

Fanny

ARPA Catenet Monitoring and Control

IEN 105

25 May 1979

David Flood Page

Bolt, Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02138

(617) 491-1850

1. INTRODUCTION.

This document describes the proposed design for the Monitoring and Control system for the ARPA catenet. By 'catenet' is meant the system of connected networks, plus the hardware and software devices (gateways) that connect them.

Section 2 gives the background to the design and some issues in catenet monitoring and section 3 gives an overview of the design. Processing details, data formats and an implementation plan will be described in a separate document.

ISSUES IN CATENET MONITORING AND CONTROL

This section provides the background for the design of the ARPA Catenet Monitoring and Control Center, and identifies the design decisions that must be taken.

A catenet monitoring and control system is intended to do for the catenet what the existing Network Control Centers do for the networks. Briefly, this is to provide fault detection and isolation capabilities, and to give a picture of the catenet operation so that the performance of its components can be monitored.

Catenet monitoring differs from individual network monitoring in that the Catenet will not be under the control of a single administrative entity but will be divided into separate areas. Each of these areas will have its own monitoring facility, and the details of operation of the facility will vary from area to area. An area may typically consist of a single network so that the monitoring facility will quite likely be part of the Network Control Center (NCC) for that network. So that any one of the areas can obtain an overall picture of the Catenet, the monitoring facilities in each area should be able to exchange information with each other, though this may not always be possible.

The access granted to other Catenet Monitoring and Control Centers (CMCCs) by a given CMCC to the Catenet components in its area will obviously affect the ability of these other CMCCs to obtain a complete picture of the Catenet operation. It may well be that the only component that can be directly accessed from outside an area is the CMCC itself. Alternatively, a request to a CMCC could result in the requesting CMCC being given permission to access the relevant components directly.

The networks in the catenet are already monitored by their NCCs so that a CMCC is going to be concerned mainly with the gateways. These may already be monitored to some extent by the networks to which they are connected, but since existing NCCs do not know about internet traffic, a CMCC would be able to obtain a more complete picture. Depending on the facilities required of a CMCC, it could be limited to monitoring gateways only, or it could include data from every catenet component, or be somewhere in between. Obviously, the more information a CMCC receives, the better fault isolation it can do, especially if it has a map of the catenet instead of just a list of gateways. This kind of operation is more costly, however, and requires more communications bandwidth. One approach would be to have one higher level CMCC using inputs from several gateway-only centers, and perhaps from the associated NCCs as well.

Since the Catenet components will be implemented in different ways from area to area of the Catenet, the functions that they will provide for a CMCC may vary between components of the same type. A CMCC must therefore have some means of finding out what grade of service it can

expect from these components. While some functions will not be available because they are simply not implemented, others may be unavailable for more dynamic reasons such as authorization or resource availability. The mechanism will need to cope with this changing situation and also with the possibility that a component may not respond to this sort of enquiry.

The functions to be provided by a CMCC can be grouped into a few main areas. These are:

- Fault Monitoring.
- Performance Monitoring.
- Accounting Information.
- High level Fault Isolation.
- Fault Diagnosis Control.
- Maintenance Control.
- Experiment Control.

Fault monitoring - this is the information needed to keep the Catenet up and running. In the main it will be information about hardware and software failures in the gateways. It is assumed that problems internal to a network will be discovered, and dealt with, by the NCC concerned.

Performance monitoring - this will consist mainly of throughput statistics, suitably digested. This kind of information is used to observe congestion and routing problems in the Catenet, and also by system engineers to determine the effect of modifications to the system. For Catenet level monitoring, some routing information will also be included.

Accounting information - the existing networks usually do not know anything about Internet traffic. Putting appropriate processing in the gateways and CMCCs is the most economical way of finding out about this traffic as it enters or leaves a network.

High Level fault isolation - We envisage that some time in the future, automatic fault isolation algorithms will be developed which will be able to use Catenet level information, and detect subtle problems such as impending bottlenecks as well as the more obvious hardware and software failures. These algorithms are not currently well defined but it may well be worth considering what kind of database they might need, and to include suitable data collection facilities in a CMCC. The data base so constructed could of course be used by people for the same purpose, given appropriate software tools.

Fault diagnosis control - This and the other control functions are the most likely to vary widely in their degree of implementation. They are all functions initiated by someone operating the CMCC which will cause the gateways to alter their behaviour. Fault diagnosis operations will be used to try and locate a fault and will include things like causing a gateway to loop one of its network interfaces.

Maintenance control - this consists of more routine operations such as reloading of software. Obviously this, like the fault diagnosis control, will be very much up to those controlling a particular area and will probably not be useful or even available to those outside it.

Experiment control - would be used by system engineers to test out a new facility or a modification to an existing one. It may well be that this is not considered a function of the CMCC proper and may be implemented in a separate facility. However if that is done then it would be desirable for the CMCC to know that the experiments are going on.

In order to describe the level of service provided by a Catenet component we must define a list of functions which might be implemented or available. These are:

For a CMCC:

- Give Catenet area status (ok, slow, network isolated)
 - as a regular report
 - on request
- Provide detailed gateway information by simulating access
- Allow direct access to gateways

For gateways:

- Provide throughput statistics
- Provide fault event reports
- Provide routing information
 - the entire routing table on request
 - routing updates as they are created
- User accounting statistics
- Report on status of network connections
- Generate fake traffic
- Change function authorization
- Load new software
- Fault isolation functions, e.g. trace, timestamp,
loop network interface

The grade of service provided in any area is obviously going to depend on political as well as technical considerations. Since this whole thing requires a certain amount of agreement to work at all, we can also hope that agreement can be reached on the level of service to be provided.

The issues peculiar to Catenet monitoring, then, are those arising from the variety of component implementations to be monitored, and the lack of a single central jurisdiction over these implementations. The decisions which will need to be made in designing the CMCC are the following:

1. The level of operation, from gateway only to full Catenet.
2. The exact types of access to other parts of the Catenet which may or may not be possible.
3. The precise functions which may or may not be available from the gateways.
4. Details of the mechanism by which the CMCC can find out what Catenet access and which gateway functions are available.
5. The user interface to the CMCC functions.
6. The fault isolation functions needed, including the nature of data collection needed for accumulating the high-level fault isolation data base.

3. SYSTEM OVERVIEW.

3.1. Introduction.

This system deals only with the networks and gateways under the ARPA umbrella. Some of the issues mentioned in section 2 are not dealt with here, but will be considered some time in the future. These are:

1. Access to other parts of the Catenet.
2. Accounting information.
3. Fault isolation probing.
4. The high-level fault isolation data base.

The ARPA CMCC will initially monitor only the gateways. Higher level catenet monitoring involving network data will be considered at some future date. The CMCC will, however, maintain a map of the catenet so that it can do a certain amount of high level fault diagnosis.

The implementation will be done in stages, beginning with basic monitoring functions and later adding control and enquiry facilities. The proposed functions are listed below but we expect that the list will change as experience of using the system suggests other ideas.

Figure 1 shows the proposed system structure. The system operation is described below under the headings:

- Monitoring Operations.
- Control Operations.
- Communication with the Gateways.
- User Interface.

The user interface processes will all be Internet facilities.

For reliability purposes there will be two systems running in parallel. The two systems will not need to know about each other. An estimate of the traffic generated by the proposed system is under 200 bits per second for the current size of the ARPA catenet, so having two systems is not going to require very much bandwidth. We will however have to ensure that the gateways do not spend too much time reporting instead of handling traffic.

Should the catenet grow to a size such that we get a problem with congestion at the monitoring site, the data collection part of the system could be made distributed. We would then have a number of subcenters which would collect data from the gateways, reduce this data and forward it to the original monitoring site.

3.2. Monitoring Operations.

Monitoring operations are of three types:

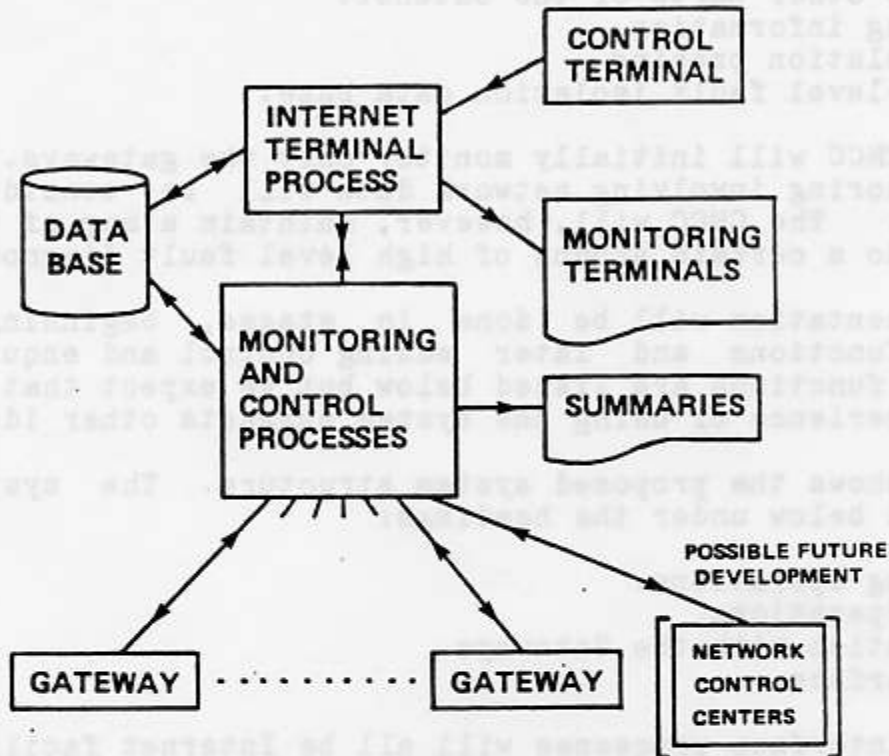


Figure 1: System Structure

1. Regular reports.
2. Trap messages.
3. Enquiry responses.

Regular reports will be turned on by a command from the main center and will continue until turned off. There will be three kinds of report:

1. Throughput statistics
2. Queue activity
3. Status report

Trap messages are issued by the catenet components in response to some event such as the failure of a network interface. The main center formats these messages and puts them into a log file. Selected messages from this file are then displayed on a terminal. Trap messages are issued when:

1. A network interface goes down.
2. A queue becomes full.
3. A neighbour gateway ceases to respond, or starts responding again.

Enquiry responses are messages sent in reply to an enquiry sent out from the main center as a result of a user command. The enquiry types currently planned are:

1. Routing information.
2. Queue status.
3. Network connection status.

3.3 Control Operations.

These will all be performed as a result of a user command.

1. Turn specified reports on or off, or alter their frequency.
2. Send enquiries.
3. Perform fault isolation probing.
4. Alter the catenet map stored in the monitoring center. This will not normally be necessary since the CMCC will keep the map up to date automatically.
5. Load a subcenter or gateway with new software.

3.4. Communication with the Gateways.

Since not all gateways will be implemented in the same way, it will be necessary for the monitoring system to find out from each gateway whether that gateway can provide the kind of report that the system is asking for. The mechanism being considered will be similar to the DO - WONT - WILL - WONT mechanism in the Telnet protocol. This mechanism

will also deal with the case where the gateway will not even respond to the enquiry. Gateways that do implement enquiries will be able to be interrogated by any program and not just the monitoring system. This will be useful for experiments and debugging.

3.5. User Interface.

The user interface has three components:

1. A control terminal.
2. Displayable summary reports.
3. Displayable monitoring reports.

The control terminal is used to initiate the control functions described above. Only one terminal can control the CMCC at a time.

Summary reports will be generated, and automatically distributed by the CMCC to authorized individuals and agencies. They will consist of hourly and daily summaries of the regular reports currently being collected by the CMCC.

Monitoring reports give a log of the events in the gateways as announced to the CMCC by trap messages. They will also contain short-term traffic and status summaries which will be produced by the CMCC. It will be possible to select the information appearing on a monitoring terminal so that, for example, one specific gateway can be investigated. It will be possible to have several monitoring terminals, each displaying different selected reports. The selection is done by the monitoring terminal user.