J. Davidson
BBN

IEN 60
Section 1.4.5

17 October 1978

BOSTON AREA MEETING OF THE INTERNET WORKING GROUP TO DISCUSS INTERACTIONS WITH GATEWAYS

A small working group meeting was held on 17 October 1978 at BBN in order to firm up the proposed agenda for a large working group meeting to be held at the SRI internet meeting at the end of October. A list of attendees is included at the end of this report. The topic we met to consider was labelled "interactions with gateways". These notes were dictated at the completion of that meeting.

The following list identifies some of the components of the CATENET which we think must interact with the gateway at one time or another if the gateway is to be a completely functioning tool in the CATENET:

OTHER GATEWAYS
ACCESS CONTROLLERS
GMCCS
STATISTICS GATHERERS
DEBUGGERS
HOSTS (EXPERIMENT)
HOSTS (STREAM)

It is presumed that in addition to the basic chores which a gateway performs, it may also want to interact with any or all of the CATENET components shown above. (In addition, of course, it may want to interact with its local network or with some hosts on its local network in performance of local net activity. We are not concerned with these at this time.)

The title of this working group session is of interest. It is obvious that each of the components shown in the picture may have a reason for interacting with the gateway. However, it is also possible that the components may have a reason for interacting with each other at various times. Thus, eventually other working groups will probably have to be convened to consider, for example, "interactions with GMCCs", "interactions with access controllers", "interactions with experimental hosts", etc. In fact, we got into certain of those types of interactions at this meeting and had to work hard to extricate ourselves from them. I think they rightly belong to another meeting.

An overview of the goals of this local area meeting was given by John Davidson. He identified some general issues which surround each interaction of one CATENET component with another. For the gateway case, these correspond to the following questions:

1. What hooks must be provided in the core module of the gateway in order to allow the inclusion of all the mechanisms required for interaction with external components?

2. What programs must be provided in the gateway to perform the interaction with these external components?

3. What data must be collected or must be collectable within the gateway? And, must this data be preprocessed by the gateway before being transmitted to a remote component?

4. What protocols should be used between the communicating interacting parties? How can the protocol be made extensible so that a minimal implementation need not constrain other implementations which want to perform more activities?

5. How is one interaction impacted by any other interaction. For example, how is routing impacted by access control or how is source routing impacted by access control, etc.?

An even more general issue is whether or not a single mechanism should be provided in each gateway that describes the kinds of interactions the gateway supports. For example, should a mechanism be included in every gateway by which an arbitrary net component can ask the question, "does the gateway know how to interact with access controllers?" "Does it know how to interact with debuggers?" "Does it know how to interact with.......?" The choice here is either to include the general mechanism or to require that kind of query mechanism to exist within each type of specific interaction. It is probably the case that providing the general mechanism will enable the gateway to provide less mechanism within each type of interaction if the type of interaction is truly not supported on the gateway. This question is still unresolved, however.

The topics that were considered at the local area meeting included:

1. Monitoring and control of gateways.

2. Statistics gathering.

3. Access Control.

4. Stream setup.

5. Gateway to gateway congestion control, flow control, and routing.

6. Debugging.

I. MONITORING AND CONTROL

Mike Brescia reviewed the issues which he has been working on recently. Among the questions he's been trying to resolve are:

1. What information should be reported? Obviously, information on traffic and on errors should be reported.

2. What should the format of the reports be?

3. What should the destination of the reports be? And, how can this destination be varied if needed?

4. What should the frequency of reports be?

5. What protocols should be used to exchange reports with the GMCC? (And, as discussed in the opening remarks, what measures should be taken to insure that this protocol is extensible?)

6. What testing should the gateway be able to support? A question on testing arose in that it was not obvious whether source routing would be enough of a mechanism to allow a GMCC to determine the health of each of the components in the CATENET. There was some desire to put an echo program in the gateway and some very strong interest in having an echo process in various serving hosts throughout the CATENET, so that data could actually be sent through a gateway, rather than just to a gateway, and then be returned.

Mike promised to have a handout showing suggested reporting messages and formats ready for the SRI meeting.

We determined that as far as protocol was concerned it was okay to use an "unreliable" protocol, and, despite the fact that the gateway will implement an end-to-end, "reliable" protocol for exchanging routing and flow control messages with other gateways, it is not necessary to use that protocol for the reporting mechanism.

There was some discussion of whether the errors noted by a gateway could be used as an input to the routing algorithms employed by the various gateways. For example, if a GMCC sees a given gateway experiencing a lot of errors, can it cause the other gateways to route around it? This appears to be a reasonably complex thing to do, but is an example of the kinds of interactions which might occur between various components of the CATENET as they interact with other components. It was not clear whether normal gateway routing mechanisms would take care of these kinds of issues.

Most of the above issues were related to monitoring. As far as control is concerned, we have determined that a minimal gateway implementation does not want to provide much mechanism for being controlled by a GMCC. At the very least, however, it appears that the gateway should provide some module, which is addressable by the GMCC, to which advisories and requests can be sent. The gateway doesn't have to do anything with these advisories and requests, but it should at least have a program available at the address which can discard the messages so that they are (1) recognized and (2) not reported as error messages. Alternatively, it may be smart enough just to discard the messages directly if it has no control facility.

Other questions related to monitoring and control are those such as, "should the gateway be able to report other gateways being down? That is, if a gateway sees that its messages to another gateway, (e.g., its routing messages) are not getting through, should it provide this information to a GMCC which is trying to reconstruct the state of the CATENET?

One of the most interesting concerns was raised by Dave Clark who commented that the performance of one CATENET component in relation to the performance of the rest of the components may be difficult to assess. For example, he asked "How will I determine if other gateways think that my gateway is unresponsive or failed?" This is information which he probably cannot discover simply by metering his own component. What he would like to have, in general, is a fundamental capability to read the traffic statistics that have been gathered by a GMCC, and then to write programs, perhaps on his own service host, which can analyze the information to determine precisely the relationships he is interested in. He recommends that the GMCC make the raw data available to hosts in a fairly low-level manner, i.e., provide the data before they have been processed by any data reduction programs.

While this issue is truly a concern, it appears to be the topic of another working group meeting which is labelled "interactions with GMCCs".

II. STATISTICS GATHERING

At the local area meeting, we determined that statistics gathering was perhaps more a function of a particular local net than of the composite CATENET, in that the types of statistics that were gathered which were not sent to the GMCC under the terms of monitoring and control are really of interest primarily to the implementors and/or administrators of the gateway in question; they can, therefore, be considered individually for each individual gateway.

III. ACCESS CONTROL

This area was introduced by Radia Perlman and John Davidson. The types of issues which were addressed included:

1. What should the granularity of access control be? Should it be on a per-host, or a per-net basis?

2. Confidence levels. What degrees of confidence can we have in our access control mechanisms if the possibility for spoofing and/or masquerading within the CATENET exists?

3. Is bilateral accounting the only kind of accounting which will work (in view of the ability to spoof or to masquerade)? In bilateral accounting, a network only charges its adjacent networks for messages let into the net. They therefore don't have to authenticate the source net specified in the internet header.

4. Are half gateways required so that policies of not letting messages out of the net can be distinguished from policies of not letting a message in?

5. How can a network (A) let messages into a net (B) but specify that they not be allowed to go through net (B) to a further net (C).

6. How can a source indicate that it wants to avoid some net because of costs or security considerations.

7. What impact does access control have on routing and what impact on the addresses placed in an internet header? (This last question is of interest. The concern is that a gateway is known by two (or more) addresses potentially. If a message is not allowed to get into one of the networks to which the gateway is connected, what happens if the message is sent to that gateway using the other network's address?)

8. What are the implications of one-way nets? For example, what impact do they have on the return path to be used in source

routing? It was noted that one-way nets might exist if network (A) is allowed to send through (B) to (C), but (C) is not allowed to send through (B) to (A). Also, the notion of one-way nets is applicable when one of the networks must operate in "quiet mode".

Several mechanisms and techniques were discussed which might be of use in providing certain of the features of access control which we see necessary.

1. We wondered if Link or end-to-end encryption techniques needed to be used to hide the information transiting the CATENET or to provide reliable identification of correspondents within the CATENET.

2. Will the gateways have to remember information from message to message in order to eliminate certain of the access control interactions from occurring on a per-message basis? This may mean that virtual calls or virtual circuits need to be set up.

3. Can public key cryptosystems be of any use? For example, could they be used to help authenticate the source and destination host specifications? Each gateway could "decrypt", with a public key, the enciphered source address, to verify that charges and access controls will be applied correctly.

4. Should the source and destination gateways try to use some end-to-end techniques between themselves?

5. Do we need to have identical access controllers?

6. Do we need to have identical gateways (and do these gateways need to be proven correct)?

7. Do we need to have reliable, redundant access control facilities at each point throughout the CATENET to account for possible failure of a single access controller?

Radia has produced a report on the way in which access control considerations affect routing, and will distribute it at the SRI meeting. It is currently on-line at ISI as IEN 58.

IV. STREAM SETUP

This area was introduced by Richard Binder. Among the issues to be resolved are:

1. What level should stream setup be performed at? For

example, should the setup be performed by a user protocol, by a TCP protocol, by a level of protocols between TCP and internet, or by the internet protocol itself? Should a gateway be told by a "Type of Service" indicator that a stream is desired? Note the different kinds of mechanisms required in the gateway depending on the decision made here.

2. What will be the impact on routing of the stream setup? For example, will alternate routing be allowed if the stream is set up using a given collection of gateways?

3. What are the ways in which the stream can be identified throughout the CATENET? For example, if a stream is set up by a gateway attached to SATNET, should the local identifier used by the gateway to identify that stream be passed back through all the intermediate gateways to the source host?

4. What does stream setup in a CATENET really mean? Higher priority? Lower delay? Are the gateways supposed to throw away other traffic in order to provide congestion-free links to adjacent gateways in the streams?

5. What amount of time and space will be used within the gateway in order to support this new requirement?

6. Is it possible to use a special server host to try and set up the stream?

We believe that this topic is perhaps a great deal more complex than it appears at first glance. It should certainly be the subject of (perhaps several) other working groups.

V. GATEWAY CONGESTION AND FLOW CONTROL

Ginny Strazisar provided a status report on the progress being made in the implementation of the routing algorithm presented at the London internet meeting. This algorithm provides the capability for routing packets on shortest paths through the CATENET and for routing packets around failed gateways and networks. Testing of this routing algorithm in progress and it should be operational in the ARPANET/PRNET gateways and in the ARPANET/SATNET gateways by the end of this year. Work has begun in designing and implementing flow control mechanisms, especially in developing a source quenching mechanism.

VI. DEBUGGING

It was felt that debugging is properly in the domain of the individual administrators or implementors of distinct gateways and therefore was not discussed in detail at this gathering.

VII. CONCLUSIONS

The working group was able to identify a number of issues in each of several areas as discussed above. Some concern was generated at the end of the meeting that we do not fully understand the goals of the overall internet project in that we question the ability of the CATENET to support such activities as access control, in the absence of certain fundamental mechanisms which are not currently being proposed for the internet. Perhaps, a general discussion of the future of our internet project, especially in relation to the development of public packet switching nets etc. is in order.

VIII. LIST OF ATTENDEES

Mike Brescia
Radia Perlman
Richard Binder
Jon Cole
Dave Clark (MIT)
Ginny Strazisar
J. Noel Chiappa (MIT)
John Davidson