Gateway Monitoring Protocol

IEN 131

1 February 1980

David Flood Page

Bolt, Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02238

(617) 491-1850

GATEWAY MONITORING PROTOCOL

1. Introduction

This document details the protocol for the gateway monitoring functions described in IEN 105, 'ARPA Catenet Monitoring and Control'. It does not deal with the control functions or fault isolation; these will be covered in a separate document.

The protocol described here contains a number of report types. We realize that to implement them all may impose an unacceptable load on a gateway; therefore the system is designed to cater to gateways not implementing the complete protocol.

The protocol is described in two parts:

- Communication mechanism
- Data formats

2. Communication Mechanism

2.1 Negotiation

Because a gateway may not implement the complete protocol, the Catenet Monitoring and Control Center (CMCC) is able to discover, each time it makes a request of a gateway, whether the gateway can satisfy that request. The method used is similar to the DO - DONT - WILL - WONT mechanism in the Telnet protocol. Briefly, this works as follows:

When the CMCC wants to obtain information from a gateway, it sends a DO message to the gateway. If the gateway is able to make the required response, it returns a WILL message accompanied by the data requested. If it cannot do this, it sends back a WONT message detailing why it could not satisfy the request. If the gateway does not even implement this negotiation mechanism, or if the message is lost in transit, then the CMCC will receive no reply. In this case it will try up to two more times at 30 second intervals. If it still gets no reply, then it acts as if a WONT message had been received.

If the CMCC wants to stop the gateway from sending information, then it sends a DONT message. The gateway then responds with a WONT reply. The CMCC will try up to three times, at 30 second intervals, to get this acknowledgement.

A gateway will need to implement this negotiation mechanism in order to participate in the Monitoring and control system. This is true regardless of how many of the report types are implemented in the gateway.

2.2 Requesting Reports

Gateways may be requested to send out a series of reports at regular intervals, as well as just sending back a single response. So a DO REPORT request contains, in addition to the report type, the number of reports required and the interval between reports. The number of reports may be a special value (65535) meaning 'until further notice'. When the CMCC wants to turn off this kind of report then it sends a DONT message to the gateway. The gateway will then cease reporting and send back a WONT message. The CMCC will send up to 3 DONT messages until it gets the WONT response. If it still receives no answer then it gives up.

If a gateway is sending out regular reports, and it receives a new request from the same source as the original request to send the same report, then the new request is considered to supercede the old one unless the new request is for a single report. In this case the gateway should make the single response, but continue sending the regular reports. If the new request is for more than one report, then the gateway should reset the sequence number (see below) and forget about the original request. The question of dealing with requests from different sources is in part an authorization question, and is not dealt with in this document; however, gateways should in general be prepared to satisfy requests for single reports from any source at any time.

A gateway may be unable to send out more than one report in response to a single enquiry; i.e. it may insist on being polled. If such a gateway receives a request for multiple reports, it sends back a WONT REPORT reply, indicating that the number of reports in the request was unacceptable. The CMCC will then send a single report request, and will continue sending these requests at appropriate intervals.

Each request sent out from the CMCC contains a report identification number. This number is returned by the gateway in the WILL REPORT or WONT REPORT message. When a request results in more than one report message, those after the first have a sequence number instead of the report id. Gateways will reset the sequence number when they receive a DO REPORT, except in the case of a single report request as described above. When a regular report is requested, the WILL REPORT reply may or may not contain the first report message. If it does not, then it should consist only of the WILL REPORT header, with no extra data.

The following is a list of the report types.

Type

0 - Gateway description.
1 - Echo.
2 - Throughput transit matrix.
3 - Interface up/down status for all interfaces.
4 - Queue activity.
5 - End to end traffic statistics.
6 - Individual interface status.
7 - Routing table.

2.3 Requesting Traps

Besides the reports, a gateway may issue traps, which are messages announcing some event in the gateway. A gateway may be directed to start or stop sending the various kinds of traps, using DO - DONT - WILL - WONT TRAP messages in the same way as REPORT messages are used, except that normally the WILL TRAP message will not be accompanied by data.

The following is a list of the trap types:

Type

1 - Interface up/down.
2 - Neighbor gateway up/down.
3 - Queue full.

Here, up/down on an interface refers to the ready line. For a neighbor gateway it is determined according to the gateway-gateway protocol in force.

3. Data Formats

Bits within a field are numbered starting at 0 and ordered left to right, so that an octet with bit 0 set on has the numeric value 128. Octets within numeric fields of more than 8 bits are ordered so that the most significant octet comes first. For example, a 32 bit numeric field with a value of 65536 would be expressed as 0,1,0,0 in octets. For other fields of more than 8 bits, the first octet contains bits 0-7, the second 8-15, and so on.

Monitoring Packets have the following format:

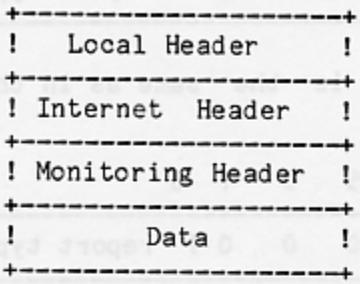

The data may be absent.

The monitoring header has the following format:

Bits	Contents
0	0 - Report or trap 1 - Negotiation message.
1	0 - Report 1 - Trap
2-3	For a negotiation message: 0 - DO 1 - DONT 2 - WILL 3 - WONT For a report or trap: zero.
4-7	Reserved for future use.
8-15	Report or trap type.
16-31	For a negotiation message: Report Id. For a report: Sequence number. For a trap: data depending on trap type.

A DO REPORT message has the header:

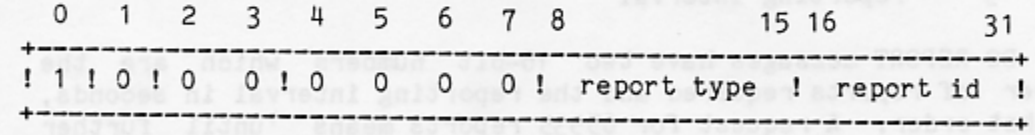

and the corresponding WILL REPORT message has:

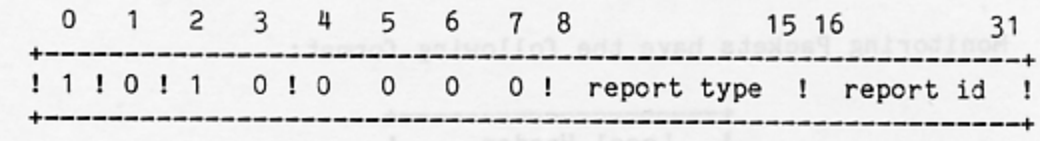

where the report id is the same as in the DO REPORT. A DONT REPORT will begin with:

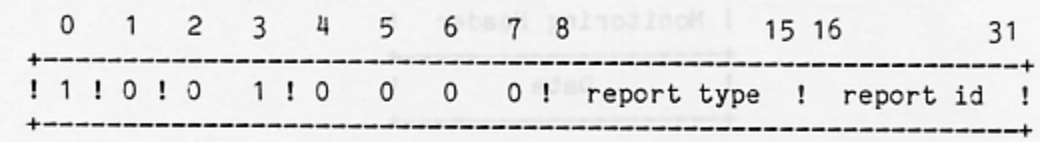

and a WONT REPORT begins with:

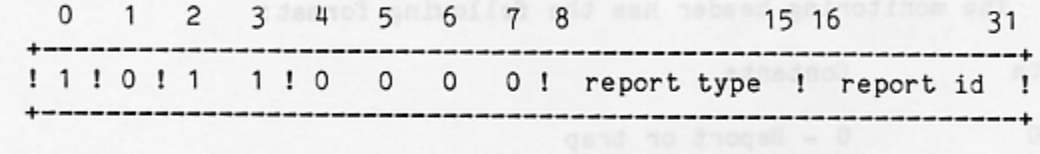

Headers for trap negotiation messages are similar except that bit 1 is 1 instead of 0.

Trap messages have a header of only 2 octets:

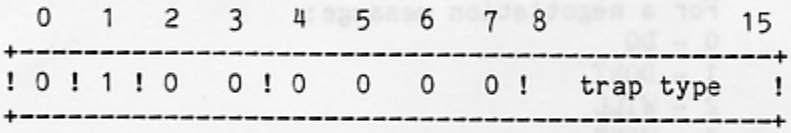

DONT messages have no data. The WONT header is followed by a single octet which indicates which field(s) in the request the gateway objected to. Bits are set on according to the offending field, as follows:

Bit	Field
0	report or trap (i.e the gateway has not implemented any reports, or traps)
1	report/trap type
2	number of reports (i.e. a gateway insists on being polled)
3	reporting interval

DO REPORT messages have two 16-bit numbers which are the number of reports required and the reporting interval in seconds, in that order. A request for 65535 reports means 'until further notice'. In addition, a type 6 report request has one extra octet at the end containing the interface number.

The first response in any set of reports may also be the WILL REPORT negotiation message and if so, the first 4 bits of the monitoring header will have the value 1010 (negotiation, report,

WILL). Subsequent reports arising from the same request have a header beginning with 0000 (report/trap, report, zero). If the first response is the WILL REPORT without any data, then its length must be 4 bytes, i.e. it consists only of the monitoring header.

Trap messages may or may not have any data, depending on the trap type.

3.1 Report Formats

3.1.1 Gateway description - type 0

The first item is the gateway name as four 8-bit ASCII characters. The next item consists of two octets containing the number of interfaces in the gateway, and the number of neighbors the gateway has, in that order. This is then followed by two sets of 32 bit numbers, whose size is given by the above octets. The first set lists the addresses of each interface in the gateway, and the second set lists the addresses of the gateway's neighbors.

3.1.2 Echo - type 1

There is no data in this message type. The gateway simply returns the message to the place that sent it.

3.1.3 Throughput matrix - type 2

The report is a conceptual matrix with rows corresponding to output interfaces and columns to input interfaces. The interfaces are numbered from 0 to N-1 and there is an extra column for packets dropped at the interface.

The matrix is expressed as N * (N+1) 32-bit counts, where N is the number of interfaces. Each packet entering the gateway via interface IN and leaving via interface OUT causes the count at position (OUT * N) + IN to be incremented.

3.1.4 Status of all interfaces - type 3

The header is followed by a bit array in which the bit in position i is 1 if interface i is up, 0 if it is down. Interfaces are numbered starting at zero, as in the throughput matrix. The ordering of the interfaces is defined in the Gateway Description message, 3.1.1.

3.1.5 Queue activity - type 4

The header is followed by a set of reports, one for each interface number. Each report in the set is 16 bits long and has the following format:

Bits	Contents
0-7	Length of input queue for this interface.
8-15	Length of output queue.

Interface numbering is as in the interface status message.

3.1.6 End to end statistics - type 5

The report has a set of counts, one for each source/destination combination. The format of each entry is:

Bits	Contents
0-7	Source network number.
8-15	Destination network number.
16-47	Count of packets source-destination.

The counts are cumulative and so is the list of source/destination combinations, i.e. the report will contain counts for every source/destination pair that has been recorded since the gateway started up.

3.1.7 Individual interface status - type 6

A distinction here is made between error free and error handling interfaces. The first four octets are the same in each case, except for a code indicating the interface type. For an error free interface, these four octets are the whole report. For a VDH error handling interface there are another three 32-bit counts of :

packet framing errors
packets received with bad checksum
packets retransmitted

The format of the first four octets is:

Bits	Contents
0-7	Interface number
8-11	Status: 0 (down), 1 (up)
12-15	Interface type: 0 - error free, 1 - VDH.
16-31	Number of times this interface has gone down.

The down count is only reset at gateway startup time.

3.1.8 Routing tables - type 7

This is a table of variable length entries each containing a network number, the minimum distance to that network from the gateway, and the addresses of each neighbor on the minimum distance path. The format of each entry is as follows:

8 bits number of neighbors
8 bits network number
8 bits distance to network
8 bits unused (allows 32-bit alignment of addresses)
32 bits first neighbor address
32 bits second neighbor address
(as many more neighbor addresses as necessary).

3.2 Trap Formats

Traps all have a 16-bit header starting with 0100 (report/trap, trap, zero). Data for the traps is as follows.

3.2.1 Interface up/down - type 1

Bits	Contents
0-7	up (1) or down (0).
8-15	interface number.

3.2.2 Neighbor gateway up/down - type 2

Bits	Contents
0-3	up (1) or down (0).
4-7	old gateway (zero) or new gateway (1).
8-15	unused (for 32-bit alignment of next field)
16-47	Neighbor gateway internet address.

A new gateway is one not previously heard from, which will therefore cause an addition to the gateway's routing tables.

3.2.3 Queue full - type 3

Bits	Contents
0-7	Interface number for queue.
8-15	Input (zero) or output (1) queue.