IEN #86

EXTENDED INTERNET ROUTING

Radia Perlman
Bolt, Beranek, and Newman
April 5, 1979

INTRODUCTION

The catenet differs from most networks because in the catenet not all links are functionally identical. Because, for various reasons, various types of packets should not be routed on certain of the links, no single topology defines the connectivity state of the catenet for all types of traffic.

There are three factors to the design of catenet routing to meet these needs:
1) categorizing a packet
2) routing a categorized packet
3) preventing spoofing
These will be defined and discussed in this paper.

CATEGORIZING PACKETS

A packet's category is a number that is n bits long, where n is the number of networks plus the number of other factors that should be considered, such as the delay class or reliability desired. A "1" for a bit means the packet is allowed to traverse that network, or the factor (such as delay) corresponding to that bit is of importance.

In this paper we will assume that packets are routed according to:
1) delay
2) reliability
3) authorization.

Authorization can be based on such things as:
1) source net
2) destination net
3) source host
4) destination host
5) "stamp of approval".

The "stamp of approval" would be some sort of code given out by an access controller for a network that wishes to restrict traffic into or through itself (and does not wish to rely solely on other information in the internet header). A user who wishes to use one of these networks must contact the relevant access controller and receive the code. Either access controllers would be cooperative, and a single access controller could give the stamp of approval for several networks at once, or the user would have to contact an access controller for each fussy network he wished to use.

A gateway on a fussy network checks the authorization of a packet, and drops the packet if it is not authorized for that net. Other gateways do not check for authorization, but instead route based only on the n bit category number. If whoever placed the routing category in the packet header claimed access to a network incorrectly, the packet would be dropped at the gateway into or out of the net the packet should not traverse.

If requiring the source host to fill in the entire category is deemed feasible, that would give the system the most flexibility, because a user could decide for some experiment to avoid certain nets, even though the user would have access to those nets. Then, in general, gateways would not need to keep lists of who is allowed in various networks--only the gateways on the fussy network would need to compute whether the packet is, indeed, allowed on the network. When a gateway drops a packet, it should notify the source, in case the source was not aware that access to that net was restricted (for instance there could be a demo on some network during which traffic through it would be restricted, but usually access to the network was unrestricted).

ROUTING

Routing must be done on a per category basis. If the number of categories is small (16 or less), then ARPANET routing (in which nodes pass information to their neighbors about how far they are from all destinations, with the modification of reporting infinite distance to downstream neighbors) can be used. Gateways would report a distance vector for each routing category.

An alternate strategy is for gateways to report link state information instead, and have each gateway's link state report get transmitted to all the gateways on the catenet. In this way, each gateway would have full knowledge of connectivity information for the catenet. Then, each gateway would compute a separate shortest distance matrix for each routing category. Nets that correspond to "0"s in the category number would be links of infinite distance, whereas nets that correspond to "1"s, would have unit distance. To compute for delay, or reliability, each net would have some constant associated with its delay or reliability characteristics. Then instead of using "1" and infinity in the connectivity matrix, these constants would be used as a multiplier of the cost of traversing links. In this way gateways can weight against using low reliability networks for packets that need high reliability, but low reliability networks would be used if there were no reasonable alternative. This strategy would also apply to computing distance vectors in the ARPANET routing scheme.

If the number of categories is too large for gateways to pass around a separate distance vector for each category, or for gateways to compute and store a separate distance matrix for each

category, then the gateways must use a link state scheme, and they would have a cache of distance matrices for categories that are currently in use. If a packet with a new category comes in, the gateway would throw away the matrix associated with the category that has been least recently used, and compute the matrix associated with the new category.

The number of categories does not have to be 2^n, with n being the number of networks. It can be much smaller. If most networks do not care what kind of packets traverse themselves, then those bits would always be "1". Some networks might be willing to "team up", by having the same access control policy. Thus those bits would always either be both "0" or both "1".

SPOOF PROTECTION

Assuming the information in the internet header is correct, it would do a source no good to put an illegal category number in the packet. The result would be that the packet would be routed via a link that would drop it. (For example, a gateway on a network that did not want any packets from some particular other network would check the source net in the header and drop packets from that network. The gateway would not merely look at the category and believe that.)

The part of the header that can be forged to the source's advantage is the source host and network. Forging an incorrect destination would be meaningless--the packet will go to whatever destination is in the packet. To prevent a source from forging a different source address, gateways should drop packets unless either:

1) the internet source is on that network, and matches the local source, or
2) the internet source is on a different net, and the local source is a known gateway.

This requires that all gateways trust each other to do the test, and that for all nets, either it is impossible to forge a local source, or direct access to the net is limited to trusted users.

This is not a completely safe scheme, of course, since the above assumptions do not always hold. Thus networks that are extremely fussy must provide some access controller somewhere on the catenet that would issue permission to use the net. Presumably any packet addressed to any access controller would be allowed access through to the access controller, since some very important user might happen to be located, in an emergency, on a network usually denied access to most of the catenet.

The permission from the access controller would include some sort of key for encrypting the packet. Thus someone overhearing an allowed packet would not be able to copy the "stamp of approval"

and use it to gain access to a restricted network, since the stamp of approval would be a function of the data in the packet, and would change for each packet. Gateways into the fussy network would check the correctness of the stamp, and drop incorrect packets. (Gateways should not notify the source when they drop a packet due to a bad stamp of approval, since the user is presumably malicious. When a packet arrives with an illegal category number the user can not always be presumed malicious.)

IMPLEMENTATION RECOMMENDATIONS

With more than 16 potential routing categories, the current routing scheme used by gateways becomes unwieldy, since routing information about all potential routing categories must be exchanged by the gateways. Thus if the current routing scheme is to be preserved, the meaning of each of the 4 bits (16 categories) must be carefully chosen for maximum usefulness. For instance, one of the bits can be for all the military nets, and it can be assumed that packets are either allowed into all military nets or none. There are no difficulties with adapting the current routing scheme to accomodate a small number of different routing categories. Gateways will report to each of the neighbors, their distance vectors to all destination networks, for each routing category. To compute their own distance vector for a category, the routing algorithm is modified slightly to include a table, for each link, of whether that link is legal for that category. Then, to compute their distance to a destination for a certain category, they take the minimum sum (as before) of their distance to a neighbor plus that neighbor's distance to the destination, but only use neighbors that are legal for that category.

If more than 16 potential routing categories are desired, then a link state routing algorithm, as described in IEN 25, must be used. Gateways would exchange information about which links are physically up. Then an individual gateway would only compute a distance matrix for a given category when it receives a packet for that category. The gateway will keep track of the last time a matrix for a category was used, and will replace the least recently used matrix with a matrix for a new category when it needs the room. When a connectivity change occurs, the gateway will either erase all matrices and recompute them when needed, or recompute matrices for all categories in its cache at once. It is probably preferable to simply erase all matrices.

The only spoof protection that should be implemented initially is the check for forgery of the source address. Access controllers add a lot of overhead and should not be implented unless (and until) they are really necessary. The design of an access controller for the catenet, and the interaction with the catenet user, is beyond the scope of this paper.