Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2733172 C 2011/10/25

(11)(21) 2 733 172

(12) BREVET CANADIEN
CANADIAN PATENT

(13) C

(22) Date de dépôt/Filing Date: 2004/05/07

(41) Mise à la disp. pub./Open to Public Insp.: 2005/11/07

(45) Date de délivrance/Issue Date: 2011/10/25

(62) Demande originale/Original Application: 2 466 567

(51) Cl.Int./Int.Cl. *H04L 12/26* (2006.01),
H04L 12/58 (2006.01)

(72) Inventeurs/Inventors:
BOWMAN, DON, CA;
BEDI, HARMEET SINGH, CA

(73) Propriétaire/Owner:
SANDVINE INCORPORATED ULC, CA

(74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : SYSTEME ET METHODE POUR DETECTER LES SOURCES DE MESSAGES DE RESEAU INFORMATIQUE ANORMAUX

(54) Title: A SYSTEM AND METHOD FOR DETECTING SOURCES OF ABNORMAL COMPUTER NETWORK MESSAGES

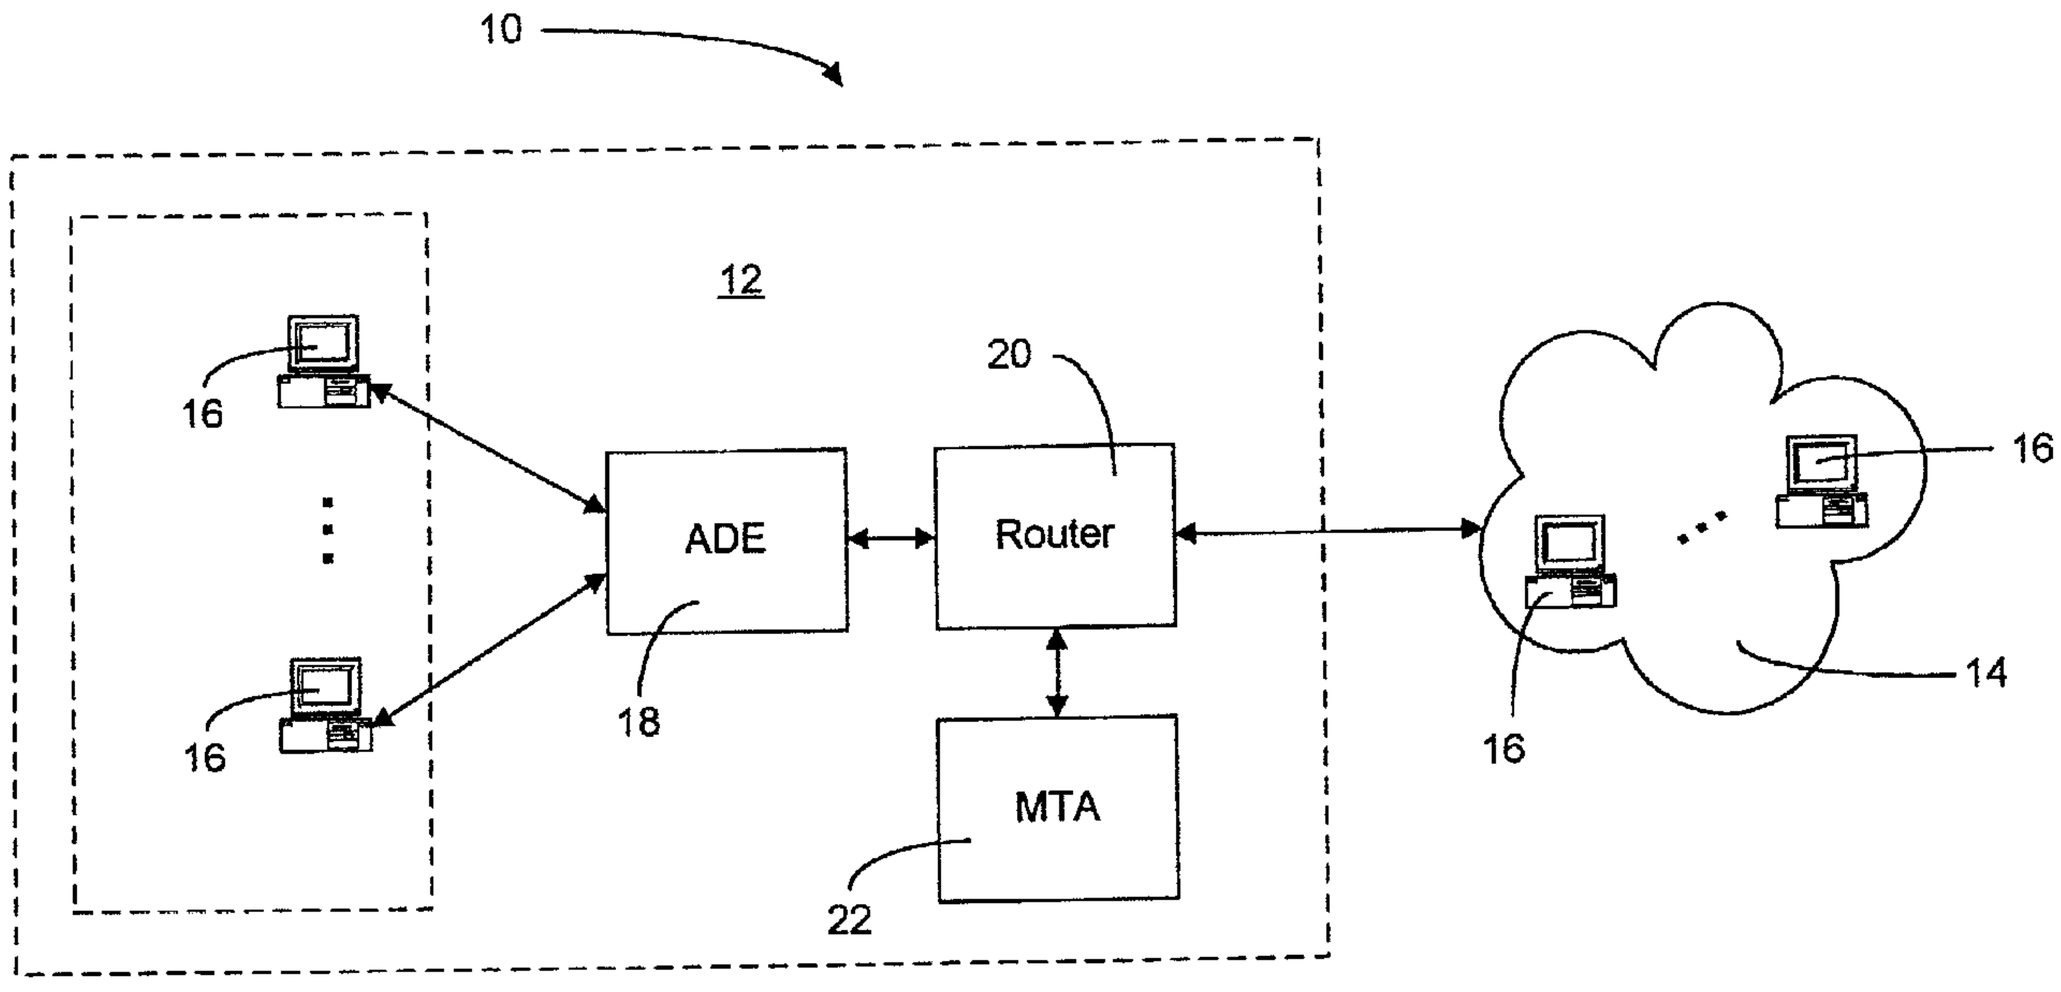

(57) Abrégé/Abstract:
A system for detecting a source or destination of abnormal message traffic on a network, the system having: an abnormality detection engine configured to track messages between a source and a destination and store the source and destination of a message as a pair, wherein the abnormality detection engine includes: a unique counter for each source and destination pair wherein the counter is incremented based on at least some of the messages between the source and the destination; and at least one abnormality detector configured to determine the credibility of the source or destination based on the counter and at least one threshold.

Abstract

A system for detecting a source or destination of abnormal message traffic on a network, the system having: an abnormality detection engine configured to track messages between a source and a destination and store the source and destination of a message as a pair, wherein the abnormality detection engine includes: a unique counter for each source and destination pair wherein the counter is incremented based on at least some of the messages between the source and the destination; and at least one abnormality detector configured to determine the credibility of the source or destination based on the counter and at least one threshold.

A SYSTEM AND METHOD FOR DETECTING SOURCES OF ABNORMAL COMPUTER NETWORK MESSAGES

Field

The present invention relates generally to a system and method for detecting abnormal patterns of computer message traffic, the intent being to determine if a host should be ignored as it appears to be sending bulk email, viruses, worms or the like.

Background

With the mass growth of the Internet there has occurred a rising flood of unwanted messages. Many of these messages are what are typically referred to as "spam". Spam is the electronic equivalent of junk mail. In addition to junk mail, other messages may include programs such as viruses or worms. One of the intents of a worm is to control a host computer for the purpose of sending more spam. Spam consumes a large amount of network resources as well as wasted time for the users having to deal with it.

There have been many solutions developed to deal with spam and unwanted messages. The most common being the use of filtration software. Filtration software examines the content of a message and determines if the message is wanted or not. Typically filtration software maintains a database of sites known for sending unwanted messages as well as databases of keywords that help to identify an unwanted message. Such a scheme is costly in the use of computer time, as it must scan every message for content and check with a database. Further, it is simple to avoid filtration software by changing the address of the sender and modifying the words of the message. Finally, filtration software may exclude wanted messages based upon what is falsely considered a valid keyword or address match.

An advancement in filtration software is to use Bayesian or heuristic filters to statistically identify unwanted messages based on the frequencies of patterns in the message. These types of filters are weak when dealing with shorter messages, as they do not have enough data to make an intelligent decision.

Another alternative is to create lists of IP addresses that are known to be used by senders of unwanted messages. These are known as "blacklists" and aid in blocking messages from the listed addresses. The problem with this approach is that the blacklisted

senders move addresses readily and the person who is reassigned the previous address may still be on the list, thus being incorrectly identified as a spammer.

Thus, there is a need for a means of detecting unwanted messages in a cost effective and efficient manner. The present invention addresses this need.

Summary

The present invention is directed to a method for detecting sources of abnormal message traffic on a network, said method comprising the steps of:

a) utilizing an abnormality detection engine to detect said abnormal message traffic; and

b) reporting on said abnormal message traffic.

The present invention is also directed to a method of wherein said abnormality detection engine consists of one or more of components selected from the set of: a fanout detector, a fanin detector, an error response detector; a bandwidth variation detector; or a message content detector.

The present invention is also directed to a system for detecting sources of abnormal traffic in a network, said system comprising an abnormality detection engine, said abnormality detection engine accepting messages to and from said network and providing a report as output, said abnormality detection engine comprising one or more abnormality detectors, selected from the set of: a fanout detector, a fanin detector, an error response detector, a bandwidth variation detector; or a variation in message content detector.

The present invention is further directed to a computer readable medium, for detecting sources of abnormal message traffic on a network, said medium comprising instructions for:

a) utilizing an abnormality detection engine to detect said abnormal message traffic; and

b) reporting on said abnormal message traffic.

The computer readable medium, wherein said abnormality detection engine consists of instructions for one or more of a fanout detector, a fanin detector, an error response detector, a bandwidth variation detector; or a variation in message content detector.

In another aspect, there is provided a system for detecting a source or destination of abnormal message traffic on a network, the system having: an abnormality detection engine configured to track messages between a source and a destination and store the source and destination of a message as a pair, wherein the abnormality detection engine includes: a unique counter for each source and destination pair wherein the counter is incremented based on at least some of the messages between the source and the destination; and at least one abnormality detector configured to determine the credibility of the source or destination based on the counter and at least one threshold.

In some cases, the traffic of the system is email. In some other cases the traffic comprises Hypertext Transfer Protocol messages.

In some cases, the abnormality detector of the system for detecting a source or destination of abnormal message traffic is configured to detect abnormal messages received by a single destination by comparing the number of distinct source and destination pairs having the same destination. In some other cases, the abnormality detector is configured to detect abnormal messages from a single source by comparing the number of distinct source and destination pairs having the same source.

In some cases, the at least one abnormality detector may be an error response detector configured to detect an abnormal amount of error response codes. In some embodiments, the at least one abnormality detector is a bandwidth variation detector configured to detect a steady rate of messages. In some other embodiments, the at least one abnormality detector is a variation in message content detector configured to detect if messages originating from a single source have largely the same content.

In a further aspect, there is provided a system for detecting a source or destination of abnormal message traffic on a network, the system having: an abnormality detection engine configured to track messages between a source and a destination and store the source and destination of a message as a pair, wherein the abnormality detection engine comprises: a unique counter for each source and destination pair wherein the counter is incremented based on at least some of the messages between the source and the destination; and at least one abnormality detector configured to determine the credibility of the source or destination based on the counter and at least one threshold wherein the at least on abnormality detector is at least one of an error response detector configured to

detect an abnormal amount of error response codes, a bandwidth variation detector configured to detect a steady rate of messages, or a variation in message content detector configured to detect if messages originating from a single source have largely the same content.

In some cases, the traffic of the system is email. In some other cases the traffic comprises Hypertext Transfer Protocol messages.

In some cases, the abnormality detector of the system for detecting a source or destination of abnormal message traffic is configured to detect abnormal messages received by a single destination by comparing the number of distinct source and destination pairs having the same destination. In some other cases, the abnormality detector is configured to detect abnormal messages from a single source by comparing the number of distinct source and destination pairs having the same source.

Brief Description of the Drawings

For a better understanding of the present invention, and to show more clearly how it may be carried into effect, reference will now be made, by way of example, to the accompanying drawings which aid in understanding an embodiment of the present invention and in which:

Figure 1 is a block diagram illustrating how the present invention may be utilized;

Figure 2 is a block diagram of the functional components of an Abnormality Detection Engine;

Figure 3 is a flowchart of the logical structure of the fanout detector;

Figure 4 is a flowchart of the logical structure of the error response detector;

Figure 5 is a flowchart of the logical structure of the bandwidth variation detector; and

Figure 6 is a flowchart of the logical structure of the variation in message content detector.

Detailed Description

The present invention is referred to as an "Abnormality Detection Engine", ADE. It is not the intent of the inventors to restrict the use of the invention simply to the

detection of spam, but rather to allow it to be utilized to detect any form of unwanted messages.

Referring now to Figure 1, a block diagram illustrating how the present invention may be utilized is shown generally as system 10. System 10 comprises an Internet Service Provider (ISP) network 12 and an external network 14. Messages, such as email are exchanged by hosts 16, between networks 12 and 14. Each host 16 is capable of sending and receiving messages. In the case of email, each host 16 will utilize a Mail User Agent (MUA, not shown) such as Microsoft Outlook to send and receive messages. All messages sent between networks 12 and 14 will pass through ADE 18. ADE 18 monitors messages and passes them to or receives them from a router 20. In the case of email messages a Mail Transfer Agent (MTA) 22 is utilized to forward or receive messages. In system 10, MTA 22 is shown as being part of network 12 but it may also reside within network 14.

System 10 is meant merely to indicate how the present invention, residing within ADE 18 may be deployed. As one skilled in the art will recognize, any number of configurations may be utilized to make use of the present invention. By way of example, ADE 18 may reside outside ISP network 12.

Referring now to Figure 2 a block diagram of the functional components of an Abnormality Detection Engine is shown. ADE 18 takes as input a data stream 30 and provides as output a stream of reporting data 32. Stream 30 comprises all messages to be monitored by ADE 18. Stream 32 may take any number of forms such as being stored in a database, being displayed to a system administrator graphically, or formatted in reports. The intent of stream 32 is to provide those interested with information on abnormal messages.

ADE 18 comprises five main components, each of which serves as detectors of anomalies in network traffic. One or more components may be enabled and configured for a specific implementation. Fanout detector 34 examines data stream 30 to determine if an abnormal amount of messages are being sent (Fanout) by a host to multiple addresses. By the term address we mean to include: an IP address, a domain name, an email address and any other means for identifying a unique source or recipient of a message. Fanout can be an indication that a host is sending too many unwanted messages. Fanin detector 36

examines data stream 30 to determine if an abnormal amount of traffic is being received from a single address. Error response detector 38 looks for an abnormal amount of error messages. Messages incorrectly addressed to an MUA are an indication of unwanted messages. Bandwidth variation detector 40 determines if a sender of messages is providing a steady rate of messages. A steady rate of messages is not typical of human use of a network and indicates a source of unwanted messages. Variation in message content detector 42 examines messages to determine if messages coming from a single source are largely the same.

Figure 3 is a flowchart of the logical structure of the fanout detector, shown as feature 34 of Figure 2. Fanout is a measure of distinct addresses. A typical MUA may utilize a few MTAs, so an indication of an increase in addresses may help in determining if a host is being utilized to deliver unwanted messages.

To describe the fanout detector in more detail, we begin at step 34a. At step 34a information on the source and destination of the current message are extracted. Typically these would be IP addresses, but they could also be domain names or email addresses. By way of example, SMTP response messages may be monitored through the use of a packet capture library to monitor TCP/IP port 25 for email. At step 34b a test is made to determine if the source and destination can be determined, if so, the fanout counter for the source and destination pair is incremented at step 34c. In the case of SMTP messages, the fanout counter would count the number of messages sent to each unique address. At step 34d a test is made to determine if it is time to generate a report on the information collected, if not processing moves to step 34e where processing for the current message ends. If it is determined at step 34d that a report should be prepared, processing moves to step 34f. At step 34f a test is made to determine if the threshold for fanout has been met. Experimentation indicates that a threshold value of 20 for each unique address is an indication of sending spam. If the threshold has not been met, processing moves to step 34h. If the threshold has been met, processing moves to step 34g. At step 34g reporting data is prepared to indicate that the destination IP address is a source of abnormal traffic. This report corresponds to reporting data 32 of Figure 2. The user may wish to reset fanout counters in a deterministic manner, for example on regular schedule, or on memory used. At step 34h it is determined if the fanout counters should be reset. If not,

processing returns to step 34e. If the fanout counters need to be reset, this is done at step 34i.

Fanin detector 36 functions in a similar manner as fanout detector 36. The distinction being that fanin detector 36 examines messages to determine if an abnormal number of messages have been received from a unique address as opposed to messages being sent. The logic for fanin detector 36 is identical to that shown in the flowchart of Figure 3, save that the counters track fanin rather than fanout.

Referring now to Figure 4 a flowchart of the logical structure of the error response detector, feature 38 of Figure 2 is shown. Error response detector 38 examines messages to determine if a message is a "reject" message. By way of example, in the case of email an MTA may reject a message and make it known to the sender. Similarly in the case of HTTP a URL may not be found, resulting in a reject message. A well behaved MUA is not likely to receive more than a few reject messages. A large number of reject messages is an indicator of abnormal messages.

Beginning at step 38a the response to a message from an MTA is read. At step 38b, if the message is not an error response it is ignored at step 38c. If the message indicates an error response, processing moves to step 38d where a counter for the MTA is incremented. At step 38e a test is made to determine if a report, shown as feature 32 of Figure 2, should be generated. If no report is required, processing ends at step 38c. If a report is required, processing moves to step 38f where a test is made to determine if a threshold has been met to require the generation of a report. Experimentation has shown that for SMTP messages an error count of ten messages from a unique address is an indication of spam. If the threshold has been met, processing moves to step 38g and a report is generated. If not, processing moves to step 38i. At step 38i a test is made to determine if the error counters should be initialized. The user may wish to initialize the error counters in a deterministic manner, for example on a regular schedule, or on memory used. If so, processing moves to step 38h to initialize the error counters, if not processing for the message ends at step 38c.

Referring now to Figure 5 a flowchart of the logical structure of the bandwidth variation detector, feature 40 of Figure 2 is shown. Beginning at step 40a, a message is read to determine the destination address of the message. At step 40b a counter

corresponding to the destination address is updated. At step 40c a test is made to determine if it is time to generate a report on bandwidth variation. If the result is negative, processing moves to step 40d and the message is ignored. If the result is positive a calculation is made on bandwidth variation. The intent here is to detect anomalies in message traffic. Typically messages from an MUA would be in bursts, consistent traffic may be indicative of a spam host. Any number of schemes may be used to determine if an abnormality in bandwidth variation exists. The use of a moving average has been found to work well. A test is then made at step 40f to determine if the desired threshold for bandwith variation has been met. If so, a report, shown as feature 32 of Figure 2, is generated at step 40g, if not, processing moves to step 40h. At step 40h a test is made to determine if the bandwidth counters should be initialized. Counter values may take up more memory than desired or a user may wish to have them reset on a regular basis. If counters are to be initialized processing moves to step 40i, otherwise to step 40d.

Referring now to Figure 6 a flowchart of the logical structure of the variation in message content detector, feature 42 of Figure 2 is shown. Beginning at step 42a, a message is read to determine the content of the message. For unwanted messages such as spam, the message content will scarcely vary. A number of algorithms may be used to detect variation in content, such as hashing the content of the message or a variety of Lempel-Ziv, Huffman encoding or the like. It is not the intent of the inventors to restrict the variation in message content detector to any one algorithm. At step 42b a test is made to determine if the message is similar to others sent from the same address, if so the counter corresponding to the address of the source of the message is updated at step 42c. At step 42d a test is made to determine if it is time to generate a report on variation in message content. If the result is negative, processing moves to step 42e and the message is ignored. If the result is positive, a test is conducted at step 42f to determine if the desired threshold for message variation has been met. If so, a report is generated at step 42g, if not processing moves directly to step 42h. At step 42h a test is made to determine if the variation counters should be initialized. Counter values may take up more memory than desired, and from time to time it may be desired to reset them. If counters are to be initialized processing moves to step 42i, otherwise to step 42e.

Another feature of the present invention, not shown, is to utilize a "white list" within ADE 18. A white list would include information on trusted sources of messages. A message coming from a source on the white list would not be examined by ADE 18.

In this disclosure, the inventors intend the term "counter" to refer to a count of the number of messages for a given address tracked by an abnormality detector, regardless of the abnormality detector in use. If the counter exceeds the threshold for an abnormality detector, a report is generated. For example, if a standard deviation were to be used to detect abnormal messages, the counter would be incremented for those messages that lie on the tails of the distribution.

Although the present invention has been described as being a software based invention, it is the intent of the inventors to include computer readable forms of the invention. Computer readable forms meaning any stored format that may be read by a computing device.

Although the invention has been described with reference to certain specific embodiments, various modifications thereof will be apparent to those skilled in the art without departing from the spirit and scope of the invention as outlined in the claims appended hereto.

Claims:

1. A system for detecting a source or destination of abnormal message traffic on a network, the system comprising:

an abnormality detection engine configured to track messages between a source and a destination and store the source and destination of a message as a pair, wherein the abnormality detection engine comprises:

a unique counter for each source and destination pair wherein the counter is incremented based on at least some of the messages between the source and the destination; and

at least one abnormality detector configured to determine the credibility of the source or destination based on the counter and at least one threshold.

2. The system of claim 1 wherein the traffic is email.

3. The system of claim 1 wherein the traffic comprises Hypertext Transfer Protocol messages.

4. The system of claim 1 wherein the abnormality detector is configured to detect abnormal messages received by a single destination by comparing the number of distinct source and destination pairs having the same destination.

5. The system of claim 1 wherein the abnormality detector is configured to detect abnormal messages from a single source by comparing the number of distinct source and destination pairs having the same source.

6. The system of claim 1 wherein the at least one abnormality detector is an error response detector configured to detect an abnormal amount of error response codes.

7. The system of claim 1 wherein the at least one abnormality detector is a bandwidth variation detector configured to detect a steady rate of messages.

8. The system of claim 1 wherein the at least one abnormality detector is a variation in message content detector configured to detect if messages originating from a single source have largely the same content.

9. A system for detecting a source or destination of abnormal message traffic on a network, the system comprising:

an abnormality detection engine configured to track messages between a source and a destination and store the source and destination of a message as a pair, wherein the abnormality detection engine comprises:

a unique counter for each source and destination pair wherein the counter is incremented based on at least some of the messages between the source and the destination; and

at least one abnormality detector configured to determine the credibility of the source or destination based on the counter and at least one threshold wherein the at least on abnormality detector is at least one of an error response detector configured to detect an abnormal amount of error response codes, a bandwidth variation detector configured to detect a steady rate of messages, or a variation in message content detector configured to detect if messages originating from a single source have largely the same content.

10. The system of claim 9 wherein the traffic is email.

11. The system of claim 9 wherein the traffic comprises Hypertext Transfer Protocol messages.

12. The system of claim 9 wherein the abnormality detector is configured to detect abnormal messages received by a single destination by comparing the number of distinct source and destination pairs having the same destination.

13. The system of claim 9 wherein the abnormality detector is configured to detect abnormal messages from a single source by comparing the number of distinct source and destination pairs having the same source.

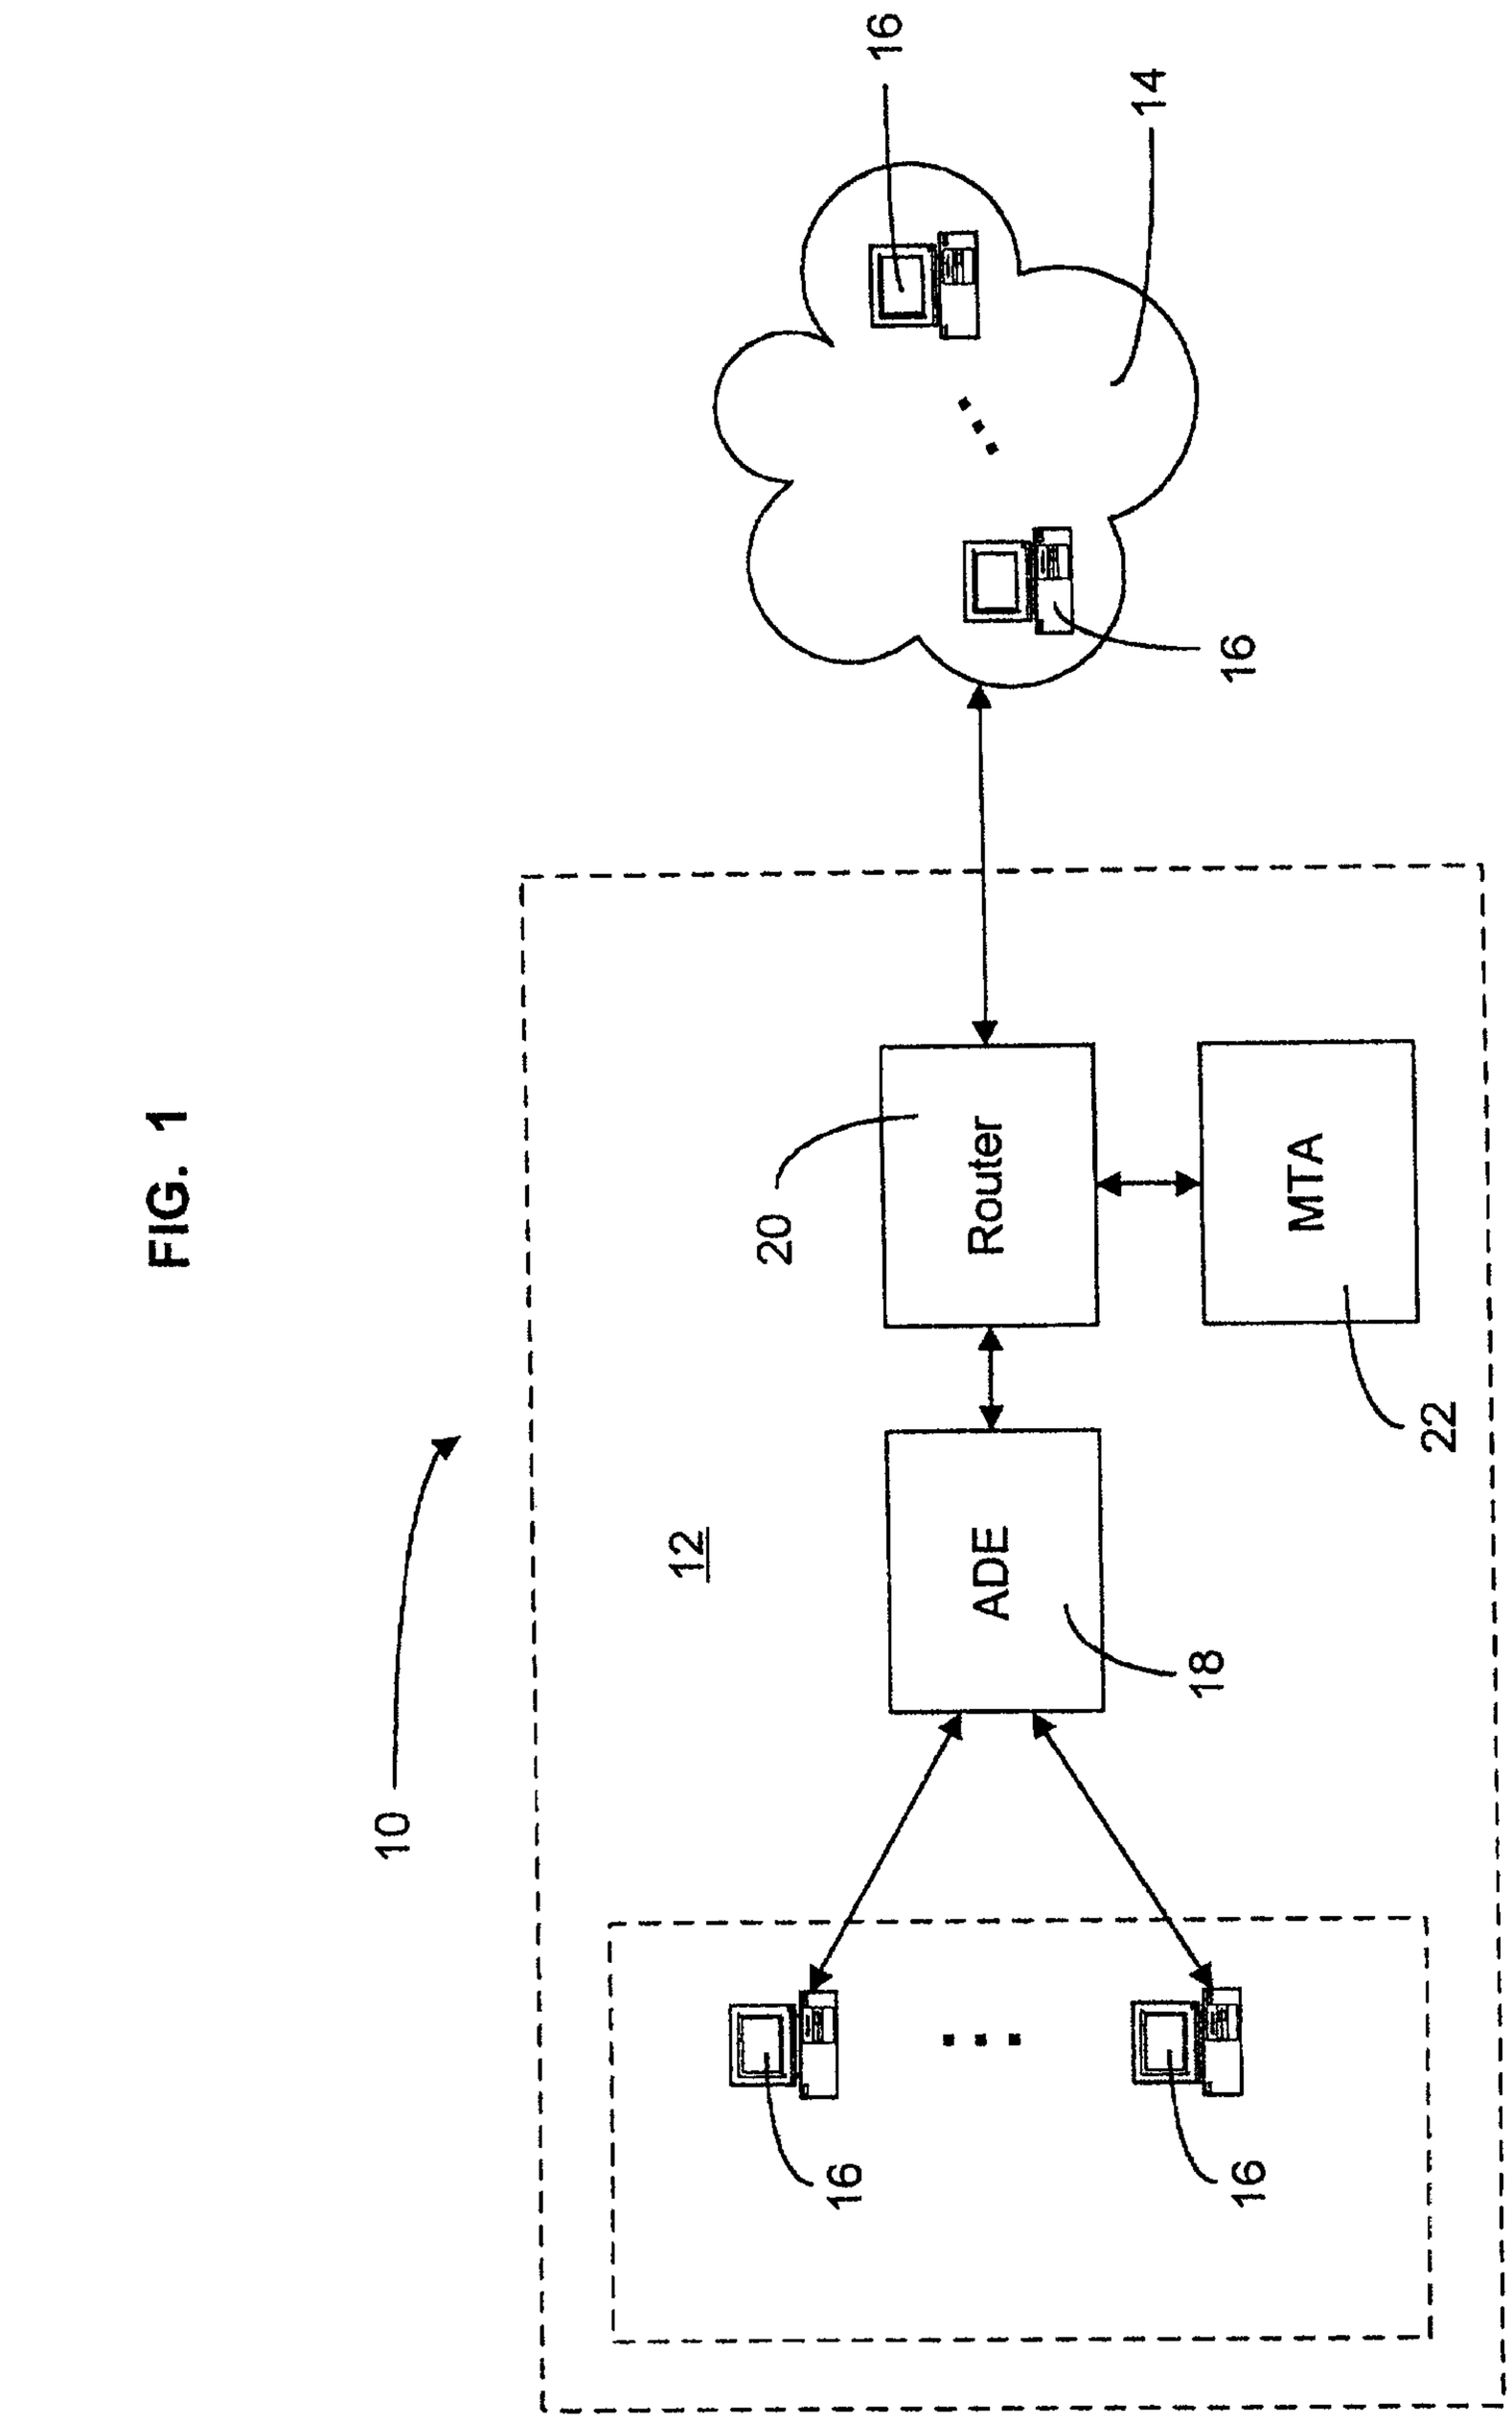

FIG. 1

FIG. 2

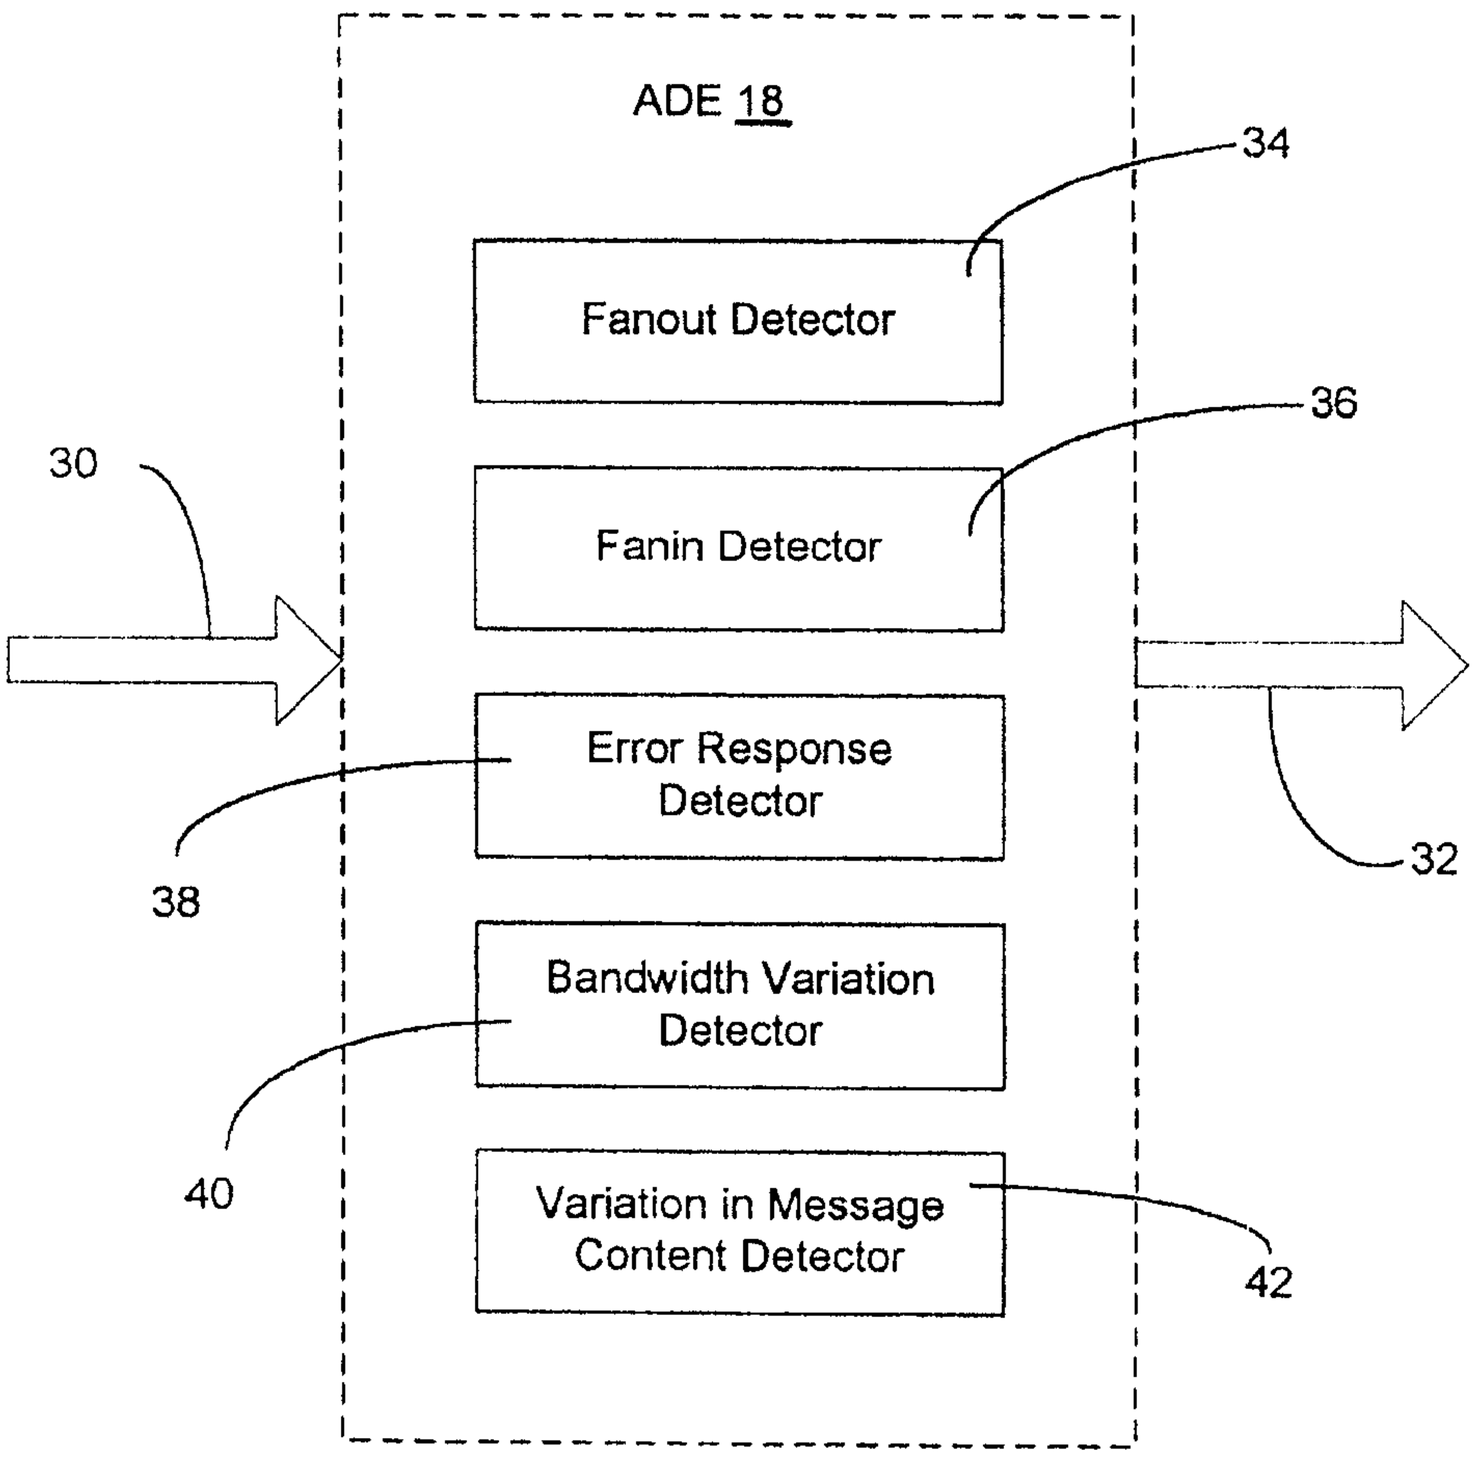
ADE 18
34
Fanout Detector
36
Fanin Detector
30
Error Response Detector
38
Bandwidth Variation Detector
40
Variation in Message Content Detector
42
32

FIG. 3

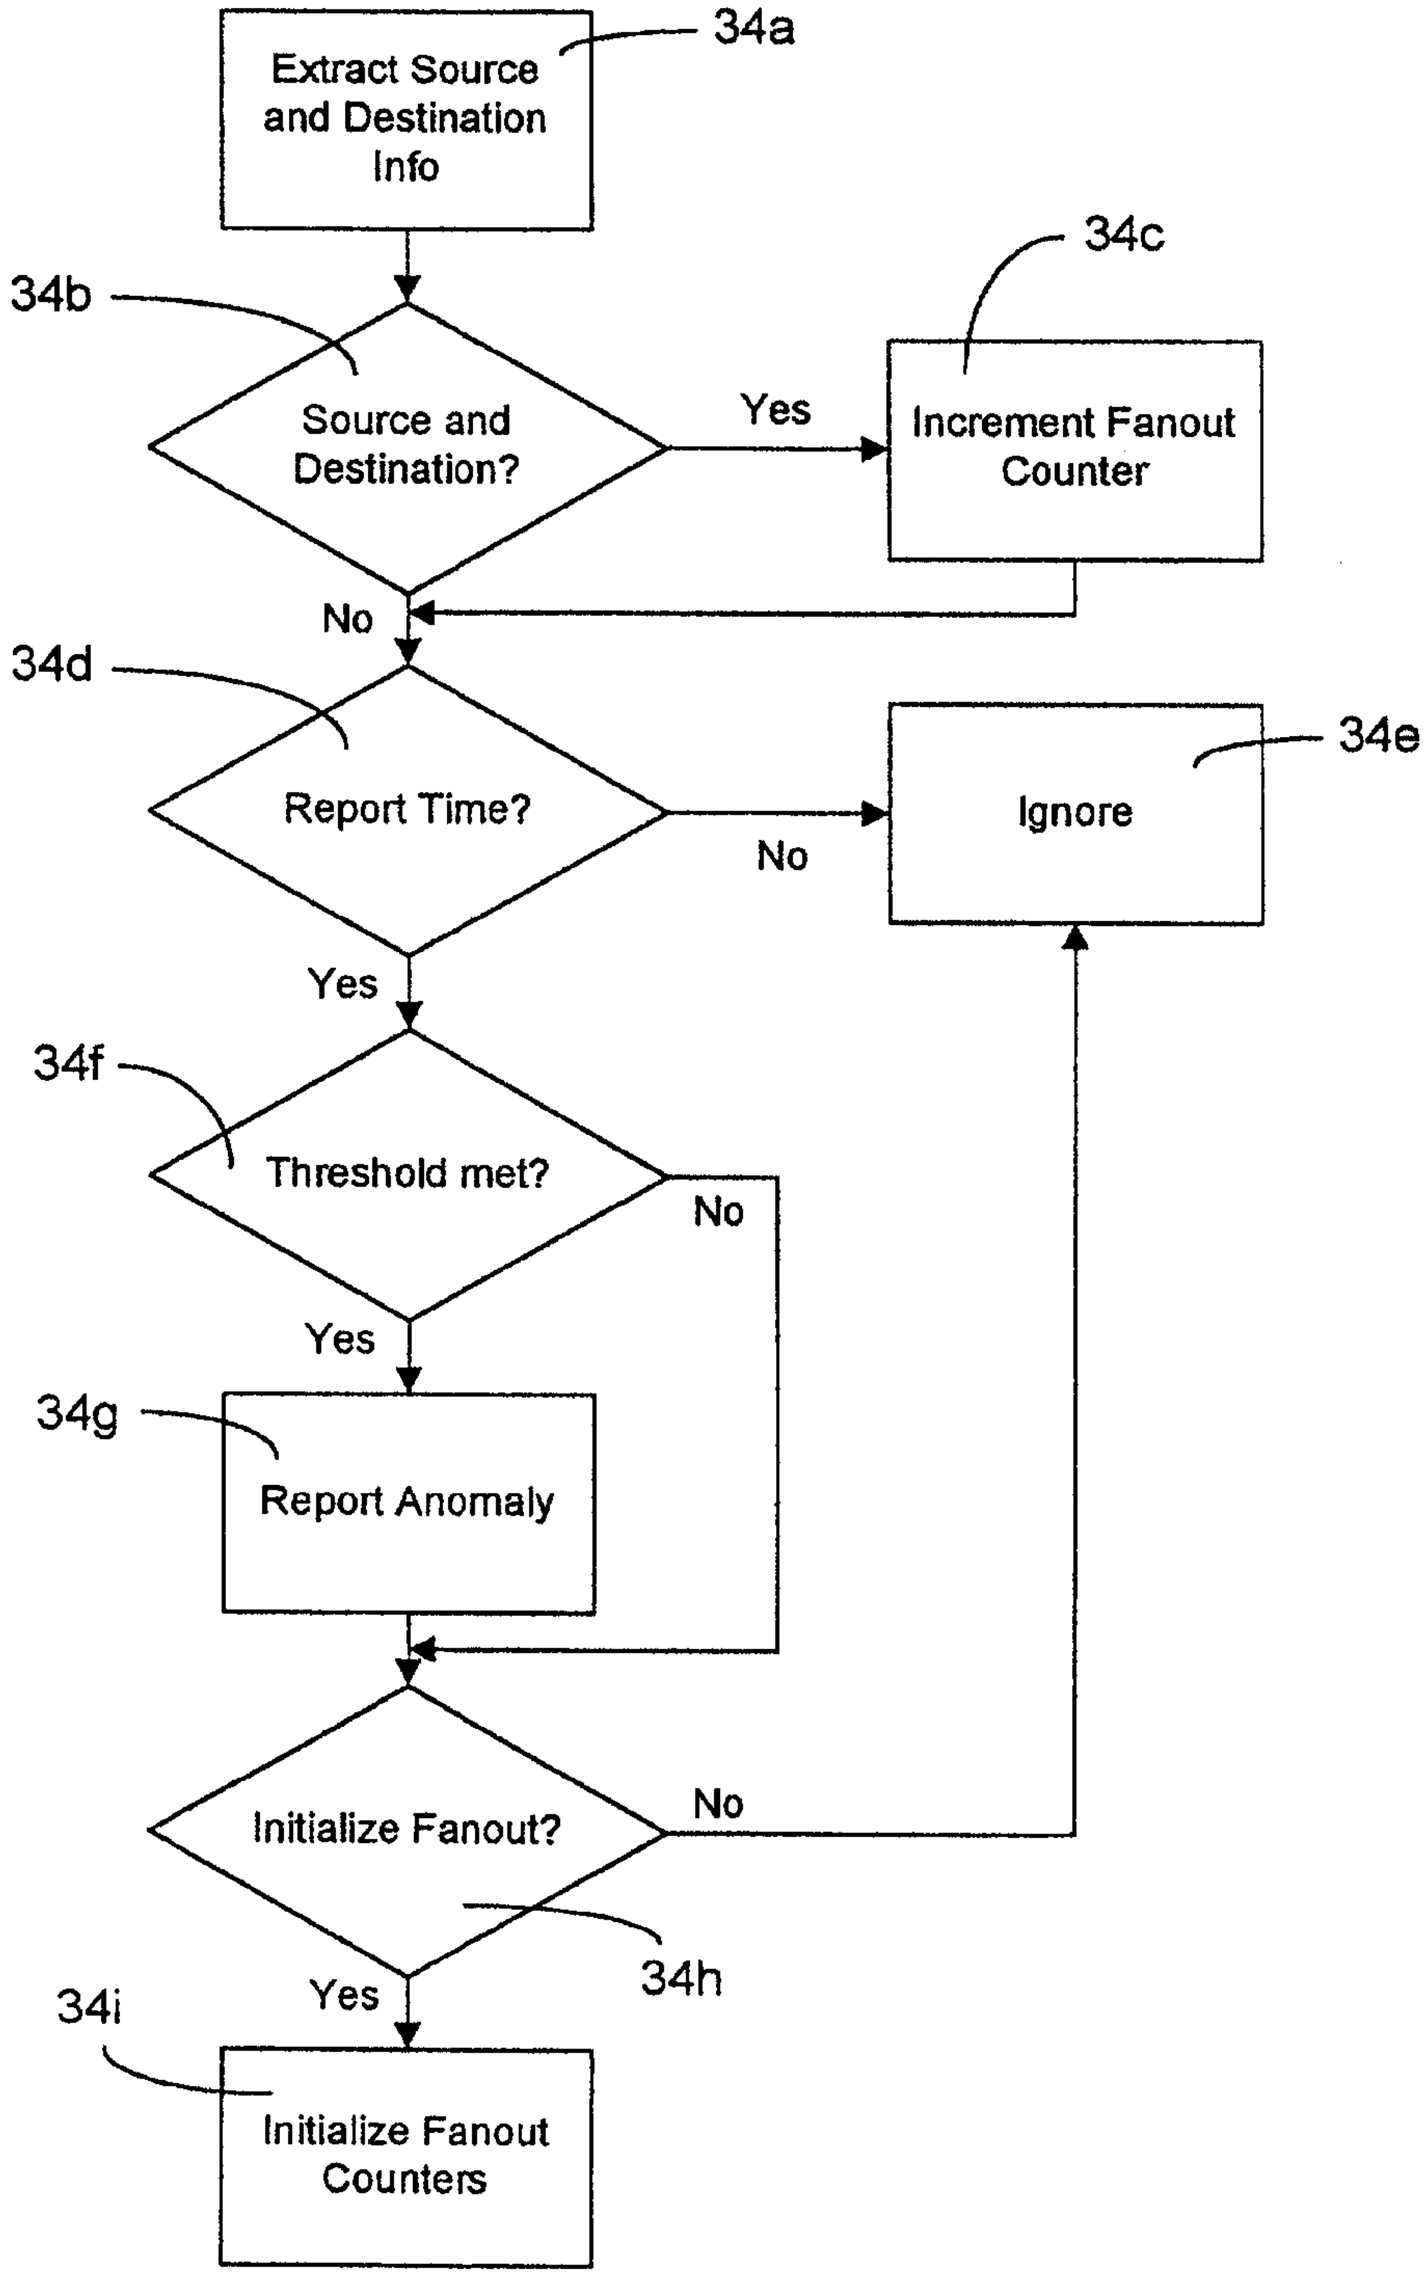
Extract Source and Destination Info
34a
34b
Source and Destination?
Yes
34c
Increment Fanout Counter
No
34d
Report Time?
No
34e
Ignore
Yes
34f
Threshold met?
No
Yes
34g
Report Anomaly
Initialize Fanout?
No
34h
Yes
34i
Initialize Fanout Counters

FIG. 4

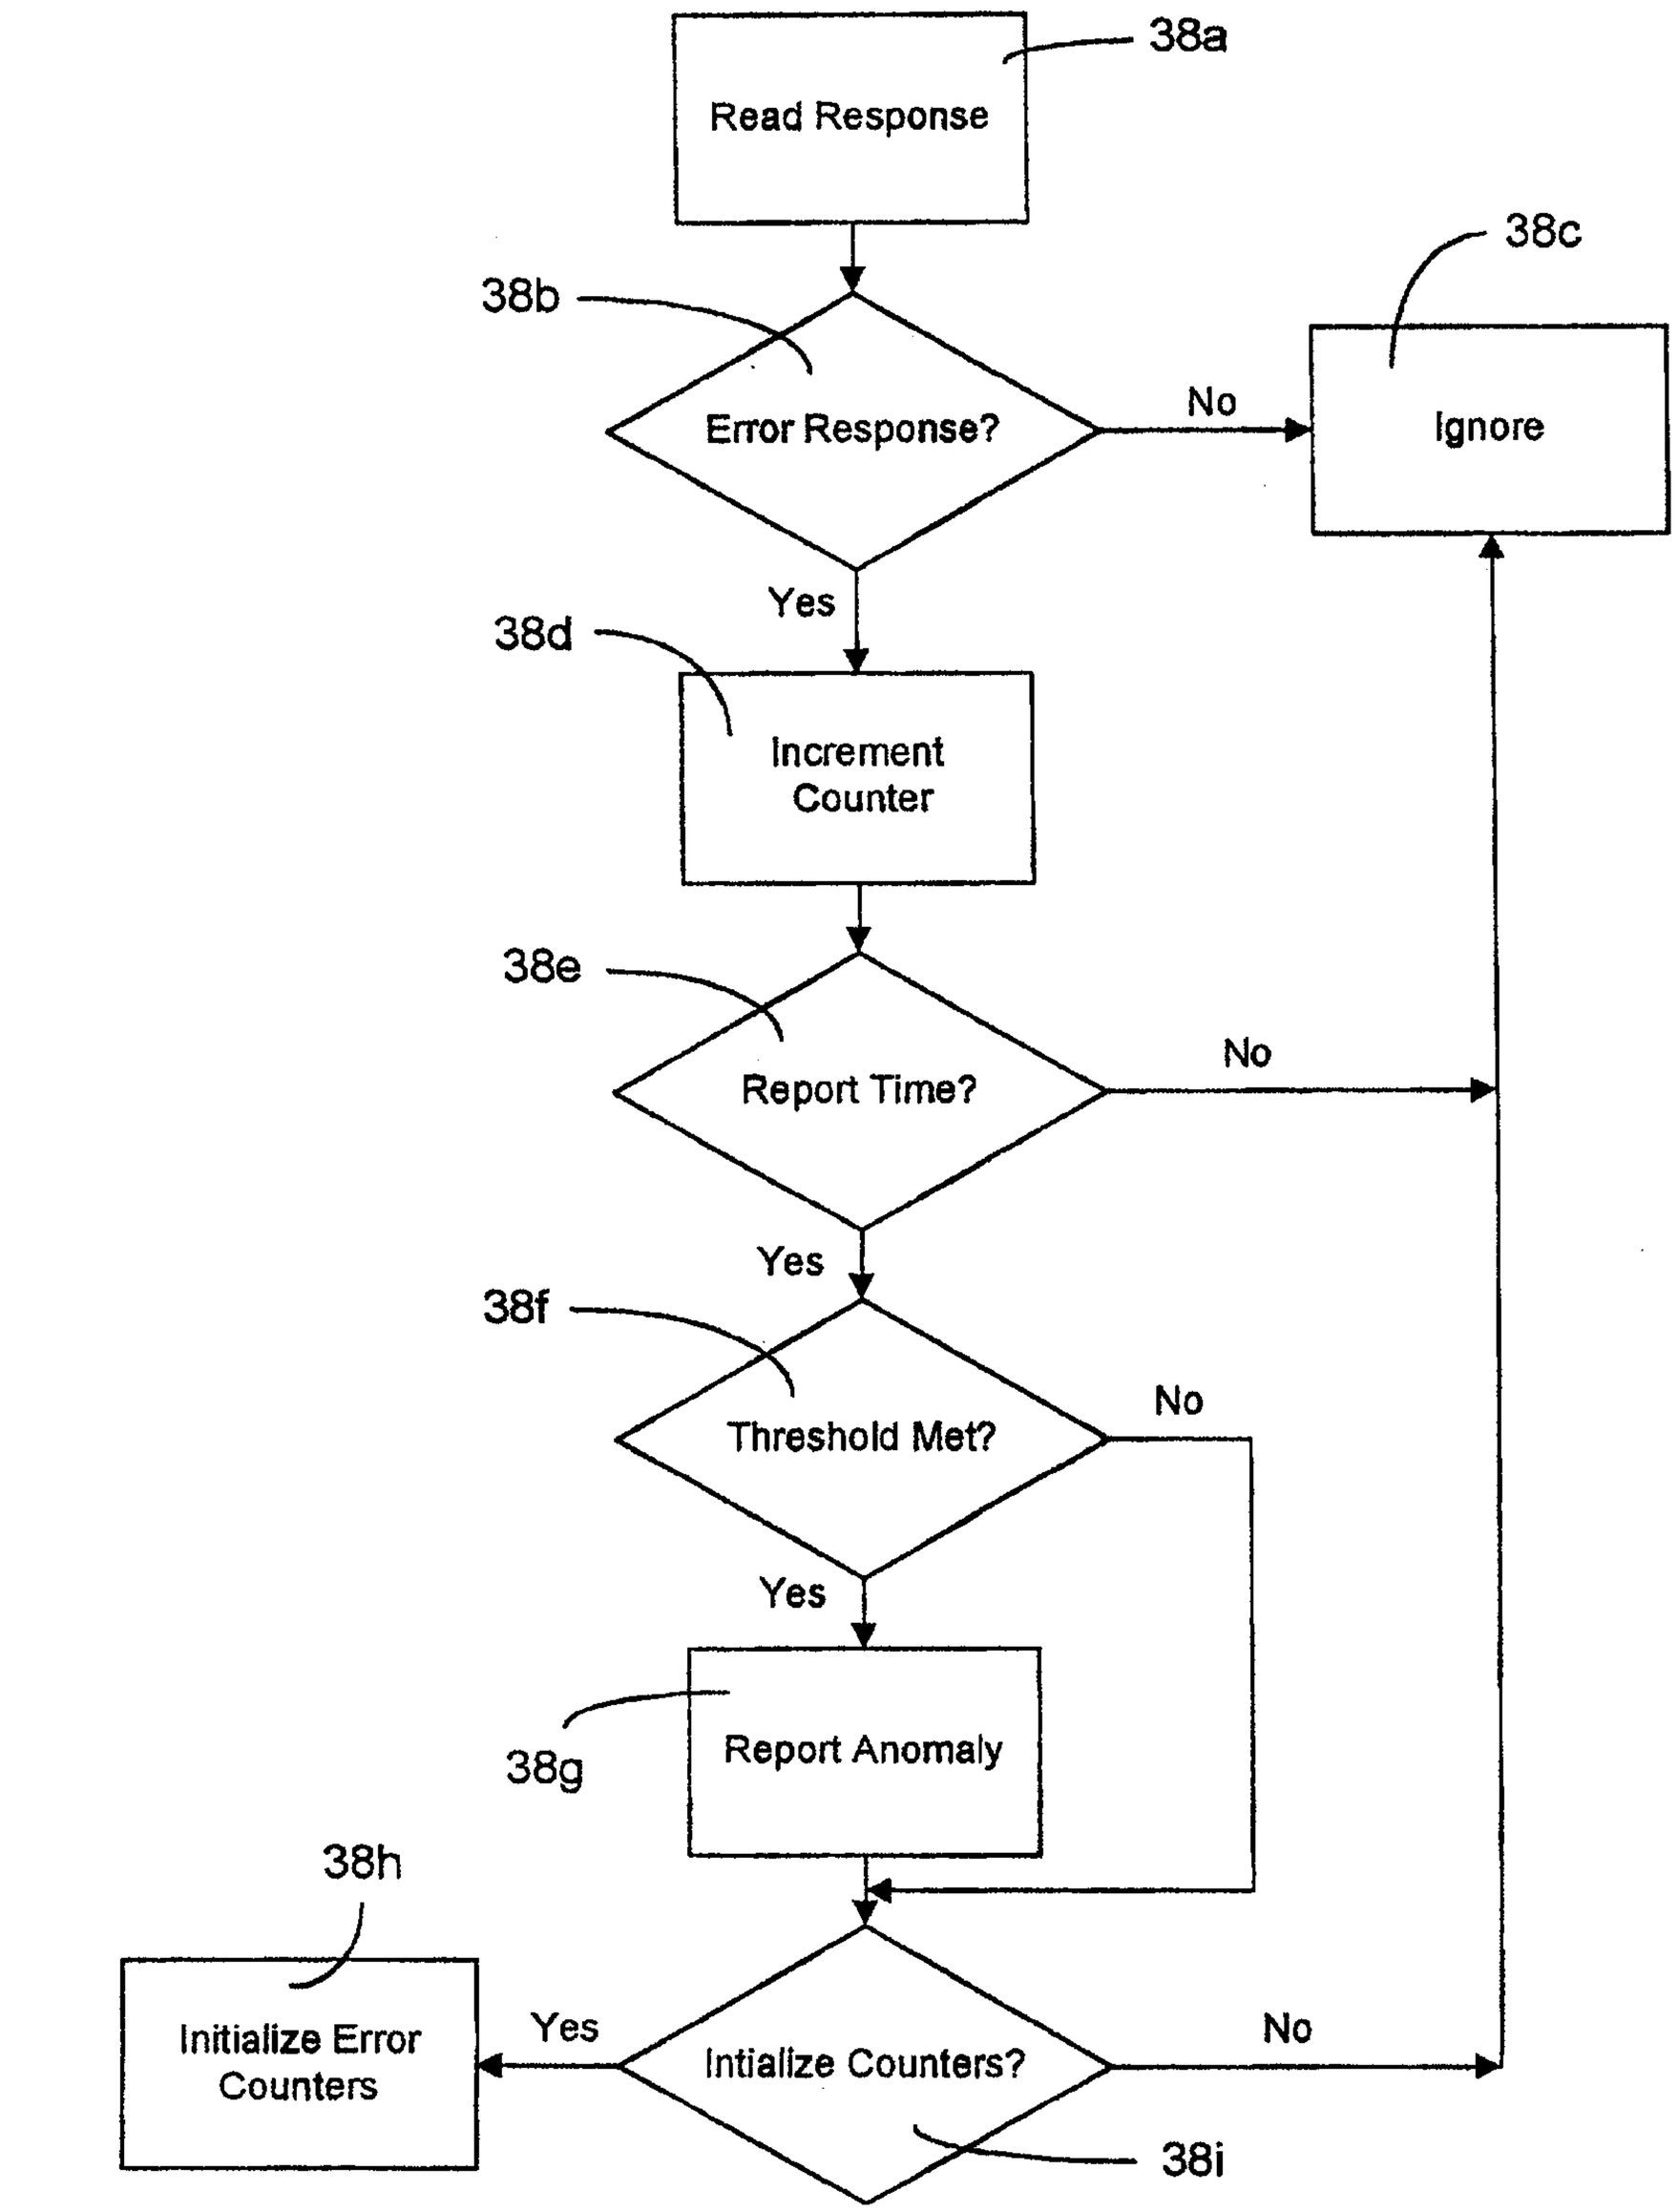
Read Response
38a
38b
Error Response?
No
38c
Ignore
Yes
38d
Increment Counter
38e
Report Time?
No
Yes
38f
Threshold Met?
No
Yes
38g
Report Anomaly
38h
Initialize Error Counters
Yes
Intialize Counters?
No
38i

FIG. 5

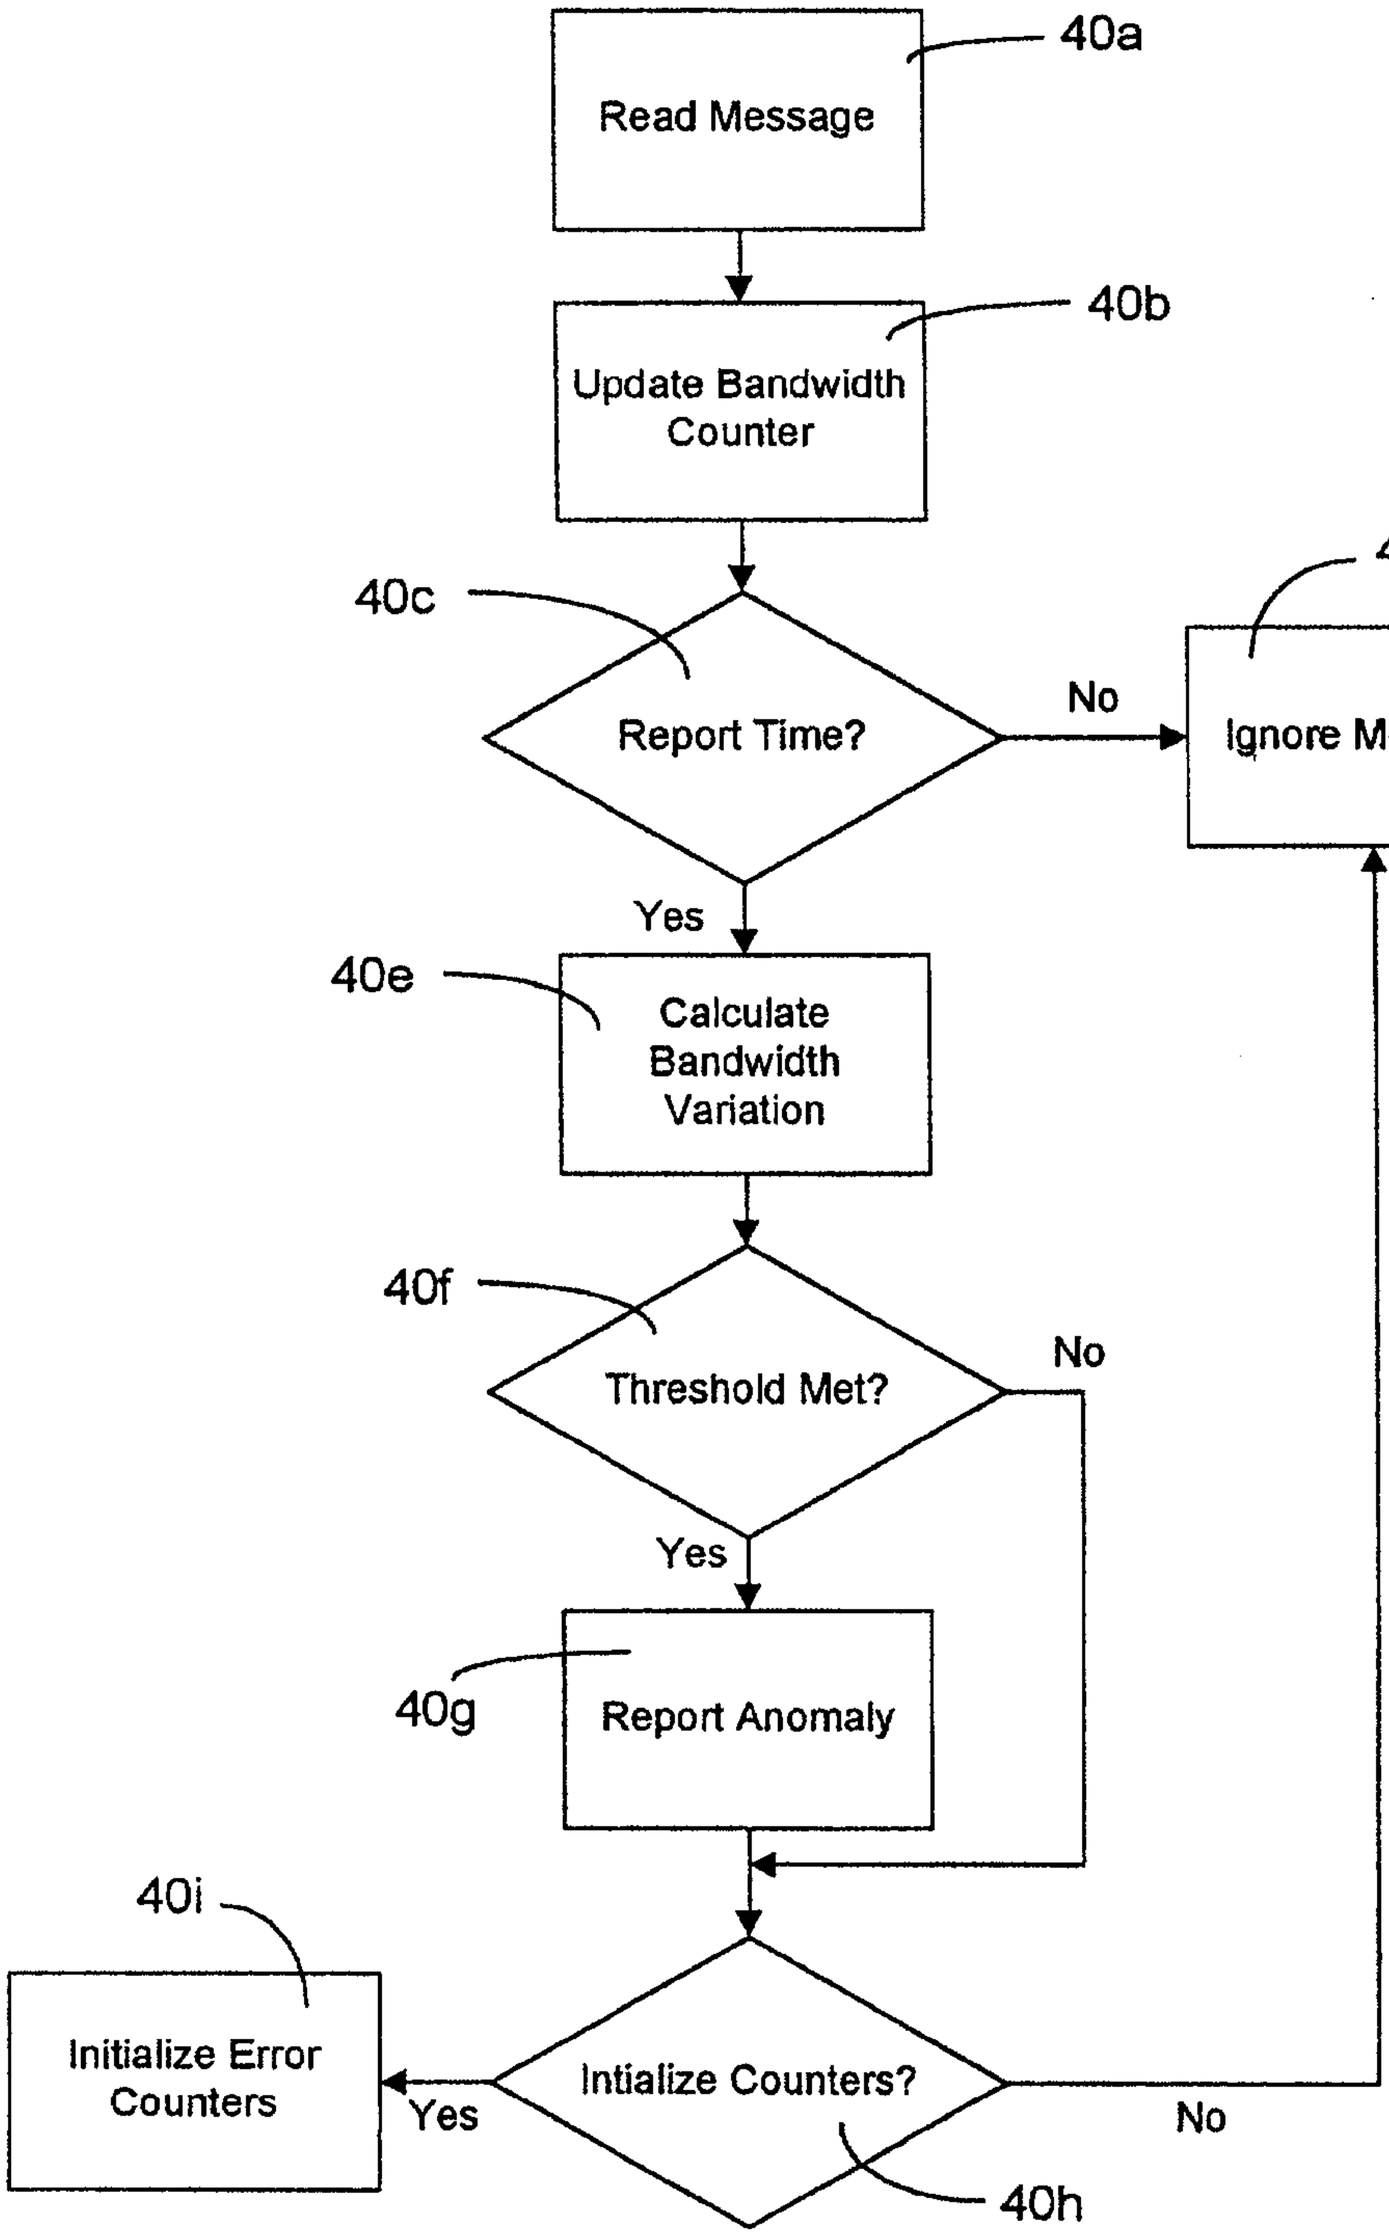
Read Message
40a
Update Bandwidth Counter
40b
40c
Report Time?
No
40d
Ignore Message
Yes
40e
Calculate Bandwidth Variation
40f
Threshold Met?
No
Yes
40g
Report Anomaly
40i
Initialize Error Counters
Yes
Intialize Counters?
No
40h

FIG. 6

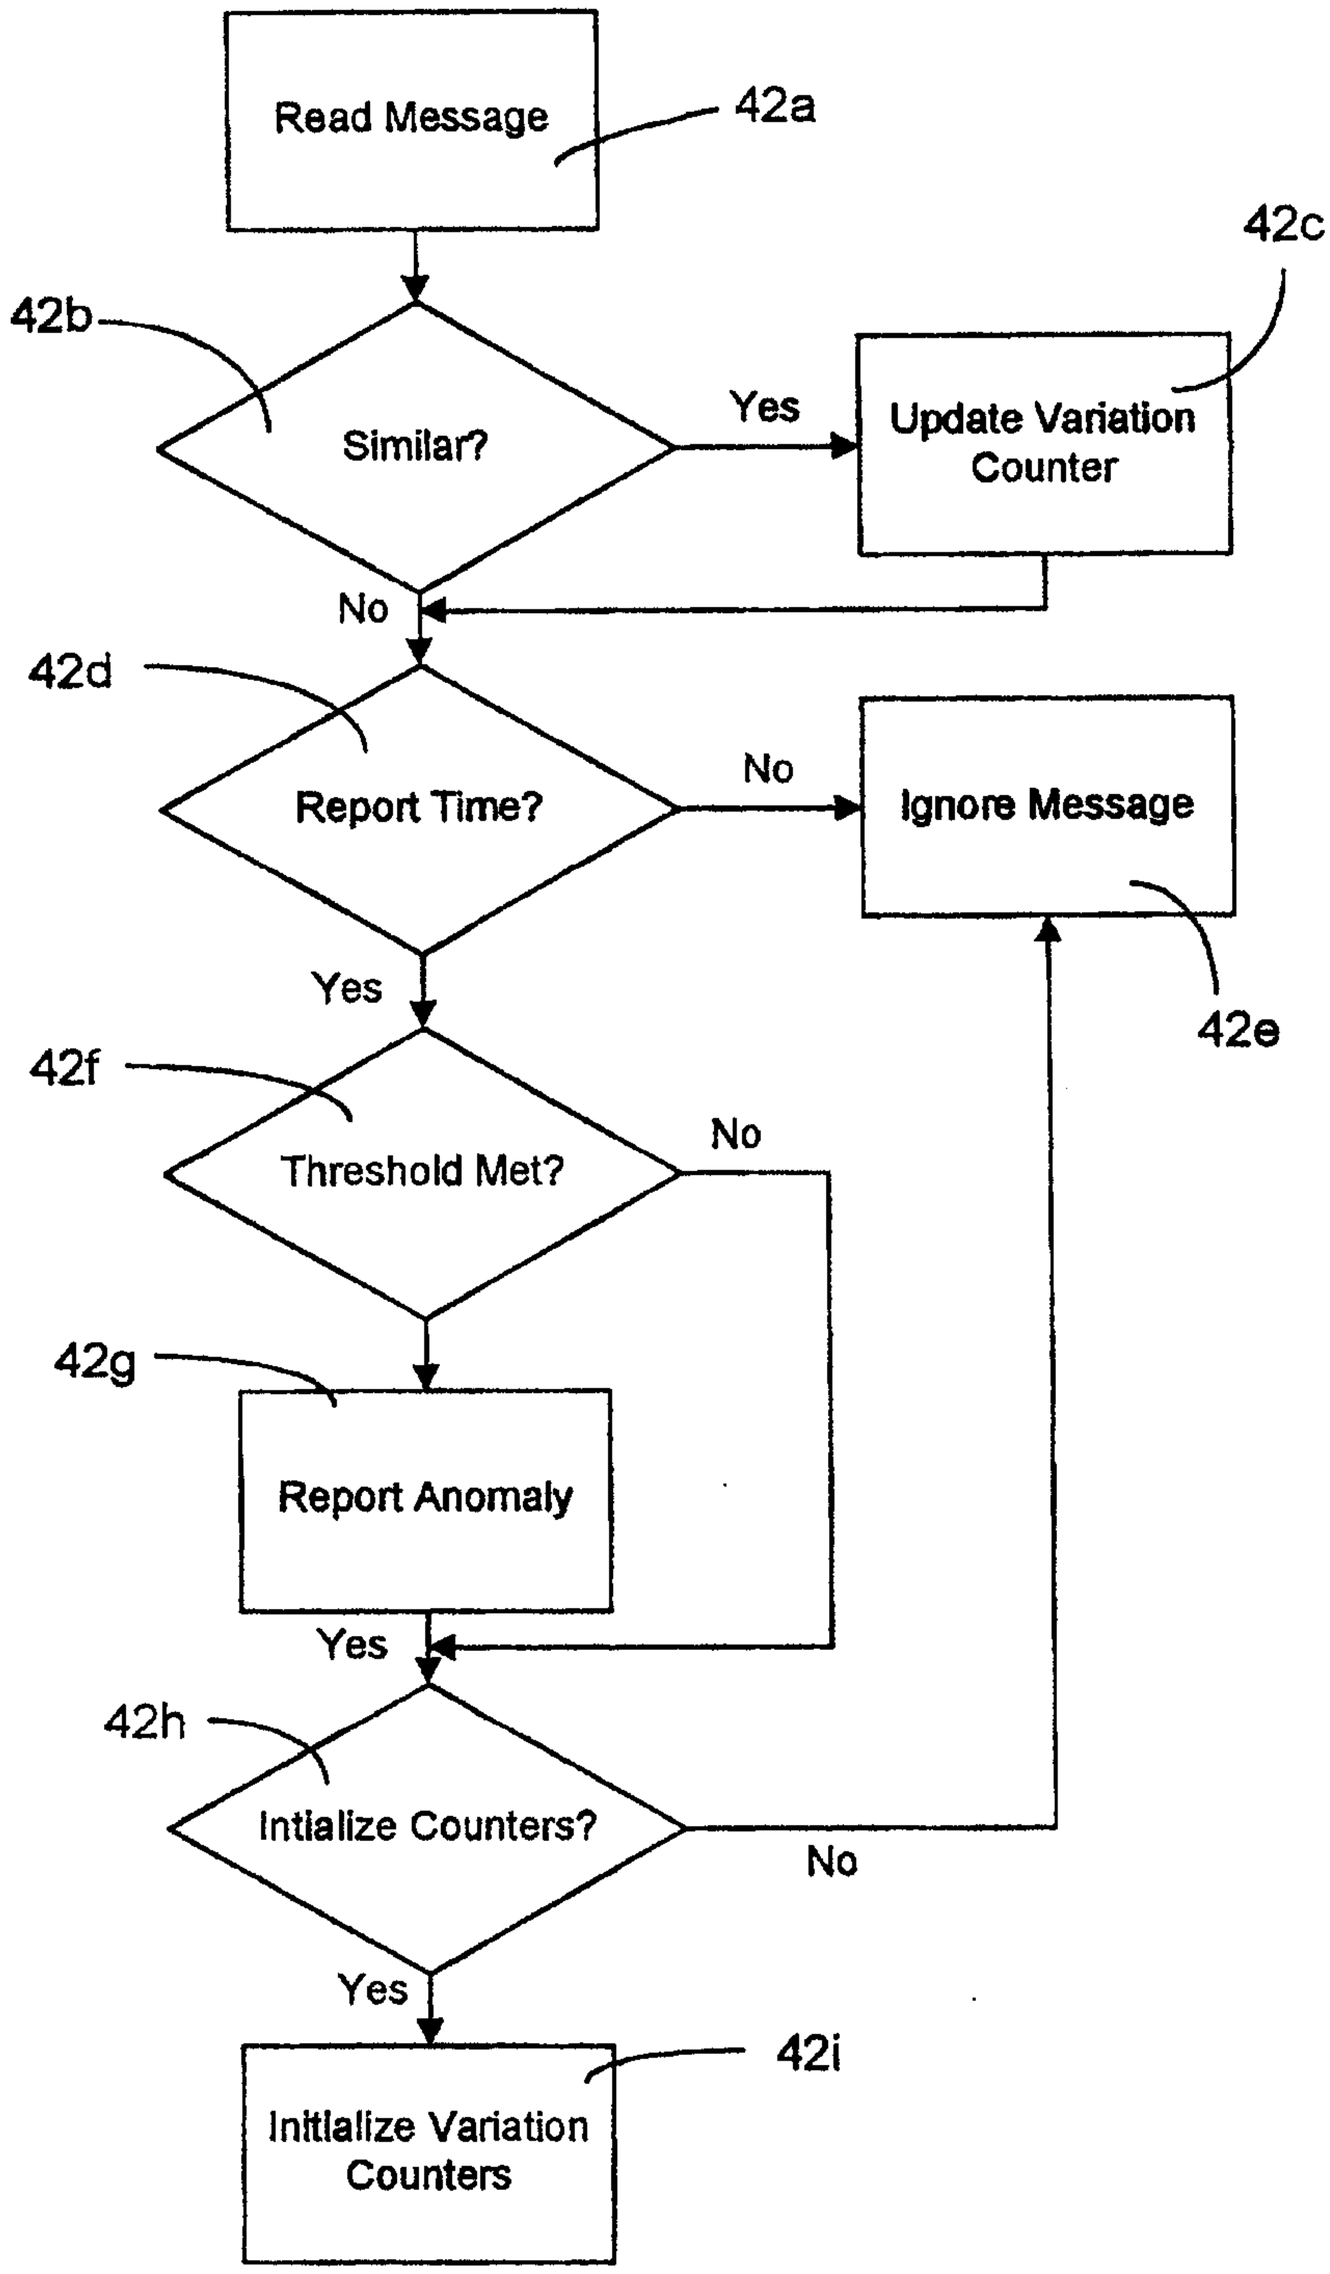
Read Message
42a
42b
Similar?
Yes
Update Variation Counter
42c
No
42d
Report Time?
No
Ignore Message
42e
Yes
42f
Threshold Met?
No
42g
Report Anomaly
Yes
42h
Intialize Counters?
No
Yes
Initialize Variation Counters
42i

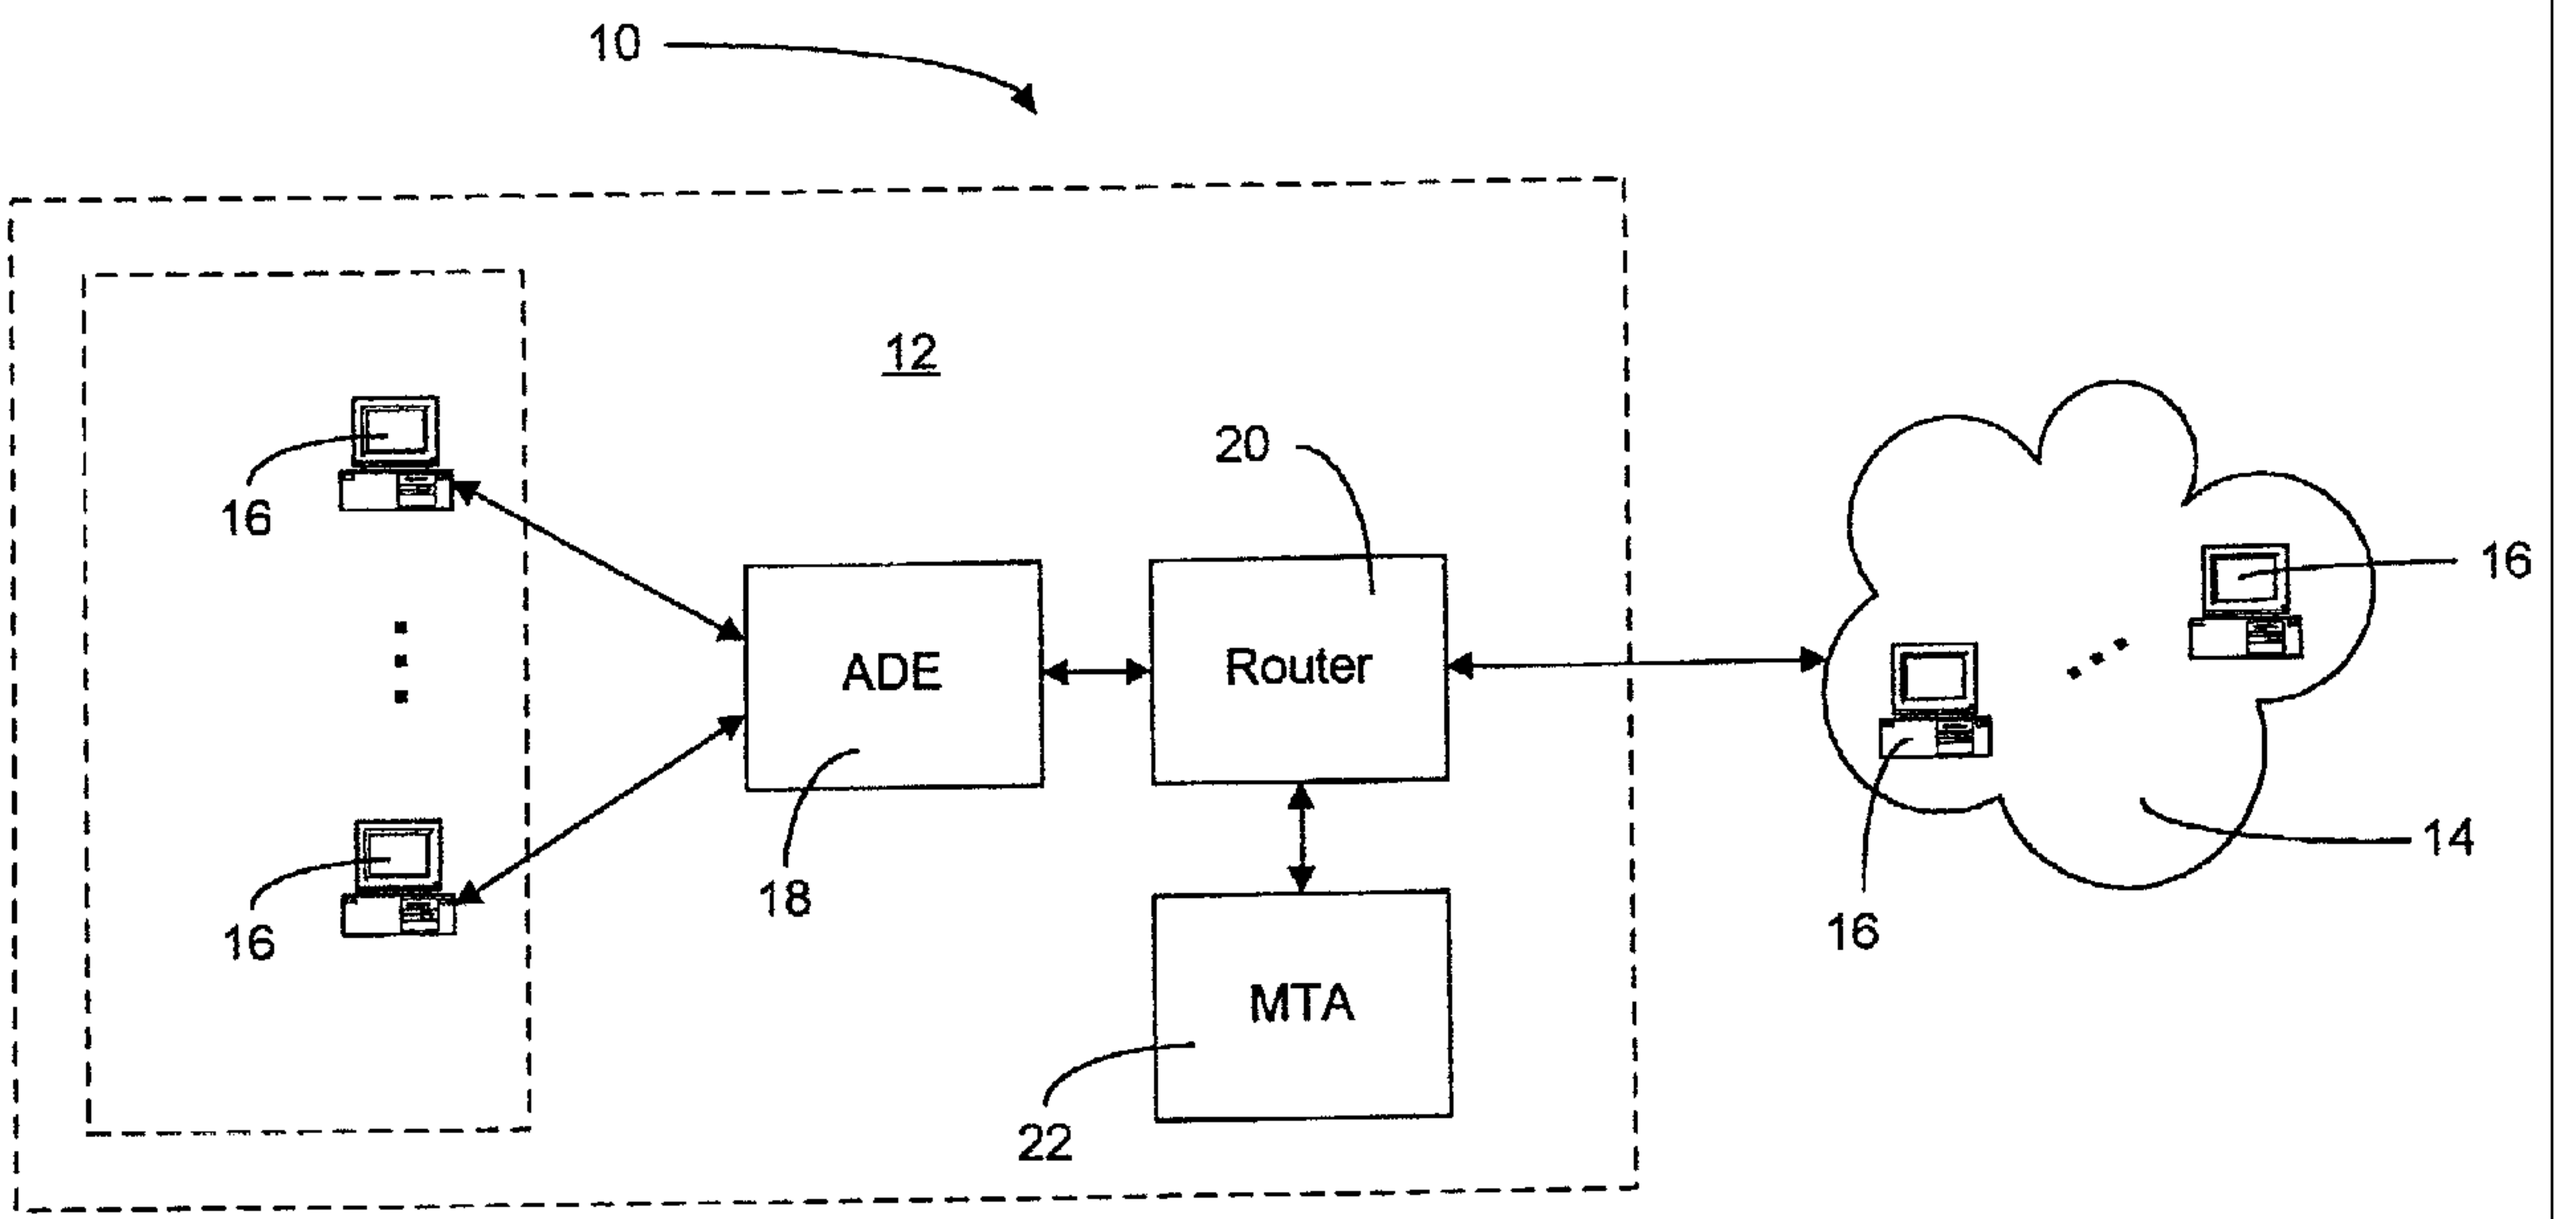
10
12
16
16
ADE
18
20
Router
MTA
22
16
16
14