of alternate routing information known to the gateways. However, in an Internet with experimental *autonomous systems* [3] being attached, and with a growing concern over the possibility of interconnecting DoD and ISO networks, it would seem prudent to allow for the most flexible options possible.

Note: the above suggestion is made in order to encourage discussion in the research community. This is not a proposal being directed at the DoD or any standards group.

References

- [1] ISO DIS 8473 "Protocol for Providing the Connectionless-mode Network Service", March 1986
- [2] RFC 791 "Internet Protocol", September 1981
- [3] RFC 975 David Mills, "Autonomous Confederations", February 1986