



(22) Date de dépôt/Filing Date: 2001/11/08

(41) Mise à la disp. pub./Open to Public Insp.: 2002/05/16

(62) Demande originale/Original Application: 2 428 238

(30) Priorités/Priorities: 2000/11/08 (GB0027280.7);
2001/08/07 (US09/923,704)

(51) Cl.Int./Int.Cl. *H04L 9/00* (2006.01),
G06F 21/00 (2006.01), *H04L 29/02* (2006.01),
G06Q 30/00 (2006.01)


(71) Demandeur/Applicant:
ORCHESTRIA LIMITED, GB

(72) Inventeur/Inventor:
MALCOLM, PETER BRYAN, GB

(74) Agent: RIDOUT & MAYBEE LLP

(54) Titre : SYSTEME DE GESTION DE L'INFORMATION

(54) Title: AN INFORMATION MANAGEMENT SYSTEM



(57) Abrégé/Abstract:

An information management system is described comprising one or more workstations running applications which allow a user of the workstation to connect to a network, such as the Internet. Each application has an analyser, which monitors transmission data



(57) Abrégé(suite)/Abstract(continued):

that the application is about to transmit to the network or has just received from the network, and which determines an appropriate action to take regarding that data. The analyser may consult policy data containing a supervisor-defined policy to govern the workstation in order to determine what action to take. Such actions may be extracting data from the transmission data, such as passwords and usernames, digital certificates or eCommerce transaction details for storage and record keeping; ensuring that the transmission data is transmitted at an encryption strength appropriate to the contents of the transmission data; determining whether a check needs to be made as to whether a digital certificate received in transmission is valid; determining whether a transaction about to be made by a user of one of the workstations needs third party approval before it is made; and controlling the transmission of messages, such as e-mails according to a policy.

ABSTRACT

An information management system is described comprising one or more workstations running
5 applications which allow a user of the workstation to connect to a network, such as the Internet. Each application has an analyser, which monitors transmission data that the application is about to
10 transmit to the network or has just received from the network, and which determines an appropriate action to take regarding that data. The analyser may consult policy data containing a supervisor-defined policy to govern the workstation in order to determine what
15 action to take. Such actions may be extracting data from the transmission data, such as passwords and usernames, digital certificates or eCommerce transaction details for storage and record keeping; ensuring that the transmission data is transmitted at
20 an encryption strength appropriate to the contents of the transmission data; determining whether a check needs to be made as to whether a digital certificate received in transmission is valid; determining whether a transaction about to be made by a user of one of the workstations needs third party approval before it is
25 made; and controlling the transmission of messages, such as e-mails according to a policy.

- 1 -

AN INFORMATION MANAGEMENT SYSTEM

This application is a division of Canadian Patent Application No. 2,428,238 filed November 8, 2001.

This invention relates to the provision of
5 extended management functionality for Internet applications, particularly in the areas of information security, transaction auditing and reporting, centralized policy, and application connectivity.

Electronic commerce ("eCommerce"), particularly
10 between businesses ("B2B"), but also between business and consumers ("B2C"), is a fast growing market where buyers and sellers communicate using the Internet, a worldwide network of linked computer systems, instead of by traditional means such as mail, telephone and
15 personal meetings. Sellers advertise products and services using digital brochures and catalogues, which can be viewed or downloaded via an Internet connection, through pages on the World Wide Web, or via electronic marketplaces typically dealing in the goods and
20 services of a particular market sector. Buyers can find suppliers, select goods, obtain quotations, place and track orders, and even make payments entirely electronically and at any time. eCommerce brings the promise of increased flexibility, choice and
25 efficiency, with dramatically reduced procurement costs.

There are two universally accepted means of interfacing users to the Internet. The first of these is the 'Web Browser' which allows users to view pages
30 on the World Wide Web by accessing individual web sites, the addresses of which are typically widely published either using traditional means, or are referenced in another web site. The most widely adopted web browser is Microsoft Corporation's
35 "Internet Explorer".

- 2 -

The second means of interfacing is using an Electronic Mail program, with which the user composes a message, known as an e-mail, which is then electronically routed to the address of the intended
5 recipient over the Internet. Well known Electronic Mail programs include IBM Corporation's "Lotus Notes" and Microsoft Corporation's "Outlook".

In a typical eCommerce scenario, a buyer might identify a particular product, together with pricing
10 and delivery information, on the sellers' web site. He may then place an order, either by filling in an electronic order form on the web site, or by sending an e-mail directly to the seller. The order would typically include a commitment to payment, perhaps in
15 the form of Credit Card details, or by some electronic payment means. The seller would then typically send a return e-mail to confirm acceptance of the order.

Web Browsers operate in accordance with recognized standards, in particular Hyper Text Transfer Protocol
20 ("HTTP"), described fully in Internet standards document RFC2616. Electronic Mail programs operate in accordance with recognized standards, in particular Simple Mail Transfer Protocol ("SMTP"), described fully in Internet standards document RFC0821 and Multipurpose
25 Internet Mail Extensions ("MIME") described fully in Internet standards documents RFC2045-2049.

While eCommerce provides enormous benefits, its adoption raises many new issues, which must be addressed in order to ensure its continued adoption,
30 particularly if it is to ultimately replace traditional methods. One of the central issues is security.

The Internet is an open communications network, which is by definition insecure since anyone can use it. Means to secure sensitive information to be
35 exchanged over the Internet (for example in an eCommerce transaction) have been provided by the

- 3 -

adoption of secure transmission protocols and messaging. Secure point to point transmission protocols, used for example between a web Server and a web Browser, include the 'Secure Socket Layer' ("SSL"),
5 defined by Netscape Communications Corporation, and its successor 'Transport Layer Security' ("TLS") defined in Internet standards document RFC2246. Secure e-mail message standards include 'Secure Multipurpose Internet Mail Extensions' ("S/MIME") described fully in Internet
10 standards document RFC2633 and "Pretty Good Privacy" a public domain secure messaging system developed by Philip Zimmerman.

In order to control access to information on servers connected to the Internet, a system of
15 usernames and passwords has been widely adopted. For example, access to discounted price lists on a particular web server may be restricted to trade users who have previously been given a username and password allowing them access. Similarly, on-line information
20 services typically make extensive use of usernames and passwords to restrict access to those who have paid for the service. By providing each user with a unique username and a changeable password, the service can ensure that only paid subscribers can access the
25 system, and allow users to prevent access by others to their personal data stored by the service.

In eCommerce applications, a major problem is the issue of identity and trust. When a supplier receives an order via the Internet it is perfectly possible,
30 even likely, that he has no prior knowledge of the customer. The supplier must establish that the customer is a) who he says he is, in other words that he is not masquerading as someone else, and that b) he is to be trusted and will ultimately pay for the goods
35 or service to be supplied. These issues have been addressed in the B2C market principally by the use of

- 4 -

credit cards. The customer provides his credit card number and address with the order, which the supplier then verifies with the credit card company, and obtains authorization for the charge. The entire process is
5 typically carried out on-line without human intervention. This method is largely effective where a supplier ships goods to the cardholder address, since a potential thief would not only need to steal the cardholders details, but would also need to intercept
10 delivery of the goods. It is much less effective in the case of services where no physical delivery is involved.

Clearly, the use of credit cards in eCommerce, though widespread, is restricted to small-scale
15 transactions potentially involving amounts, say, up to \$10,000. For those transactions above such amounts (which in aggregate monetary terms far exceed those below them), a mutually trusted third party must be used to establish both identity and trust.

20 Central to establishing identity is the use of Digital Certificates. The customer can be issued with a Digital Certificate by a trusted third party, which is then used to electronically 'sign' communications. On receipt of a signed message, the recipient (in this
25 case the supplier) can positively establish a) the identity of the sender, b) that the message has not been altered, and c) that the sender cannot subsequently deny he sent the message. Recognized standards for Digital Certificates are described in ITU
30 document X.509, and their use in Internet communications in Internet standards documents RFC2312, RFC2459, RFC2510, RFC2511, RFC2527, RFC2560, RFC2585 and RFC2632.

Chargeable, Third party services, such as that
35 provided by Valicert Inc., can be used to verify that a

- 5 -

Digital Certificate has not been revoked, for example after the certificate has been compromised in some way.

Once authenticity of messages is established, the supplier can use another third party to establish trust, or the same third party can be used to establish both authenticity and trust. For example "Identrus", a consortium of the world's major banks, provide a system such that when a supplier receives a message signed with an Identrus issued Digital Certificate, he can independently verify that the customer is a valid account holder in good standing with a recognized bank. Ultimately the system is to be extended such that the bank will additionally warrant the transaction, thereby guaranteeing payment to the supplier. It will be appreciated that the terms 'customer' and 'supplier' can apply to any two parties engaged in Internet communication.

It can be seen that appropriate combinations of the systems described provide a secure foundation for use of the Internet and the services and functions available through it. However, there are a number of problems with conducting eCommerce using only these systems. These problems are discussed below.

In the secure transmission protocols and messaging referred to above, data is usually encrypted before transmission and decrypted by the intended recipient prior to viewing. Thus, should the data be intercepted during transmission, it will be safe from viewing by unauthorised third parties unless they know or can ascertain the secret encryption key of the encryption algorithm.

The encryption and decryption of data at each end of a secure link or message requires significant processing power. Additionally both transmitting and receiving parties must be in possession of the same encryption key of the encryption algorithm, at the same

- 6 -

cryptographic strength, in order for the system to operate successfully. This often presents a problem, for example where regulations for the import or export of data into or out of a computer system prohibit the
5 use of higher strength algorithms, forcing the link or message to be encrypted at a lower cryptographic strength, or preventing secure communications at all. Consequently, secure links and messaging are typically used only when necessary.

10 In the case of communications over the World Wide Web, the requirement to secure transmissions is determined and initiated by the web Server. If, for example, the server is about to transmit an order form for completion by the user it may initiate a secure
15 link such that the order information will be encrypted when transmitted back to the server. Similarly, once the order is complete the server may terminate the secure link and return to normal unencrypted communication.

20 Typically, the only indication the user has that a link has been secured is an icon (usually depicting a padlock), which appears in the browser window. Once the icon has appeared, the user can then typically interrogate the browser to determine the strength of
25 the encryption algorithm being used, and can decide whether or not to enter, and subsequently transmit sensitive information, such as his credit card and address details.

In practice however, users frequently do not check
30 that the link is secure, far less that it is of suitable cryptographic strength to protect the information being transmitted. In order to address this problem, e-mail applications such as Microsoft Corporation's "Outlook" provide the ability to encrypt
35 all e-mails by default.

- 7 -

The wide adoption of usernames and passwords has created a management problem for many Internet users due to the sheer number that need to be remembered, particularly when good security practice requires

5 passwords to be frequently changed. Similarly, users will often need to use a variety of different usernames since someone else may already have taken their 'favourite' at a given site. Facilities to remember, and to automatically complete username and password

10 fields on subsequent occasions, have been provided in web Browsers such as Microsoft Corporation's "Internet Explorer", and by add-on 'helper' utilities such as Gator.com's "Gator". These facilities typically maintain a file of usernames, passwords and the web

15 page to which each applies. These files are encrypted to ensure that only the appropriate user can access them. If such username and password files are lost or become unavailable, such as when the authorised user has forgotten the encryption key or can no longer be

20 contacted to provide it, or when the file is accidentally or maliciously lost, destroyed, or corrupted then access to Internet accounts and services may be lost, and each site must be approached individually to replace or recover the necessary

25 username and/or password. This can be a very expensive problem for corporations in terms of lost access and administrative time. Additionally, such remembered usernames and passwords are only available for use on the machine on which they were originally used. If the

30 user moves to another machine, or uses multiple machines then the stored usernames and passwords are unavailable to him from those other machines.

All businesses, and many individual users, have a legal obligation to maintain accurate records of the

35 transactions they undertake, but for eCommerce transactions this can prove difficult. Businesses must

- 8 -

keep records for auditing purposes, for example to prove the terms upon which goods were ordered in the event of a dispute. Such records are considerably more difficult to maintain in an eCommerce environment, requiring the user to retain, for example, copies of orders sent by e-mail, or to print out the web page receipt from a web site purchase. For the user, this is labour intensive and there is no guarantee that any such created records are complete or reliable.

One automated solution of keeping records of eCommerce transactions is provided by Max Manager Corporation's "Max Manager" application. Max Manager captures receipt pages at known web sites, extracts transaction information from those receipt pages, and then stores locally both the receipt page and the extracted transaction information on the machine on which the application is running. However, in order to operate, Max Manager must be supplied with the exact address and layout of the receipt page. Max Manager determines that an eCommerce transaction has taken place either by detecting the address of the receipt page, or by comparing the current page being viewed by a browser with the layout of the receipt page that it has been supplied with. Once it has identified a receipt page, the relevant transactions details are extracted from the receipt page by using the known layout of the page as a template for matching purposes.

A significant drawback with Max Manager is that it may only be used to extract data from those pages for which it has been supplied with details. Moreover, if the layout of the receipt page is changed then Max Manager cannot meaningfully extract any data from the page until it is supplied with a new template for the changed layout. Since web sites change frequently, Max Manager must be constantly updated to take account of such changes. This is impractical on a large scale and

- 9 -

inevitably leads to transactions being missed, or worse reported incorrectly.

Problems also stem from the fact that computer terminals are distributed, often resulting in terminals and users being located at different locations. In multi-user environments, user machines may be physically connected to each other, for example using a Local Area Network ("LAN"), which provides a gateway for connection to the Internet. They may also be connected to local servers such as Microsoft Corporation's "Exchange Server", which acts as a central collection and distribution point for e-mail messages, and Microsoft Corporation's "Proxy Server", which acts both as a cache to improve performance of frequently visited web Sites, and as a filter to prevent access to certain web Sites which may have been designated as undesirable. However, in so far as the exchange of information is concerned, except in the case of a message sent between two local users, each user operates entirely in isolation from others at the same location. This presents a significant management problem for corporate and other organizations, which have no means of centrally controlling employee activity and cannot benefit from the significant cost savings that might be made from the sharing of information. For example, two users in an organization may independently receive e-mail messages digitally signed by the same sender. Both recipients must separately validate the Digital Certificate, incurring two validation charges, at least one of which was unnecessary.

The present invention provides additional functionality to the systems mentioned above to alleviate their inherent problems and to provide a single integrated system for information exchange.

- 10 -

According to one aspect of the invention, an information management system comprises an information management system comprising a plurality of client workstations adapted for connection to a computer network, each workstation having a memory; a data repository arranged to receive data from each of said client workstations; an application stored in the memory of each workstation for transmitting outbound data to said network and receiving inbound data from the network; policy data defining rules for the recording of data that may comprise part of a transaction conducted between a client workstation and a third party across the computer network; an analyser, the analyser being operable in conjunction with the policy data to analyse at least one of the outbound data and the inbound data, to identify the existence of a transaction occurring between a client workstation and a third party by analysing the outbound or the inbound data, and to cause transaction data that is all or part of the outbound data or the inbound data related to an identified transaction to be stored in the data repository.

According to another aspect of the invention, a method of managing information comprises the steps of providing a plurality of client workstations adapted for connection to a computer network, each workstation having a memory; providing a data repository arranged to receive data from each of the client workstations; providing an application stored in the memory of each workstation for transmitting outbound data to the network and receiving inbound data from the network; providing policy data defining rules for the recording of data that may comprise part of a transaction conducted between a client workstation and a third party across the computer network; and analysing, at least one of the outbound data and the inbound data to

- 11 -

identify, with reference to the rules of the policy data, the existence of a transaction occurring between a client workstation and a third party; and storing transaction data that is all or part of said outbound
5 data or the inbound data related to an identified transaction in the data repository.

According to a further aspect of the invention, there is provided a computer program product for controlling a plurality of computers in a private
10 network to manage information, the network having a data repository arranged to receive data from the plurality of computers, and policy data defining rules for the recording of data that may comprise part of a transaction conducted between a computer in the private
15 network and a third party across a public network, the program product comprising a recording medium readable by a computer, having program code recorded thereon which when executed on each of the plurality of computers configures the computers to analyse, in
20 conjunction with an application running on the computer that is operable to transmit outbound data to the public network and receive inbound data from the public network, at least one of the outbound data and the inbound data to identify, with reference to the rules
25 of the policy data, the existence of a transaction occurring between the computer and a third party; and to control the computer to store transaction data that is all or part of the outbound data or the inbound data related to an identified transaction in the data
30 repository.

The information management system provides many advantages in the eCommerce environment to on-line trading companies, who may benefit by being able to

35

- 12 -

regulate the transactions made by their staff according to their instructions encoded in the policy data, automatically maintain records of passwords and business conducted on-line, avoid paying for
5 unnecessary checks on the validity of digital certificates, and ensure that the transmission of data by their staff is always carried out securely.

Preferred embodiments of the present information management system will now be described in more detail,
10 by way of example, and with reference to the drawings in which:

Figure 1 is a schematic illustration of the present arrangement of systems and resources making up the Internet according to the prior art;

15 **Figure 2** is schematic illustration of the preferred embodiment of the invention implemented in a corporate environment;

Figure 3 is a schematic illustration of the operation of a web browser in accordance with the
20 preferred embodiment of the invention;

Figure 4 is an illustration of a typical input window generated by a web browser;

Figure 5 is a schematic illustration of the operation of an e-mail client in accordance with the
25 preferred embodiment of the invention;

Figure 6 is a flowchart illustrating the operation of a plug-in module, according to a preferred embodiment of the invention, for capturing username and password values transmitted by a user to a remote web
30 site;

Figure 7 is an illustration of example policy data specifying control conditions for recording data;

Figure 8 is a flowchart illustrating the operation of a plug-in module for recognizing credit card numbers
35 contained in data transmitted to or from a web server or e-mail client;

- 13 -

Figure 9 is a flowchart illustrating the operation of a plug-in module for establishing the validity of a digital certificate received by a user;

Figure 10 is an illustration of example policy data for determining whether or not a digital certificate should be verified;

Figure 11 is a flowchart illustrating how the example policy data shown in Figure 10 is used to determine whether or not verification is required for a digital certificate;

Figure 12 is a flowchart illustrating the operation of a plug-in module, according to a preferred embodiment of the invention, for identifying transmissions from a user or to a user that comprise part of an eCommerce transaction;

Figure 13 is an illustration of example policy data intended to be used with the process illustrated in Figure 12 to identify a transaction;

Figure 14 is a flowchart illustrating the operation of a plug-in module, according to a preferred embodiment of the invention, for recording transmissions identified as comprising part of a single transaction thereby forming a record of the transaction;

Figure 15 is a flowchart illustrating the operation of a plug-in module for approving or rejecting identified transactions on the basis of a predetermined policy setting; and

Figure 16 is an illustration of example policy data for determining whether an identified transaction requires approval, and for identifying an appropriate approver;

Figure 17 is a flowchart illustrating the operation of a plug-in module for determining an appropriate level of encryption for a transmission and

- 14 -

allowing that transmission to be transmitted only if that level is provided;

Figure 18 is an illustration of example policy data specifying the required encryption strength for various data types;

Figure 19 is an illustration of example policy data for controlling the re-direction of out-bound messages;

Figure 20 is a flowchart illustrating the operation of a plug-in module for re-directing out-bound messages to a third party for review before transmission, utilizing the policy data shown in Figure 19;

Figure 21 is a flowchart illustrating the operation of a plug-in module for controlling the upload of information to a site external to the company, utilizing the policy data shown in Figure 19;

Figure 22 is an illustration of example policy data for controlling the forwarding of messages to recipients inside or outside the company;

Figure 23 is a flowchart illustrating the operation of a plug-in module using the policy data shown in Figure 22;

Figure 24 is an illustration of example policy data for controlling whether or not an out-bound message is digitally signed; and

Figure 25 is a flowchart illustrating the operation of a plug-in module utilizing the policy data shown in Figure 24.

The preferred system described herein provides users of the Internet with an automatic way of managing the flow of information on a computer system. It provides facilities to manage the security level at which transmissions occur, facilities to record on-line transactions and to refer transactions that are about to be made to third parties for approval, and means to

- 15 -

stop transactions from occurring if approval is refused; it also provides facilities to extract and record pertinent data from any transmissions that are received or are about to be transmitted, and to
5 intelligently manage the transmission of e-mails.

The preferred system provides solutions for many of the problems encountered by eCommerce companies trading over the Internet; consequently the following exemplary discussion will be directed mostly to the
10 implementation and use of the system by a company of reasonable size conducting at least some of its business over the Internet. It will be appreciated that anyone however, including companies of any size or description and private individuals, who use the
15 Internet may benefit from the functionality provided by the preferred system.

The functionality of the preferred system is implemented through modules of code which are 'plugged-in' to the web browser or e-mail client. These 'plug-in' modules may be used to control and alter the
20 behaviour of the web browser or e-mail client in operation.

Many existing web browsers and e-mail clients may already be readily integrated with such plug-in
25 modules. In the case of Microsoft's Internet Explorer, the plug-in is known as a 'Browser Helper', and is more fully described in the document "Browser Helper Objects: The Browser the Way You Want It" by Dino Esposito, published by Microsoft Corporation in January
30 1999. In the case of Microsoft's Outlook and Exchange e-mail Clients, the plug-in is known as an 'Extension' and is more fully described in the document "Microsoft Outlook and Exchange Client Extensions" by Sharon Lloyd, published by Microsoft Corporation in March
35 1998. The use of the 'Browser Helper Object' and

- 16 -

'Extension' plug-ins made in the preferred system will be described in more detail later.

The use of browser or e-mail client plug-in modules to implement the functionality of the preferred system has the additional advantage that, since
5 encryption of message content is usually carried out by the browser or e-mail client itself, examination of transmission content, to extract password information or to determine the desired level of encryption for
10 example, may take place before the content has been encrypted ready for transmission, or indeed after it has been received and unencrypted.

Figure 1 shows the relationship between service providers, typically companies selling goods and
15 services over the Internet 10, and users wishing to purchase such goods or services. Users equipped with web browsers 22, 24 and 26, can connect via the Internet and retrieve web page information from web servers 14 and 18. Alternately users with e-mail
20 applications 20, 30 and 32, can send and receive e-mail messages with abc.com and xyz.com via e-mail servers 12 and 16.

In a corporate set up, such as that which is illustrated in the bottom right hand corner of Figure
25 1, web browsers 24 and 26 of a corporate user are connected to the Internet via proxy server 28. The proxy server 28 is used to cache web pages and control access to web sites. Similarly, the corporation has e-mail clients 30 and 32, connected to the Internet via
30 e-mail server 34 which acts as a central collection point for e-mails coming into the corporation and which controls distribution of the e-mails to individual users. It will be appreciated that while Figure 1 describes abc.com and xyz.com as sellers, a corporation
35 can be both a buyer and a seller, and as buyers abc.com

- 17 -

and xyz.com would be described as corporate users for the purpose of this description.

In the case of e-mails sent and received by personal e-mail application 20, it should be noted that
5 the mail will typically be collected and distributed by a remote e-mail server provided by the supplier of the Internet connection service to which the personal user subscribes.

While many of the features and functions of this
10 system provide considerable benefit to an individual user, the system provides the maximum advantage when operating in a multi-user environment where transaction information is gathered from many users. Figure 2 shows a schematic diagram of the preferred
15 configuration of the system in a multi-user environment. The preferred system comprises a Central Management server 40 connected to a database 42 and operator consoles 44. The Central Management Server 40 is also connected to Back Office Application Plug-ins
20 comprising Application Interfaces 50, 52 and open Application Program Interface 54, and to Gateway components 60, 62 and 64. Gateway component 62 is shown as connected to User Application Plug-ins, located on one or more user's machines, comprising Internet
25 Explorer Plug-in 70, Netscape Navigator Plug-in 72, Microsoft Outlook Plug-in 74, and Lotus Notes Plug-in 76. These plug-ins are used provide the functionality of the preferred system in the hosting program in which they are integrated. Four possible hosting programs
30 are shown, Internet Explorer, Netscape Navigator, Microsoft Outlook and Lotus Notes, but any other program with the capability to connect to the Internet may also be used, providing its behaviour can be modified to implement the functionality of the
35 preferred system.

- 18 -

Connection to the Internet 10 is made via the User application plug-ins and their respective hosting programs.

5 The gateway components 70, 72 and 74 are optional but are preferred as they enable the entire system to scale, with each gateway storing and forwarding information, thereby allowing any number of users to be connected.

10 Information from the multiple application Plug-Ins 70, 72, 74 and 76 for the different applications on multiple user machines is gathered by the central management server 40 and stored in an associated database 42.

15 The Back Office Application Plug-Ins 50, 52 and 54 allow the system to interface with third party management applications such as order processing and accountancy systems. This allows transaction information to be automatically entered and processed by such systems.

20 Operator Consoles 44 are provided for administrative purposes, and in particular for the approval of transactions. While logically depicted as directly attached to the central management server in Figure 2, such consoles could be run on any networked machine. Where an e-mail or web browser plug-in
25 determines that a particular transaction requires approval, a request is sent to the Central Management Server and queued pending approval by an authorized operator.

30 The operation of the system is controlled by policy data, which stores the corporation's regulations regarding security, authorisation, and the actions that users are permitted to perform, as well as operating information. Preferably, the policy data is stored in
35 a policy file on the central Management Server for access by any of the Operator Consoles 44, Back Office

- 19 -

Application Plug-ins or User Application Plug-ins. The system administrator or network supervisor may define one or more policies or settings of the policy file and may assign individual users or groups of users to
5 different policies, thus controlling a user's ability or even a workstation's ability to interact with the Internet without the need to set parameters and controls directly at each user's machine. A user in the accounts department of a company for example may be
10 assigned to an 'accounts policy'; any subsequent change to that policy will then automatically result in a change in the capabilities of all users assigned to that policy.

It is preferred that the capability to edit or set
15 the policy data is restricted to the network supervisor or other authorised person or persons. This may be achieved by designating one or more supervisor workstations in the network enabled with access to edit the policy data such as Operator Consoles 44.

20 Preferably, the policy is tree-like in structure, allowing settings to be forced down to individual policy nodes of the tree, and global changes to be rapidly made, for example if the CEO wishes to cause all purchases to require his approval if company cash
25 flow should become a problem. Such a policy based system greatly reduces the latency inherent in both traditional purchasing systems and in current eCommerce purchasing environments.

Each user of the network will have his or her own
30 representation of policy data. Preferably, only the branches and leaves of each user's policy that differ from a master network policy are stored as this allows space in memory to be saved. Although the policy data is preferably stored in file form on the Central
35 Management Server, it is not intended that storage of the policy data be restricted to file form only. Any

- 20 -

other representation or encoding of policy settings may be employed within the preferred system.

The implementation of the system in a web browser or in an e-mail client will now be described in more
5 detail.

Use of the preferred system in a web browser

Figure 3 shows the simplified operation of a web browser. The web browser is launched at step S100 in response to a start request from either the user or
10 automatically from the start-up file of the user's computer. The start-up file contains commands to automatically run specified programs when the computer is booted-up. After the web browser has been started it typically requests a 'home page', the default web
15 page for viewing, in accordance with a pre-determined setting. This is shown at step S102.

The request is sent to the appropriate web server 90, the exact Internet address of which usually being determined by Domain Name Services; the web server 90
20 then replies with the appropriate data defining the web page. This process is represented respectively as steps S104 and S106 which result in step S108.

The data defining the web page consists of HTML script, and other possible data types such as XML or
25 ActiveX, and Javascript which encodes executable programs. The Browser interprets this data, displaying and/or executing it as appropriate at step S110.

The browser then typically waits for user input at step S112. Such input may include filling displayed
30 fields, clicking on a hyper-link, or entering the URL address of a new web page. Ultimately, such actions lead to a further request being sent to the web server 90 at step S114 and step S116. The request may simply be another web page address, or it may contain
35 additional data such as that typed in to displayed fields by the user.

- 21 -

Figure 4 shows a sample web page display, in which a GUI is presented to the user in order to receive the users name and e-mail address. It will be seen from reference to Figure 4 that the user has entered his
5 name as 'Fred Smith' into the name request field provided, and his e-mail address as 'fsmith@xyz.com' into the e-mail address field.

When the user clicks the 'Submit' button provided on the request window, the details the user entered are
10 included in the command sent to the web server 90. Such a command might be:

```
http://www.sample.com/sample2.htm?UserID=Fred  
+Smith&email=fsmith@xyz.com&submit=submit  
15
```

It can be seen from the above that the user's name is incorporated into the command as the value of a variable called 'UserID' and his e-mail address is incorporated as the value of a variable called 'email'.
20

The command is assembled in step S114, and transmitted to the web server 90 in step S116 where the user name and e-mail address information may be used, for example, to send product information to the user via e-mail, or to gain access to other web pages.

25 The plug-in module provided by the preferred embodiment of the invention in the form of a Browser Helper Object (BHO) provides additional functionality to augment that of the standard web browser. The BHO is implemented to respond to a number of significant
30 events that occur as the web browser is operated and directed by the user to interact with various web sites and pages.

The BHO is implemented to monitor navigation requests and data submitted to the web server from the
35 browser and identify data that is unique to the user. It may do this simply by searching the outbound data

- 22 -

stream for the presence of pre-determined words or phrases. In the above case shown in Figure 4, the two variable definitions 'UserID' and 'email' may be searched for, and the data that follows them extracted
5 and stored. Alternatively, the BHO may search for the '?' symbol, which indicates the end of the URL address being connected to and indicates that what follows is data. The BHO may also monitor the inbound data stream received from the web site being connected to.

10 Also, the BHO may be implemented to monitor the operation of the web browser itself. As the web browser operates it generates 'events' to notify co-dependent software modules or objects that something significant has just occurred or that an action has
15 just been completed. The name of the event is usually descriptive in its own right of what has just occurred; additional data that describes the event in more detail is normally available. The BHO is implemented to trap these events and to take action in dependence on them.

20 One such event that the BHO is implemented to respond to is called 'BeforeNavigate2' which the web browser fires when the user requests the browser navigates to a new page. The event is issued and may be recognised by the BHO before the requested page is
25 downloaded, allowing the BHO to take any pertinent action before the user views the page. One such action might be to record the page and any data submitted in response to that page in a database. Another such action might be to identify the URL of the requested
30 page from the event and prevent the page being downloaded.

Another event that the BHO traps is the 'DocumentComplete' event, which is fired by the web browser when a new page has been fully downloaded from
35 the web site into memory. The page is encoded in the form of a Document Object, conforming to Microsoft's

- 23 -

Document Object Model (DOM). The DOM provides comprehensive access to the data comprising the page, allowing the BHO to extract data items which are of interest to it. For example, the BHO may request data
5 from the DOM to determine whether the page forms part of an eCommerce transaction. It may do this by searching objects in the DOM for terms such as 'Receipt' or 'Account Number'.

The BHO may also use the DOM to determine the
10 field names or field types of data being requested on a web page. The data entered by the user into such fields may then be extracted from the DOM and stored or acted upon. Field names are typically descriptive of what is stored; passwords, for example, are often held
15 in a field called 'password' and so may be searched for on a web page. Credit card numbers may be searched for in a similar way. Usually, password fields are of a type such that any entered data is displayed as asterisks. This too may be determined from analysis of
20 the DOM and used to identify pertinent data.

User data would not normally be present in a web page downloaded from a web site, but would be entered by the user into an HTML form. Usually, the potentially sensitive user data is transmitted to the
25 web site via the web server when the user selects a 'Submit' button. At this stage, the BHO can trap the 'Submit' event issued by the web browser, and access the DOM to extract the user data, and, if necessary, prevent the data from being transmitted.

30 Encryption and decryption on a secure link will take place after point C and before point A in Figure 3 respectively. Thus, the BHO may analyse the data before it is encrypted or after it is decrypted. This is advantageous since there is no need for the BHO to
35 perform any encoding or decoding of data itself. This does not affect the ability to determine if the link is

- 24 -

secure or not, since a secure link can be identified by the protocol identifier "https" at the beginning of the current URL. It is preferred that examination of the transmission's content take place before encryption or
5 after decryption occurs.

Discussion of the operation of an e-mail client

With reference to Figure 5 of the drawings, the operation of a typical e-mail client, and the
10 implementation of the preferred embodiment in an e-mail client will now be described.

Figure 5 shows the simplified operation of an e-mail client. Receive and Send operations typically operate independently, and these operations are shown
15 separately on opposite sides of Figure 5, beginning at steps S120 and steps S130 respectively.

An e-mail client's 'receive message' operation is initiated at step S120. This may be done automatically at predetermined intervals in order to keep the user
20 informed of any new messages that he has received, or it may be done in response to the user manually selecting a 'receive messages' icon. Starting this operation causes the e-mail client to poll the e-mail server 95 and download any new messages to the user's
25 machine. In step S122, an e-mail message is received by the e-mail client. Typically, when a new message is received, it is added to an 'Inbox', with the received message headers (senders name, date and title for example) arranged in a list. The user then clicks on
30 the appropriate entry in the list to read the full message causing it to be displayed on his computer screen. The e-mail message is displayed at step S124.

In the case of an outgoing e-mail, the user selects a 'compose e-mail' option as step S130. In
35 response, the e-mail client provides an interface comprising a text editor in which the user can enter

- 25 -

the text of the body of the message and other information such as destination address, subject and so on. The user composes the message in step S132 and then opts to send it, by selecting an icon or menu option provided by the e-mail client to issue a 'send command'. The e-mail is sent to the e-mail server for transmission to the recipient in step S134. If any encryption is applied by the e-mail client it is applied in step S134 before transmission.

10 In the preferred embodiment, additional functionality is provided for the e-mail client via a plug-in module. Preferably, the e-mail client is one of those provided by Microsoft, such as the Microsoft Exchange client, or the Microsoft Outlook client, and 15 the plug-in module is encoded as an Exchange client extension. These are described in the document "Microsoft Outlook and Exchange Client Extensions" by Sharon Lloyd mentioned earlier.

An Exchange Client Extension is a component object 20 that complies with the Microsoft Windows Component Object Model (COM) and that utilizes the Exchange IExchExt interface. This interface provides a number of additional interfaces for modifying the operation of the Exchange e-mail client, such as the 25 IExchExtCommands interface, which allows existing client behaviour to be replaced or modified and new commands to be added to the client's menus; and the IExchExtEvents interface which allows custom behaviour to be implemented to handle client 'events' such as the 30 arrival of new messages, reading, writing, sending messages and reading and writing attached files. The IExchExtMessageEvents, IExchExtSessionEvents and the IExchExtAttachmentEvents interfaces are also provided and provide additional functionality for the more 35 specific tasks that each of the interfaces names suggest.

- 26 -

In the preferred embodiment, the Exchange Client Extension that forms the plug-in module is implemented to respond to client 'events' that are fired by the client program as it performs operations and completes
5 actions. The 'events' in question are provided by the COM interfaces mentioned above. The monitoring of the e-mail client by the plug-in module may therefore be seen to be analogous to the way in which the BHO plug-in module monitors the operation of the web browser.

10 The e-mail client plug-in module is implemented to respond to the 'OnDelivery' event for example which is fired when a new message is received from the underlying mail delivery system and before it is visible to the user. The 'OnDelivery' event contains
15 information to access the different parts of the e-mail message that have been downloaded and which are held in memory. The message header, the message body and any message attachments are encoded in memory as properties of the message object which may be separately accessed
20 through Mail Application Program Interface (MAPI) calls.

Via the information supplied as part of the 'OnDelivery' event, the plug-in module may access the message header and extract the identity of the sender
25 for example. Furthermore, the plug-in module may utilise information obtained from MAPI calls to scan the body of a received message for keywords or pertinent data. It may search for evidence of an eCommerce transaction, by identifying significant words
30 such as 'receipt' or 'account number'. The message may then be stored for auditing purposes. In the case of an unapproved sender, or harmful message content, the message may be deleted unseen.

The analysis of a received e-mail occurs therefore
35 at point A in Figure 5 prior to being viewed by the user. Preferably the e-mail is examined before the e-

- 27 -

mail is even placed in the Inbox. Where a message is not automatically decrypted before being placed in the Inbox, for example where the user is required to enter a decryption key, the message is examined immediately following decryption, but before viewing. Digital Certificates may be included as attachments to the e-mail and can readily be examined prior to viewing, allowing any appropriate actions, such as validation, to be performed.

Another significant client event that the plug-in module is implemented to respond to is the 'OnWriteComplete' event which is fired when the user has selected the 'send command' and requested the e-mail client to transmit a new e-mail message to the mail delivery system. This event is fired, at point B in Figure 5, before the transmission and before any encryption takes place. The new message which is to be transmitted is similarly stored in memory as an object which may be accessed by MAPI calls. The plug-in module may use the MAPI calls to scan the content of the outgoing e-mail for sensitive data, such as credit card number, and consequently cause the message to be recorded or even blocked.

Plug-In Module Operation

The preferred implementation of the web browser and e-mail client plug-in modules has been described above with reference to Figures 3 and 5 above. Next, the functionality provided by the plug-in modules will be described in detail and with reference to Figures 6 to 18.

Identifying and Recording Usernames, Passwords and other Information

The preferred system provides means to automatically identify, collect and store data

- 28 -

contained in transmissions to and from a user's workstation, in particular, the usernames and passwords entered by a user to access web site pages, File Transfer Protocol ('FTP') sites and other such sites on the Internet.

Systems that presently provide facilities to record passwords currently only do so when a user clicks on the 'remember password' option provided on the GUI. The password is saved into a protected local file on the user's machine that is only opened when the user is authenticated on that machine, for example, by entering his username and password at boot time. The remember password option causes the system to remember the user's password when he next visits, pre-filling in the password field with that password so that the user does not have to enter it each time it is required. The drawback of the password file being stored locally is that if the user moves to another machine he does not have access to the stored password file and will have to re-enter the password himself.

The preferred system identifies passwords automatically, without the need for instruction from the user, and stores the identified passwords and usernames in a data repository. Preferably this is central database 42. This allows the passwords of any user to be recalled regardless of the terminal at which the user logs on, providing the terminal has access to the central database.

Any identified passwords and usernames are stored in the database along with the field names in which they are stored on the original web site and the address of the Internet site to which they were transmitted to and at which they are used. Site information can be retrieved straightforwardly as it is included in the HTTP request submitting the password

- 29 -

and username information to that site, and in the representation of the web page held in memory.

Preferably, for the sake of security, the information stored in the database is encrypted, so
5 that only a select number of people, such as network supervisors, system administrators or company directors have access to it. They may access the database either through a workstation in the network, by entering a username or a password to identify themselves, or
10 through a supervisor workstation, such as Operate Consoles 44.

This storage of usernames and passwords along with address details presents a significant advantage to companies that use on-line facilities. With existing
15 technologies, if a user forgets his authenticating password preventing access to the protected file, or leaves the company without having disclosed it, the Internet service cannot be accessed. A similar situation occurs if the protected file is damaged,
20 deleted or otherwise lost. Each Internet service must then be approached in order to replace or recover the lost password, which can be expensive both in terms of lost access and administrative time. With the preferred system, the password information may be
25 recovered from the central database, so that access to web sites is not lost.

Figure 6 is a flow diagram which schematically illustrates the operation of a plug-in module implemented to extract user name and password
30 information from data to be transmitted to a web server.

In step S150, the plug-in module begins and parses the data about to be transmitted to the web server from the browser. This occurs at point 'C' in the process
35 illustrated in Figure 3. Control then passes to step S152 where the plug-in module determines whether or not

- 30 -

the data which is to be transmitted contains username or password information.

The passwords and usernames may be identified in the manner described above with reference to Figures 3, 4 and 5, by identifying field names in submitted command or by using the DOM, for example, to search for field names, field types, or the display method used to identify the data on web pages. They may also be retrieved from the HTML of web pages, the pop-up windows or GUIs (Graphical User Interfaces) presented by remote servers or providers on the World Wide Web or even by scanning the content of e-mail messages.

Identifying passwords and usernames in transmitted commands or in the DOM of a web page from their field names relies on those field names describing their purpose with obvious labels such as 'password' or 'username'. In cases when the field names are not themselves meaningful, the nature of data being transmitted, may be inferred from the field type of the data, that is 'String', 'Integer' and so on, or the display method used to enter the data. Fields that are intended to receive a password can be identified from the representation in by searching for a 'password' field type in the DOM. Text boxes on a web page into which password data is to be entered, for example, typically display each character input as an asterisk; this property may be determined from the DOM and used to infer that data input into the text box is a password even if there are no other indications. Although, the password is displayed as a string of asterisks, the representation in memory still contains the character information that was entered by the user. The password may then be retrieved simply by extracting the input from the field.

Alternatively, passwords and usernames may be identified by referring to those stored by other

- 31 -

programs such as Microsoft's 'Internet Explorer' when the remember password option is selected by the user. Such passwords are stored in a local protected file on the user's computer. This file is 'unlocked' when the
5 user is authenticated on the computer, and so may be accessed to obtain password and username information by the browser plug-in module of the preferred system.

If the plug-in module does not detect a username or a password in the data which is to be transmitted,
10 control passes to step S158 at which point the module exits and control is passed back to point 'C' in Figure 3. The browser may then transmit the data to the web server. If however, a username or a password is detected by the plug-in module at step S152 then
15 control passes to step S154, where the values of the identified username or password and the URL or other identifier of the web page to which the data is to be transmitted are extracted. Control then passes to step S156 where, these values and the URL or other
20 identifier are stored in a predetermined system database 42. Once storage has taken place, control then passes to step S158, where the module exits and control is passed back to point 'C' in Figure 3. The browser may then transmit the data to the web server.

25 The preferred embodiment need not be limited only to the storage of passwords or usernames which have been used as an example because of the immediate advantage that their storage provides. Other types of data, in particular those relating to eCommerce
30 transactions, such as credit card information and digital certificate may also be usefully extracted and stored to provide a database or record. The system for extracting information from transmissions may also be applied to e-mail systems.

35 The information may be extracted in the manner described above, via the DOM or via MAPI calls to the

- 32 -

COM representation of e-mail content, or may be extracted from the language in which a web page is encoded. Web pages are typically encoded in Hyper-Text Mark-Up Language (HTML), a human readable text based
5 language which may be searched for known key words or indicators using conventional text matching techniques.

In the preferred embodiment, recording of the data may involve recording just password and username information, recording the URL of a web page being
10 viewed or of an e-mail account, recording any data submitted to the fields of a web page, and recording the HTML of a web page, so that the web page may be retrieved later and viewed.

The plug-in modules provided by the preferred
15 system operate in conjunction with policy data, which may be recorded in a file, database, or software code for example. The policy data provides a user of the preferred system to instruct the operation of each of the plug-in modules thereby controlling their
20 functionality.

A sample representation of policy data, illustrated in Figure 7, shows how a user may control the operation of the plug-in module to record password and username information along with other types of
25 data.

The tree-like structure of the policy data is clearly seen in Figure 7 which shows just one major branch of the policy data entitled 'Recording'. The Recording branch is separated into two sub-branches
30 called 'Browser' and 'Email' which contain instructions for the operation of web browser and e-mail client plug-in modules respectively.

The browser branch contains three sub-branches called 'DataToRecord', 'WhenToStartRecording' and
35 'WhenToStopRecording'. The DataToRecord branch specifies the type of data that is to be extracted from

- 33 -

transmissions to or from the user's workstation and a web server. Four types of data are referred to in the illustration, these being the URL of the web page being viewed, the HTML of the web page being viewed, data
5 that is entered by a user into fields provided on the web page and submitted to a web site, and any passwords and usernames that are entered by the user. These are referred to on four distinct sub-branches of the DataToRecord branch, entitled 'URL', 'HTML',
10 'SubmittedFields' and 'Passwords'. A Yes/No option on each of those sub-branches specifies whether or not the indicated data is to be recorded.

The WhenToStartRecording branch sets out a number of conditions which specify the point at which the data
15 specified in the DataToRecord branch is to be recorded. Five conditions are illustrated in this example each one set out on a separate branch and respectively entitled 'WhenBrowserIsOpened', 'IfCreditCard
NumberSubmitted', 'IfPasswordSubmitted',
20 'IfKeywordsReceived' and 'IfKeywordsSent'. Whether or not these conditions are to be used in order to determine when to start recording is indicated by Yes/No markers on each branch.

Similarly, the WhenToStopRecording branch lists
25 three conditions that may be used to determine the point at which the recording of the data specified in the DataToRecord branch is to be stopped. These conditions are 'WhenUserClosesBrowser',
'WhenUserChangesSite' and 'WhenUserChangesPage'. The
30 Yes/No status of each of the conditions and for each of the type of data to be recorded may be easily set by a user of the preferred system in order to control the operation of the plug-in module.

The e-mail branch of the policy data is divided
35 into a DataToRecord branch and a WhenToRecord branch. Each of these branches are sub-divided into branches

- 34 -

that deal with Sent Mail and Received Mail. The type of data that may be recorded is set out in the DataToRecord branch and for sent mail may be the message text, any attachments to the message, and in
5 the case of messages signed with a digital signature, a copy of the digital certificate accompanying the signature. Conditions for recording sent mail and received mail are set out in the WhenToRecord branch and may be based on the identification of a credit card
10 number, a keyword or a digital certificate in the mail that is sent or received.

The tree-like structure described is the preferred form for the policy data since it allows the data to be easily organised and referred to. It also allows
15 different users to be assigned to different branches of the tree in order to receive different policies. Although, the tree like structure is preferred, other arrangements may also be possible. The branches shown in this diagram are only intended to be illustrative.

20

Identification of Credit Card Numbers

The preferred system also searches for Credit Card numbers or other account information in the data to be transmitted to the web server or e-mail client by
25 searching for a string of numeric digits, typically between 14 and 16 long. It then determines if the string of digits passes one of the tests universally used by credit card companies to validate card numbers.

If a credit card number is found in the transmission,
30 then the preferred system may take a number of actions depending on the settings in the policy file, such as refer the transmission to a third party for approval, renegotiate a higher level of encryption to safe-guard the transmission if the credit card number belongs to
35 the company, or prevent the transmission from taking place altogether.

- 35 -

The most common test to identify a credit card number is defined in ANSI standard X4.13, and is commonly known as LUHN or Mod 10.

The Luhn formula is applied to credit card numbers in order to generate a check digit, and thus may also be used to validate credit cards, in which case the check digit is figured in as part of the formula. To validate a credit card number using the Luhn formula, the value of every second digit after the first, starting from the right hand side of the number and moving leftwards, is multiplied by two; all of the digits, both those that have been multiplied by two and those that have not are then added together; any digit values that have become greater than or equal to ten as a result of being doubled, are summed as if they were two single digit values, eg: '10' counts as '1+0 = 1', '18' counts as '1 + 8 = 9'. If the total of the sum is divisible by ten, then the credit card is a valid credit card number.

Figure 8 is a flow diagram which illustrates the operation of a plug-in module which detects credit card numbers, using the Luhn formula described above, in data which is about to be transmitted. If a credit card number is identified, the plug-in module implements further policy-based checks, based on the outcome of which, the plug-in module may determine to transmit the data containing the credit card number or to prevent transmission.

The module begins operation at step S160, which follows point 'C' in the process illustrated in Figure 3 in the case of browser implementation, or point B in the process illustrated in Figure 5 in the case of an e-mail client implementation. Control passes from step S160 to step S162 in which the module scans the data about to be transmitted to the web server or e-mail

- 36 -

service and extracts from it a first string of digits that is likely to be a credit card number.

This is achieved by scanning the data comprising the transmission for strings of digits over a particular number of digits long. Credit card numbers are usually made up of more than 12 digits, and are not usually more than 16 digits long. Thus any strings of digits in this range may be identified as possible credit card numbers.

Following the extraction step S162 control passes to decision step S164 where a routine end of file check is made. If the data contains no candidate credit card number and the end of the file check is reached before any first candidate number could be found, then at step S164 control passes to step S178 where the transmission of the data is allowed to proceed without any further checks being made. The module then exits at step S180. Control resumes in the web browser operation shown in Figure 3 from point 'C' or in the e-mail client operation shown in Figure 5 from point B.

If a first potential credit card number is found in the data in step S162, then it is extracted and stored in memory. The end of file has not been reached yet so control passes from step S162 to step S164 and then to S166 where a checksum for the stored candidate number is calculated using the Luhn formula. Control then passes to decision step S168, where the checksum is queried.

If the checksum indicates that the candidate number is not a valid credit card number then control passes back to step S162 where the next potential credit card number is extracted from the data. If a second credit card number is not found, then the end of the file is reached and control passes to step S178 where the transmission is allowed to proceed, and then to step S180 where the module exits.

- 37 -

However, if the checksum indicates that the candidate number is a valid credit card number then control passes to decision step S170 where the settings of the policy data are queried for the appropriate
5 action to take. The action may be determined from such factors such as the number itself, the identity of the user transmitting the number and the address to which it is to be sent. The policy data might for example specify that credit cards are not to be transmitted, or
10 that a higher encryption strength is required for the transmission before it may proceed.

This policy checking step allows credit card transactions to be controlled at a higher level than the user making the transaction. Thus financial
15 decisions may be swiftly and easily implemented, and may be enforced automatically without the need for monitoring. An organisation may, for example, wish to restrict the capability to make credit card transactions on the organisation's account to
20 particular authorised people or it may wish to restrict transactions on a particular account altogether.

In step S170, the credit card number, and other details of the transaction are compared to the settings in the policy file and it is determined whether or not
25 transmission may take place. If for any reason, with reference to the policy checks, it is determined that the credit card number should not be transmitted, control passes to step S172 where transmission of the data is halted, and then to step S174 where the module
30 exits. At this point the system could notify the user that the request has been denied by means of a pop-up message box. Control then returns to point A in Figure 3, in the case of a web browser, or to step S132, 'compose mail' in Figure 5, in the case of an e-mail
35 client.

- 38 -

If in step S172, it is determined that the credit card number may be transmitted, control passes to step S176 where transmission of the data occurs, and then to step S180 where the module exits. In this case control
5 is resumed from point C in the web browser operation illustrated in Figure 3, or from point B in the e-mail client operation illustrated in Figure 5.

Credit card numbers need not be identified in step S162 solely by scanning the content of the
10 transmission. In web browser implementations the credit card number may for example be directly identified by referring to the field names of any variables that are to be transmitted or also from the representation of the web page in memory. The discussion above about the
15 identification of passwords explains this in more detail.

The preferred system may also be configured to search outgoing transmissions for other pertinent financial details, such as account numbers. Company
20 account numbers from which funds may be deposited may be stored in a separate file. Any likely strings of characters or digits may then be extracted from the outgoing data in the manner described and compared with the entries in the accounts file to determine whether
25 or not it is a valid account number. The transaction may then be allowed to proceed or be refused in the manner described above. Although credit card numbers have been referred to, it will be appreciated that any type of card number for making payment such as debit
30 card numbers may also be used.

Also, although identification of credit card numbers has been explained with reference to data that is to be transmitted, it will be appreciated that similar techniques could be used to identify and
35 extract credit card numbers from transmissions that are being received.

- 39 -

Validation and Authentication support

On-line transactions typically require some form of authentication that the user is who he says he is, and that he is able to pay for the goods ordered. These requirements are usually met by the purchaser supplying the trader with his credit card number and the cardholder address which can then be verified by the seller with the card issuer. Increasingly however, Digital Certificates are attached to electronic transmissions by the user which, in conjunction with a digital signature, allow the receiver to verify that the transmission originated from the person named as sender. Digital certificates from certain issuing authorities, such as Identrus, may also act as a warranty that the holder will honour his commitment to pay a specified amount of money. These certificates are useful when trading on-line.

Digital signatures are a widely used means of establishing an individual's identity on-line when transmitting information or when conducting a transaction. They also provide a guarantee to a recipient of any transmitted information or transaction details that those details and that information have not been tampered with en-route by an unauthorised third party.

Digital certificates are issued to individuals, organisations, or companies by independent Certificate Authorities, such as Verisign Inc. An organisation may also act as its own Certificate Authority issuing its own digital certificates which may or may not be derived from a 'root' certificate issued by another Certificate Authority. A digital certificate typically contains the holder's name, a serial number, an expiration date, a copy of the certificate holder's public key and the digital signature of the certificate

- 40 -

issuing authority. A private key is also issued to the certificate holder who should not disclose it to anyone else.

5 Certificates are unique to each holder and may be revoked by an issuer if the holder is no longer viable; alternatively, a holder may ask to have it revoked if his private key has been compromised.

10 The public and private keys may both be used in tandem to encrypt or decrypt a message. Anyone may use a certificate holder's public key to encrypt a message so that it may only be read by the certificate holder after decrypting the message with his private key.

15 Messages may also be digitally signed using software which converts the contents of the message into a mathematical summary, called a hash. The hash is then encrypted using the sender's private key. The encrypted hash may then be used as a digital signature for the message that is being transmitted. The original message, the digital signature and the transmitter's
20 digital certificate are all sent to a recipient who, in order to confirm that the message he has received is complete and unaltered from its original form, may then produce a hash for the received message. If, the received hash, having been decrypted with the holder's
25 public key, matches the hash produced by the recipient, then the recipient may trust that the message has been sent by the person to whom the certificate was issued, and that the message has not been altered en route from its original form. Digital certificates are therefore
30 of considerable and increasing importance to companies conducting business on the Internet.

In cases where an on-line trader makes use of Digital Certificates to ensure the identity of its customers, it is necessary to check with the issuer
35 that the certificate is still valid before any transaction is authorised. Such checks can be carried

- 41 -

out on-line using an independent verification service such as that provided by Valicert, Inc. A fee is usually charged for such services.

It may be that individual employees of an
5 organisation each receive e-mails from a single client, each signed with his digital certificate, on separate occasions. Presently, there is no way for a
information about certificates received by one employee to be shared with another employee unless they share it
10 themselves manually, and as a result, individual employees might request that the same certificate be validated each time that they receive it. This is wasteful however, since once a certificate is revoked by its issuer, it is never again reinstated, so any
15 validation fees spent on an already revoked certificate are unnecessary. Additionally, the recipient may wish to make a business judgement as to whether a previously validated certificate should be re-checked or not. For example, if a digitally signed order for \$1M worth
20 of goods is received one day, and the certificate successfully validated, and on the next day another order for \$50 is received, signed with the same certificate, the organisation may consider a second validation check to be unnecessary, thereby saving the
25 validation charge.

The preferred system provides means to record information about the Digital Certificates that have been received, the status of the certificate at the last check as well as, where applicable, transaction
30 information, such as client, amount, date, goods and so on. This information is stored in a central database to which all users of the system have access. The preferred system also provides means to use the stored information to decide whether or not a validation check
35 is desirable, and to accept or refuse transmissions depending on the status of the digital certificate.

- 42 -

Thus, users of the system may receive and review transmissions without the need to establish their authenticity themselves.

Figure 9 illustrates the operation of a plug-in
5 module of the preferred system implemented to extract digital certificates from transmissions received by company employees and record them in a database along with their validity status and details of any associated transactions such as date, amount, goods and
10 so on. The module first checks to determine if the certificate is obviously invalid, and that the message has been correctly signed by it. The certificate is obviously invalid, for example, if its expiry date has passed, or if it contains an invalid 'thumbprint'. Such
15 a thumbprint may be a checksum for the certificate itself for example. The message has not been correctly signed if the signature cannot be verified from the information contained within the certificate. Details of certificate validation and message signing are more
20 fully described in the ITU and RFC documents previously referenced. The module then makes a check to determine whether or not the certificate has already been stored in the database and only records those that have not. Where a copy of the certificate is already stored, the
25 module checks the database record to determine whether it has been previously identified as revoked in which case the transmission is immediately rejected. Otherwise, the module then determines, in accordance with policy defining business rules, whether or not to
30 validate the certificate. Taking in to account the results of any such validation, it then determines whether the certificate should be relied upon, and therefore whether or not the transmission signed by the digital certificate should be rejected or accepted. The
35 module is initiated at step S190, following receipt of data containing a digital certificate. Digital

- 43 -

certificates are typically transmitted as attachments to messages and may be identified by examining the first few bytes in the header of the attachment. These bytes identify the type of file attached and so will
5 indicate whether an attachment is a digital certificate or not.

The initialisation step S190 occurs after point A in Figure 3 if the module is implemented in a web browser, and after point A in Figure 5 if the module is
10 implemented in an e-mail client. Following initialisation, the module proceeds to step S191 in which the expiry date of the certificate is checked, and the digital signature confirmed as having signed the message. If the certificate has expired, or the
15 message is incorrectly signed, the module proceeds to step S198 and rejects the transmission. Otherwise the module proceeds to step S192 in which the database is searched for a previously received copy of the digital certificate. Control then passes to decision step S194.
20 If a copy of the certificate has been found in the database then control passes to decision step S196 where the module determines whether the certificate has previously been marked as revoked. This will have occurred if a previous validity check revealed that a
25 digital certificate has been revoked. If the certificate is not already in the database control passes from step S194 to step S202 where the new certificate and the date on which it was received are stored in the database, together with additional
30 details such as the address from which it was sent and details of any transaction associated with the transmission such as monetary value, account number etc. If the certificate has already been marked as revoked in step S196 then control passes directly to
35 step S198 in which the transmission comprising the digital certificate is automatically rejected. This may

- 44 -

involve for example, transmitting an automatically generated message to the initiator of the transmission whose certificate has been found invalid to explain the refusal, and preventing the recipient of the digital
5 certificate from taking any further steps in connection with the refused transmission. The module then exits at step S200.

If, however, the certificate has not been previously marked as revoked in step S196, then control
10 passes to step S204 in which the history of transmissions signed by the certificate, by other certificates from the same company, or used to conduct transactions on the same account are considered with policy data to determine if an on-line validity check
15 of the certificate is required. Control also passes to step S204 after a new digital certificate has been added to the database in step S202.

The policy data contains instructions which when considered in conjunction with the history of signed
20 transmissions previously received and revocation checks previously made indicate whether or not the certificate used to sign a transmission should be verified on this occasion. Example policy data is illustrated in Figure 10 to which reference should now be made.

25 The policy data is stored on the AcceptanceConfidenceRating branch on the DigitalCertificates branch of the policy data representation. The AcceptanceConfidenceRating branch sub-divides into two separate branches which deal
30 individually with 'monetary' digital certificates, where a certificate has been used to sign a transmission involving a transaction with the recipient for an amount of money, and 'identity' digital certificates which do not involve a monetary
35 transaction with the recipient of the transmission. Certain certificates are issued only for use in

- 45 -

conjunction with monetary transactions. For example, the 'warranty certificate' issued by some on-line banking organisations, such as Identrus, as a warranty to the receiver of the signed transmission. Such
5 warranty certificates testify that the sender of the transmission is a customer of an Identrus member bank, and that if he or she does not make payment, the bank will accept liability.

Organisations which issue different kinds or
10 classes of digital certificate mark each certificate according to its class. Identifying a certificate as being of a particular class is then a matter of knowing how different organisations classify their certificates and searching for the appropriate indicator in the
15 certificate received.

Issuers of digital certificates may provide many different classes of certificate suited for different purposes. These may be separately dealt with by the policy data by corresponding sub-branches of the policy
20 datatree.

In the example policy the first branch entitled 'IdentityCertificates' deals with transmissions that do not involve a monetary transaction. The branch comprises four separate sub-branches. The first of
25 these, entitled 'AlwaysAcceptFrom' contains a reference to a table, 'table a', which lists the names of individuals and organisations who are considered to be reliable. The names listed in this table will be those known and implicitly trusted by the company, for whom,
30 it is not considered necessary to determine whether or not a digital certificate has been revoked by its issuer.

The second branch entitled 'AlwaysCheckFrom' contains a reference to a separate table, table b, in
35 which the names of organisations and individuals for whom digital certificates should always be checked are

- 46 -

stored. Clearly, the contents of table a and table b will depend upon the experience of the user of the preferred system and will be left up to that user to enter.

5 The third branch entitled 'CheckIfDaysSince CertificateReceivedFromCompany' specifies a period of time following the receipt of a valid digital certificate from a company, within which checks of any further digital certificates received from that company
10 are not considered necessary. In this case, the period of time is set to be 10 days.

 The fourth branch entitled 'CheckIfDaysSince LastReceivedThisCertificate' specifies a similar time period in the case of an individual digital
15 certificate. In the example shown the policy data specifies that validation checks on any given digital certificate need only be made every 30 days. Again, the number of days specified on both of these branches is left to the user of the preferred system to decide.
20 The amount of time that has passed since a valid digital certificate has been received may be determined by referring to digital certificates and associated data stored in the database. By checking digital certificates periodically, rather than each time they
25 are received allows money spent on making checks to be saved.

 The MonetaryCertificates branch also contains an AlwaysAcceptFrom and an AlwaysCheckFrom branch which refer to tables x and y respectively. Table x lists
30 all those organisations and individuals for whom no status check of a digital certificate is required; table y lists all those for whom a check is always required. The MonetaryCertificate branch also contains a CheckIfAmountExceeds branch which specifies a
35 threshold transaction amount over which all digital certificates must be checked and lastly an

- 47 -

IfRecentlyChecked branch which sets out two conditions for performing checks on a digital certificate that has been recently received and validated. The IfRecentlyChecked branch allows the user to specify
5 that digital certificates received for transactions for a small amount, in this case \$5000, received within a specified time, in this case 30 days, of a previous revocation check, need not be validated.

Figure 11 illustrates the plug-in module process
10 which interacts with the policy data shown in figure 9. This process is a sub-process of that shown in figure 8 and occurs at step S204 resulting in decision step S206 in which the plug-in module determines whether or not to perform an on-line check of the status of the
15 Digital Certificate it has received. The sub-process begins in step S220 from which control passes to decision step S222 in which it is determined if the transmission is monetary from the class of the Digital Certificate used to sign the message. If the
20 transmission is monetary then control flows to decision step S232, which is the first in a chain of decision steps corresponding to the branches in the MonetaryCertificates branch of the AcceptanceConfidenceRating branch of the policy data.

25 If in step S222, it is determined that the transmission is not monetary, control passes to decision step S224, which is the first decision step in a chain of decision steps corresponding to the IdentityCertificates branches of the
30 AcceptanceConfidenceRating branch of the policy data. In each of the decision steps in the chain a simple check is made to see whether the conditions specified on each sub-branch of the IdentityCertificates branch of the policy data is satisfied. Depending upon the
35 results of that check control flows either to step S242 in which confidence in the Digital Certificate is

- 48 -

established and no on-line check of the Digital Certificate's status is deemed necessary, or step S244 in which confidence is not established and an on-line check is deemed necessary, or to the next decision step
5 in the chain.

Thus, in decision step S224, in which it is determined whether the sender of the Digital Certificate is listed in table a, that is the 'AlwaysAcceptFrom' table, if the sender of the Digital
10 Certificate is listed in Table A then control flows from decision step S224 to step S242 where confidence is established in the Certificate and the sub-process ends returning to step S208 in figure 8. If the sender is not listed in table a then control flows from step
15 S224 to the next decision step in the chain step S226 in which it is determined whether the sender of the Digital Certificate is listed in table b, that is the 'AlwaysCheckFrom' table. Similarly, if the sender is listed in this table control flows to step S244 at
20 which an on-line check of the Digital Certificate's status is deemed necessary. Control returns from step S244 in the sub-process to step S210 in figure 8.

If the sender of the Digital Certificate is not listed in table b then control flows from decision step
25 S226 to the next decision step in the chain which represents the next condition listed as a sub-branch listed in the policy data. Thus, in decision step S228 a check is made as to whether this Digital Certificate has been validated in the last 30 days. This will
30 involve looking up the Digital Certificate in the database of stored Digital Certificates and extracting from the stored information the date on which the Digital Certificate was last checked. If the status of the Digital Certificate has been checked in the last 30
35 days, control flows to step S242 where confidence is established. If the information in the database of

- 49 -

stored Digital Certificates indicates that the Digital Certificate has not been checked in the last 30 days then control flows from step S228 to decision step S230 in which a check is made to see if another Digital Certificate has been received from the same company and if that Digital Certificate has been checked within the last 10 days. This determination again involves checking the database of stored Digital Certificates and information relating to those Digital Certificates.

10 If the other Digital Certificate has been checked in the last 10 days then control flows to step S242 where confidence in the received Digital Certificate is established. If not, then control flows to step S244.

In the case of a monetary transmission, the conditions set out on the policy data are stepped through in decision steps S232 to decision step S240. If in decision step S232 the sender of the Digital Certificate is found to be listed in table x, which sets out the names of companies and organisations for which no check of the Digital Certificate status is deemed necessary, then confidence is established and control passes to step S242. Otherwise control passes to the next decision step in the chain which is decision step S234. At decision step S234 if the sender of the Digital Certificate is founded listed in table b, that is the 'AlwaysCheckFrom' Table, then confidence is not established and control passes to S244.

20

Otherwise, control passes to decision step S236, where it is determined whether the amount for which the transaction is being made exceeds \$10,000. This determination is made with reference to the signed transaction data which will contain the monetary amount either in a manner pre-determined by the issuer of the certificate, or contained within the e-mail or associated e-mails forming the transaction or transmissions. If the transaction is found to be for an

30

35

- 50 -

amount to be in excess of \$10,000, or if the amount of the transaction cannot be determined with a reference to the data transmitted, then confidence is not established and control passes to step S244. Otherwise, control passes to decision step S238 in which it is determined whether the Digital Certificate has been checked within the last 30 days. Again, this determination is made with reference to the database of stored Digital Certificates and data relating to Digital Certificates. If the certificate has not been checked within the last 30 days then confidence is not established and control passes to step S244. If it has been checked, then control passes to decision step S240 where, if the previously determined amount of the transaction is deemed to be over \$5,000 then confidence is not established and control passes to step S244. If the amount of the transaction is under \$5,000 then it is classed as an acceptable risk to rely on the Digital Certificate, confidence is established and control passes to step S242.

These last two conditions allow the system to determine whether to check the certificate's status based on recent trading history. If for example, a transaction is about to be made by a party for a modest sum, i.e. under \$5000, and the search of the recorded transaction and certificate details reveals that quite recently, the same party made a transaction and at that time their digital certificate was confirmed as being valid, then it is arguable that checking the validity of the party's certificate again so soon after the first is unnecessary, and it is preferable to trust in the party rather than paying the validation fee a second time.

It will be appreciated that the instructions in the policy file may be set out to reflect the level of confidence the company has in its customers or

- 51 -

suppliers, based on the experience of individuals within the company, the amounts of transactions that are deemed allowable without significant risk and so on. The policy file may also be set up to implement
5 more general policies which are to be used in conjunction with a record of transaction details for the holder of that digital certificate. For example, any transaction that is offered by the holder may be compared against the record of those that he has made
10 before to see if the amount and the goods and services requested are in keeping with his trading history. If they are not, then it may be desirable to check the validity of the certificate to confirm that it is still valid and guarantees the identity of the sender. If it
15 has been revoked, then a third party may have acquired the original holder's private key and be attempting to make fraudulent transactions.

Having checked the policy data in step S204, confidence in the digital certificate will either have
20 been established or will not have been established. In decision step S206, if confidence has been established then control passes to step S208 in which the transmission containing the transaction is accepted. Control then passes to step S200 where the module exits
25 and control passes back to point A in Figure 3 in the case of a web browser, or point A in Figure 5 in the case of an e-mail client.

If in step S206 confidence in the digital certificate is not established, then control passes to
30 step S210, where an on-line validation check is performed on the digital certificate. This may involve checking to see whether the Digital Certificate has been revoked or whether, in the case of an eCommerce transaction, the issuer of the Digital Certificate will
35 confirm a warranty for the amount promised in the transaction. Control next passes to step S212, in which

- 52 -

the validity status stored in the database for that certificate is up-dated. Control then passes to decision step S214 in which, if the certificate was found invalid, control passes to step S198 where the
5 transmission is rejected, or to step S208 where the transmission is accepted. Rejection of the transmission may mean that it is deleted from the recipients mail box before it has been opened, or that the transmission is marked with the word 'rejected' or
10 some other indicator. Following either of steps S198 or S208, control passes to step S200 where the module exits. Whenever a transaction is allowed to proceed the database is updated such that it includes information about the transaction, such as date and
15 amount, in order that the information may be used in determining the need for future validation checks.

Recording of information.

The preferred system also provides an automatic
20 way in which to record transaction information for transactions carried out on-line. In this context, the word 'transaction' and 'eCommerce transaction' are intended to mean an agreement promising money or money's worth made between two parties over the
25 Internet, or even over the same network of a company. Normally, the user himself is responsible for maintaining transaction information by making hard copies of the relevant electronic records or by
actively storing copies of any electronic records in
30 files on his computer. The reliance on manual methods to ensure these records are maintained is clearly both unreliable and labour intensive.

The preferred system on the other hand scans the information content of all communications processed by
35 the system for indications that a transaction has begun or is taking place. Such indications are numerous. The

- 53 -

most straightforward is whether the link is secure or not as most web pages negotiate a secure link before a transaction is carried out and close that link afterwards. Determining whether a link is secure is
5 achieved by examining the URL of the destination web page. A secure link is indicated by an 's' after the prefix 'http'. Thus, one mode of operation of the preferred system is to record all data transmitted to a web page while a link is secure. The preferred system
10 also maintains a record of web pages that negotiate secure links but that are not eCommerce sites, i.e. those that are connected to other than to make purchases, and not record any data transmitted to these pages. Such a web page might be Microsoft's Hotmail web
15 page which provides an e-mail service.

Another indication might be simply the URL of the site. In this case the preferred system may be configured to record all data transmitted to a web page which has been identified as that of an on-line trading
20 company. Other indications might be an identified credit card number, an electronic receipt, an e-mail acknowledging the sale, use of a digital certificate, in particular a digital warranty certificate, or a purchase code.

25 Once a transaction has been identified as occurring, the preferred system may record the details of the transaction by both storing in entirety each communication between a user and the identified trader, or by scanning the transmissions and extracting
30 particular details, such as date, amount, type of goods, quantity and so on.

Recording of transaction data may be stopped when the end of the transaction is identified or after a predefined number of transmissions between the
35 purchaser and the trader have taken place. Similarly, once a transaction has been identified, the preferred

- 54 -

system may record in the database a predefined number of cached transmissions that took place immediately before the first recognised transmission of the transaction.

5 This will be useful when, for example, the first indication that a transmission is occurring is the detection of a credit card number or an electronic receipt, since these are likely to be received at the very end of a transaction. The prior transmissions may, for example consist of web pages containing information
10 regarding the goods or services being purchased, or an exchange of e-mails where specification or delivery terms were agreed. Note that it is perfectly possible for the prior transmissions to be of the same type as
15 that in which the transaction was detected, of a different type, or be a mixture of types. For example a user might visit a web site www.abc.com, obtain details of goods, and then order them in an email sent to orders@abc.com.

20 The preferred system records the transaction details in a centralised common database 42. Additionally, the database may be a local file or a service on a network. The information stored in the database may be encrypted using known encryption
25 techniques so that only a person with the necessary authorisation may access it.

Figure 12 is an illustration of the operation of an example implementation of a module for identifying when an electronic transaction is being conducted on-
30 line. Figure 14 illustrates the process by which the preferred system records an identified transaction in the database, and Figure 15 illustrates how the preferred system allows an identified transaction to be approved or rejected on the basis of a predetermined
35 approvals policy.

- 55 -

With reference to Figure 12, the operation of a module for identifying when an on-line transaction is occurring will next be described.

The module begins operation at step S250 in
5 response to receiving data or in response to a user
initiating transmission of data to a remote site. In
the case of a web browser implementation this will be
after point A or after point C respectively as shown in
Figure 3; in the case of implementation in an e-mail
10 client it will be after point A or B respectively as
shown in Figure 5.

Control then passes to decision step S252 in which
it is determined whether, in the case of a web browser,
a secure link has been negotiated between the site
15 transmitting data and the site receiving data. This
may be achieved by querying the URL address that has
been connected to, as mentioned above, or interrogating
the web browser to see if encryption is being employed.
In the case of an e-mail message this step is omitted
20 and control passes directly to step S260. Since on-
line web browser transactions usually involve the
transmission of personal information, such name and
address, credit card number or other account
identifying information, secure links are usually
25 negotiated as a matter of course. Thus the presence of
a secure link alone is a good indication that a
transaction is taking place. However, secure
connections may be negotiated for reasons other than
the transmission of transaction details. Thus, if in
30 step S252 it is determined that the connection is
secure, control passes to step S254, in which the
address of the remote site to which the connection has
been made is checked against a list of known sites
which do not provide facilities for conducting on-line
35 transactions but which do establish secure connections.
Browser based e-mail sites such as Microsoft's Hotmail

- 56 -

site is one such example. Control then passes to decision step S256 where a determination based on the previous check is made. If the site address is identified as a non eCommerce site, that is one which
5 does not facilitate the making of a transaction, then it is determined that a transaction may or may not be occurring and control passes to decision step S260 to make further checks on the transmission content. If in step S256 the site address is not identified as a known
10 non eCommerce site, then it is assumed that an on-line transaction is occurring, and the module exits at step S258.

If a secure connection is found not to have been established at step S252, or if a secure connection has
15 been established but to a known non-eCommerce site, as determined at step S256, or the transmission is an e-mail, then control passes to decision step S260. In decision step S260, the first of a number of checks on the content of the transmission are made in order to
20 determine whether or not it is part of a transaction. In step S260, the transmission is scanned to see if it contains a credit card number. The method for doing this has been described with reference to Figure 8. If a credit card number is found in the transmission then
25 it is assumed that a transaction must be occurring and control passes to step S258 at which the module exits. If no credit card number is found then control passes instead to decision step S262 where the transmission is scanned to see if it contains an account code. Account
30 codes may be (for example) stored in a separate file which is accessed by the module when performing this step or alternatively an account code may be identified from descriptive data in the transmission such as a field name like "Account Number" or similar characters
35 appearing in the text of a message.

- 57 -

If in decision step S262 an account code is found, then it is assumed that the transmission constitutes part of a transaction and control passes to step S258 where the module exits. If no account code is found
5 then control passes to step S264 where, in the case of a web browser, the URL is compared with a list of known eCommerce URLs stored in a file or in a database. In decision step S266, a determination on that comparison is made. If the URL is found to be at a known eCommerce
10 page, or within a known set of eCommerce pages, then it is determined that an eCommerce transaction is taking place and control passes to step S258 where the module exits. Similarly, in the case of an e-mail, the destination address may be compared against a list of
15 known eCommerce e-mail addresses, for example 'orders@abc.com', and if a match is found then it is determined that an eCommerce transaction is taking place and control passes to step S258 where the module exits.

20 The checks that have been described are only representative of the possible checks that could be made to determine whether or not a transmission is likely to be part of an eCommerce transaction and are not intended to be exhaustive. Furthermore, the order
25 in which the checks have been illustrated has no special significance. The order is simply dependent on the structure of the policy data as will be seen from reference to Figure 13.

In step S268, a general check is illustrated which
30 represents any further check for an indication of a transaction, in addition to those described above, which a company decides it is desirable to employ, such as searching for purchase codes or embedded codes placed in the data for example. It is preferred that
35 the web browser or e-mail client being used in the preferred system allows the user to mark transmissions

- 58 -

with an embedded code to indicate that the transmission is part of a transaction and should be recorded. Also, the embedded code could be placed in the data by the web site or the e-mail client transmitting some of the transaction data to the user's workstation.

Control passes to this step after step S266, if the site is not recognised as a known eCommerce site and if such a transaction indicator is found in step S268, then a transaction is deemed to be occurring and control passes to step S258 where the module exits. If at step S268, no such indicator is found then no transaction is deemed to be occurring and the module exits at step S258. Following exit, the data may be transmitted, following points C and B in Figures 3 and 5 respectively, or be processed following on from its receipt at point A in Figures 3 and 5.

In the example described, the aim is to start recording transmissions and possible transaction details if there is just a suspicion that a transaction is occurring. It is assumed that recording data which is not part of a transaction is preferable to not recording a transaction at all. Figure 13 is an illustration of the policy data used to identify that eCommerce transaction is occurring and to control the way in which the transaction data is recorded. The policy data is represented by a Transactions branch of the policy data tree which sub-divides into two separate sub-branches called 'Identification' and 'Termination'. The Identification branch is itself divided into five sub-branches which correspond to the determinations made in the process illustrated in Figure 12. The first of these sub-branches entitled 'IfConnectionGoesSecure' allows a user to specify whether recording should start when the plug-in module detects that the connection to the web server has become secure. The condition specified on this sub-

- 59 -

branch corresponds to decision step S252 shown in Figure 12. It will be appreciated with reference to Figures 12 and 13 that the flow of control shown in Figure 12 corresponds to the layout of the conditions specified on the branches of the policy data tree shown in Figure 13. The ExcludedSites branch of the IfConnectionGoesSecure branch contains a reference to table q in which the web sites known to negotiate secure sites but which are known not to be eCommerce web sites are listed. Table q is referred to in step S256 of the process shown in Figure 12.

The next sub-branch of the identification branch is entitled 'IfCreditCardNumberPresent' and allows the user to specify whether the detection of a credit card number should or should not be used to initiate recording of data being transmitted or received. This sub-branch corresponds to decision step S260. The PreviousPages sub-branch of the IfCreditCardNumberPresent branch lists the number of web pages, prior to the web page in which the credit card number was detected, which should also be recorded. Since credit card numbers are normally submitted at the end of the transaction, the provision of this sub-branch enables previous web pages which are likely to contain the details and request of the transaction to be retrieved and stored. These web pages are cached continuously by the preferred system so that should a transaction be identified they may be retrieved from the cache and stored in the database. This will be explained in more detail with reference to Figure 14.

The next sub-branch of the Identification branch tree is entitled 'IfAccountsCodePresent' and allows a user to specify whether or not the detection of an account code in the data being transmitted or received is to be taken as an indicator to initiate recording of

- 60 -

the data. The accounts codes are identified in step S262 shown in Figure 12 by reference to table r. The reference to this table is contained in the AccountCodes sub-branch of the IfAccountCodePresent branch. Note that this table also shows the number of previous pages to record, in a similar manner to that described above for credit card identification, however in this case the number of previous pages to record is stored in table r allowing a different number of pages to be specified for each detected account code.

The IfKnownECommerceSite branch allows the user to specify a list of URLs corresponding to sites, parts of sites, or even single pages, where eCommerce transactions are known to take place. The current page URL is matched against entries in this list to determine if a transaction is taking place. The KnownSites sub-branch contains a reference to table s in which the URLs of known eCommerce sites are stored. The determination of whether the URL of the web site is a known eCommerce site is made in decision step S266 following step S264 of Figure 12. Lastly, the IfOtherIndicatorPresent branch provides a way for the user to specify whether or not the determination of other indicators should be used as a starting point for the recording of data. Two sub-branches of this branch entitled KeyWords and PreviousPages specify possible indicators that may be detected, in this case key words listed in table t, and also the number of previous pages that are required to be stored if key words are detected.

The Termination branch of the Transactions branch divides into four sub-branches which specify conditions which are used to end the recording of data being transmitted or received. Each sub-branch sets out a condition by which the end of the transaction may be defined. The first branch entitled

- 61 -

'IfConnectionGoesInsecure' allows a user to specify that the relinquishing of a secure connection by the web browser indicates the end of a transaction so that recording should be stopped. The other sub-branches
5 specify that when the web site changes recording should stop, if a digital receipt is received recording should stop and recording should stop after the receipt of 20 web pages following identification that a transaction is occurring.

10 It must be stressed that policy data shown in this diagram, in particular, but also in the other diagrams is unique to each user. Not only may a user specify whether or not particular conditions are to be acted on by setting the Yes or No variable accordingly, or by
15 changing the number of pages that are to be recorded for example, but also the structure and arrangement of branches and conditions specified on those branches may be different from user to user. It will be appreciated that while the example policy describes recording of
20 transactions in a web browser environment, a similar policy would control the e-mail environment, omitting the secure connection option, but allowing policy to be defined for recording e-mails on detection of credit card numbers, account codes or other identifiable
25 information within them, or where e-mails are sent to known eCommerce addresses.

The full benefit of the method for identifying a transaction is realised when the method is utilised along with means for recording transmissions between a
30 user of the preferred system and a remote site. This allows a record of all transactions carried out by a user to be kept and maintained automatically. The records may be kept up to date without the need for making paper copies of each transmission received or
35 sent. Thus, a company's record keeping is made considerably easier and more accurate.

- 62 -

Figure 14 illustrates the operation of a module for recording transmissions which comprise a transaction. The module is initiated at step S270.

If the module is implemented as part of a web browser, step S270 is initiated at point A in Figure 3 after the receipt of data or after point C in Figure 3 directly before transmission of data to a remote site. If the module is implemented as part of an e-mail client step S270 occurs after point A in Figure 5 after an e-mail has been received or after point B in Figure 5 just before an e-mail composed by the user is sent to a recipient.

Following step S270 , control passes to step S272, in which the test for identifying a transaction, described above with reference to Figure 9, is performed and a determination is made as to whether an eCommerce transaction is occurring or not. Control then passes to decision step S274 where, if it is determined that no transaction is occurring control passes directly to step S276 where the module exits.

If a transaction is determined to be occurring however control passes to step S278 in which the policy is consulted against one or more of the means of detection, the identity of the sender, the amount of the transaction, or other parameters to determine which prior transmissions, if any, should be stored with the identified transmission, and in how much detail the transmission should be recorded. The policy might, for example, require that a transaction involving a large sum of money be recorded in more detail than a transaction for a small sum. An example of this in operation might be the recording of every web page accessed during the making of a transaction on an on-line trader's web site for transactions involving large sums of money, but only recording the transmission

- 63 -

containing an electronic receipt for transactions for smaller amounts.

As well as determining the amount of data to be stored, the policy file may also determine the nature of the data to be recorded. The entire transmission or web page may be recorded as a series of snap shots of the transaction, in the same way as web pages are stored in cache memories for example, or alternatively, individual items of data, such as the date, the trader's identity, the amount and so on, may be extracted from the transmission or web page, and stored either on their own or together with the snap shot data.

In this way, memory for storage can be used most effectively to ensure that the most important transactions have sufficient space to be recorded. The amount of transaction data to be recorded may also depend on the trader's identity, geographical location, trading history with the user's company, and the goods and services on offer.

In Figure 13, the example policy data shows a simple scenario in which the amount of data to record is specified in terms of the number of web pages that are to be retrieved from the pages cached in memory. The number differs depending on whether a credit card number, an account code or a keyword is identified. Furthermore, table r shows that with the recognition of different account codes the number of previous web pages to store might be different, reflecting the relative importance of the account.

Extending this simple case to a more sophisticated one may be achieved by providing a higher level of detail in the policy data. Additional branches to the policy data tree could specify company or individual names, or specific keywords relating to goods and/or

- 64 -

services; the amount of data to record depending on these keywords and names as well.

Also, the tables might be expanded to refer to the amount of different types of data that should be
5 stored. Data such as the company name, what was being sold, the quantity and so on could be extracted from e-mail text, from the HTML text defining the web page, or from the DOM representation of the web page and stored in the database.

10 All web pages or information stored in the cache may be retrieved, or alternatively the system may retrieve only pages that have details in common with the page initially identified as being part of a transaction.

15 Alternatively, a list of all stored messages may be presented to the user for the user to manually select those transmissions which relate to the identified transaction.

Following determination of how much data to record
20 control passed to decision step S280. In step S280 if earlier transmissions, are to be stored, control passes to step S282 where the transmissions stored in the local cache are retrieved. In the case of a web browser, this may be a determined number of prior
25 pages, as described above. Where the transaction was detected in a web browser, policy may also dictate that the cache is searched for prior e-mail messages relating to the transaction, for example sent to or received from the same organisation. This may be
30 determined by matching portions of the browser URL to portions of the e-mail addresses. Similarly, transactions detected in e-mail messages may cause both prior e-mails and prior web pages to be retrieved from the cache. Control then passes to step S284 in which
35 the identified transaction and any retrieved prior transmissions are stored in the system database 42.

- 65 -

In step S280, if earlier transmissions are not required, control passes directly to step S284 where the transmission identified as a transaction is recorded in the system database. At the same time as
5 the transmissions are stored in step S284, related data such as the user identity, the amount and the other party to the transaction may also be recorded in the system database to form a complete record, although this will depend on the instructions of the policy
10 data. Control then passes to S286 and the module exits.

Following on from step S276 after the module exits, the data may be transmitted, following point A in figures 3 and 5, or be processed following on from
15 being received at points C and B in figures 3 and 5 respectively.

Once a transmission has been identified as taking place all transmission between the user and the other party may be recorded until the system detects that the
20 transaction has been completed. Detecting the end point of a transaction and stopping recording may be done in a manner similar to that described above for identifying whether a transaction is taking place. The most simple implementation is to record transmission
25 information until an electronic receipt or shipping order is received. Alternatively, recording of transmissions may be caused to stop after a pre-determined number of transmissions between the user and the other party have occurred, or if a certain amount
30 of time has passed since the transaction was identified.

Transmissions may be made simpler if each time the user changes web site the cache is emptied. This keeps the memory required for the cache memory low, as well
35 as reducing the number of previous transmissions that

- 66 -

need to be searched if searching techniques are to be employed.

It will be appreciated that the methods described above can also be used to record associated
5 transmissions which occur after a transaction has been detected and recorded. For example, a transaction made using a web browser will typically be followed by a confirmatory e-mail sent from the seller to the buyer.

This e-mail can be detected as forming part of the
10 transaction, since it will contain common characteristics, such as order number, account number, goods description, price etc. It may also be sent from an address similar to the web site address, for example 'customerservices@abc.com' when the original web site
15 used to make the purchase was 'abc.com'. A time element is preferably used such that only subsequent transmissions which occur within a given time period are considered as being associated with the original transaction.

20 In addition to recording transaction information, it may be advantageous to record other information which provides management with the ability to analyse the behaviour of their users, for example to ensure that the organisation is in fact obtaining real
25 productivity benefit from its adoption of electronic commerce. Such information is not limited to the productivity of the user himself, but to the overall process, allowing for example a comparison of web purchasing sites to determine which are the most
30 efficient in terms of the purchase process, and therefore which will be of greatest benefit in reducing purchasing costs. The preferred system provides for this by recording additional information, such as the amount of time a purchase takes, the number of
35 keystrokes and mouse clicks required to complete a purchase, the amount of 'idle' time while the user

- 67 -

waits for pages to download or responses to be received. This information can be recorded with the transaction record in the database, allowing statistical analysis to be carried out across a range
5 of transactions.

The time taken for a transaction can be determined by associating a timestamp with each of the transmissions received. When it is determined that the transaction is complete, the timestamp associated with
10 the first transmission (which may have been recovered from the cache in step S282) is subtracted from the timestamp associated with the last transmission, and the result, which will be the overall transaction duration, is stored in the database in step S284.

15 Alternatively, first and last timestamps could be recorded in the database and the transaction duration calculated later. The number of keystrokes and mouse clicks can be determined in a Microsoft Windows based system by using standard Windows 'hooks' into the
20 operating system. Such techniques are described more fully in the document "Win 32 Hooks" by Kyle Marsh of the Microsoft Developer Network Technology Group, dated 29 July 1993 available on the Microsoft Corporation web site ([http://msdn.microsoft.com/library](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnmgmt/html/msdn_hooks32.asp)
25 [/default.asp?url=/library/en-us/dnmgmt/html/msdn_hooks32.asp](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnmgmt/html/msdn_hooks32.asp)).

The preferred system maintains counters of the number of keystrokes (using the WH_KEYBOARD hook) and mouse clicks (using the WH_MOUSE hook) that occur
30 between each received transmission and associates these totals with the received transmission. Keystrokes and mouse clicks which occur while another application has the focus (for example if the user temporarily switches to another application) are ignored. When it is
35 determined that the transaction is complete, the totals of keystrokes and mouse clicks which occurred between

- 68 -

the first transmission (which may have been recovered from the cache in step S282) and the last transmission are added together, and the result, which will be the total number of keystrokes and mouse clicks during the overall transaction, is stored in the database in step S284. Similarly, the web site transaction response time can be measured by noting the time at which each outbound transmission request is sent, and then subtracting it from the time the response transmission is received. Aggregating the response times between the start and end of the transaction will give the total time the user spent waiting for the web site. Similarly the preferred system also counts the user response time, that is the time between a transmission being received, and the time a response is transmitted.

The preferred system also calculates how much of the user response time is spent entering data, and therefore allows the time required for the user to 'absorb' the incoming transmission (being the difference) to be determined. The time spent entering data is determined by maintaining a 'stopwatch'. The stopwatch is reset each time a new transmission is received, and is re-started immediately whenever the user enters a keystroke or clicks the mouse. Should the user not enter a keystroke or click the mouse for a pre-determined period of time, for example 5 seconds, the system assumes that the user is now absorbing details of the previous transmission and stops the watch. The watch is also stopped when the keystroke or mouse click causes an outbound transmission to be sent.

Aggregating the time spent entering data between the start and end of the transaction will give the total time the user spent entering data on the web site. The aggregated times can be stored in the database at step S284 for future analysis.

- 69 -

The preferred system also provides means for monitoring transactions that are being made and automatically referring the transaction for approval if that is deemed necessary. This process allows a large
5 company to monitor and control the transactions being made by its employees using a single set of criteria set out in the policy data. The policy data may be referred to each time a transaction is identified in order to determine whether the user is authorised to
10 make that transaction himself or whether he needs to request authorisation from higher up in the company. The process is illustrated in Figure 15 to which reference should now be made.

The module embodying this process is initiated at
15 step S290. This initiation preferably takes place as soon as all relevant details of the transaction which need to be considered have been determined, and before the transaction is committed. In the case of an e-mail transaction, details such as goods and price are
20 typically contained within a single e-mail and can be considered prior to transmission of that e-mail. In the case of a web browser transaction, the existence of transaction may be detected before all details are known, in which case initiation does not take place
25 until they are. This does not normally present a problem as final commitment does not occur until the very end of the transaction process, well after all relevant details are known. Detection of the transaction and relevant details may be determined in
30 the manner described above with reference to Figure 12. Referring briefly to figures 3 and 5, it will be seen that step S290 occurs after point C in figure 3 in the case of web browser implementation, or after point A in figure 5 in the case of e-mail client implementation
35 once the required details are known.

- 70 -

Control passes from step S290 to decision step S292 in which the details of the transaction are compared against the policy settings to determine whether or not approval is required. The determination
5 may be based on the identity or the position of the employee making the transaction, the amount of the transaction, or the other party in the transaction. In some instances, approval might always be required, such as if the financial director of a company wishes to
10 review each transaction before it is made.

Figure 16 is an illustration of example policy data that may be used to determine whether or not a transaction requires approval from a third party and also to determine the identity of an appropriate
15 approver who is to be used. In this case, conditions in the policy data stipulate whether approval is required depending on the transaction amount, and the URL address of the other party to the transaction.

The relevant policy data is set out on the
20 Transactions Approval branch of the policy data tree. This branch sub-divides into four sub-branches. The first branch is entitled 'MaximumUnapproved TransactionAmount' and defines a threshold amount for transactions. Transactions for amounts over the
25 threshold must be approved by an approver before they are made.

The second sub-branch entitled 'MaximumUnapproved MonthlyAmount' defines a maximum amount for transactions that a user may make within a month. In
30 this case, any transaction made by the user which would cause the monthly total to exceed \$2,500 will require approval from a third party, as will further transactions made after that threshold has been reached.

35 The third branch entitled 'ExcludedSites' refers to a table containing web site and e-mail addresses of

- 71 -

all sites which always require approval from a third party before a transaction may be made. Finally, the last branch entitled 'Approvers' refers to a table in which the names of possible third party approvers are listed. Along side each name is the maximum transaction amount for which that approver has the authority to approve, and a list of excluded sites for which that approver may not approve a transaction. In the simplest of cases, approvers will be other computer users logged in on the same network as the user making the transaction such as department managers, or supervisors. The approvers will, by nature of their role, be members of the trading company who assume and who have the authority to assume responsibility for the financial transactions the company makes. It is also possible that approvers might be drawn from a group of people who are employed primarily for this role, such as people in the financial department only.

If the conditions on the first three sub-branches of the transaction branch indicate that approval is required an appropriate approver may be found by scanning through the table of approvers until an approver whose transaction limit is equal to or greater than of the proposed transaction and who is not prohibited from approving transactions to the relevant site is found.

It will be appreciated that the example policy data shown in Figure 16 is policy data that is specific to a single computer user, or group of users, in the network. Other users, or groups, may have different settings and a different list of approvers.

It will be appreciated that the conditions to determine an appropriate approver may be introduced by creating new sub-branches of the policy data tree.

The operation of the approvals process could for example be extended to any kind of transmission, not

- 72 -

just those comprising an eCommerce transaction. Such operation may be implemented by having the conditions defined or the sub-branches of the policy data specify usernames, addresses or keywords for example that are
5 to be identified in the transmission and acted upon. Thus, all e-mail transmissions to a particular company or individual may be caused to require approval, or all e-mails containing predetermined information recognised via keyword identification.

10 If it is determined in step S292 that no approval is required, control passes directly to step S294 at which the module exits. Following step S294 the transmission of the transaction is permitted and the transaction may be proceed. Control returns from step
15 S294 to point C in figure 3 or to point B in figure 5.

However, if in step S292, after consulting the policy settings, it is determined that approval of the transaction is required, control passes to step S296 in which the particulars of the transaction are used to
20 determine an appropriate approver for the transaction.

The approver may be a company employee logged on at his workstation, or at a workstation with a dedicated approver function such as Operator consoles 44, as shown in Figure 2, or may even be an automated process.
25 In the case of a large company with a number of departments, it maybe advantageous to have a group of approvers for each department, each group monitoring the department's accounts. This allows transactions to be rejected before they are made if, for example if the
30 department head decides that he wishes to temporarily suspend purchases, or purchases of a particular nature.

Control passes from step S296 following determination of an appropriate approver to step S298 where an approval request is transmitted to the
35 designated approver via systems approval queue 100. Following step S298 control passes to decision step

- 73 -

S300 where it is determined if a response from the approver has been received. A timer is started the moment a request for approval is submitted. If a response has not been received in step S300 control
5 passes to step S302 where it is determined from the timer whether or not a time out period has elapsed. Providing the period has not elapsed, control passes from step S302 back to S300 where the system continues to wait for a response from the approver. Thus decision
10 steps S300 and S302 form a loop in which the system waits until a response is received or until time out expires. In decision step S300 once a response is received, control passes to step S304 in which action is taken depending upon whether the transaction was
15 approved or rejected.

If the transaction was approved control passes from step S304 to step S294 at which the module exits and the transmission is allowed to proceed. If however the transmission is not approved then control passes
20 from step S300 to step S306 at which the module exits. Exiting at step S306 however prevents the transmission of the transaction from taking place and returns the user to point A in figure 3 in the case of web browser implementation or to step S132 "compose e-mail" in
25 figure 5 in the case of e-mail client implementation.

Also, if in step S302 the 'time-out' is deemed to have expired without a response from the approver having been received then control passes directly to step S306 at which the module exits.

30 The right hand side of Figure 15 shows the steps involved for the approver. The approval process is initiated in step S310, from which control passes to step S312 in which the approver's machine queries the systems approval queue for any new approval requests.
35 Control then passes to decision step S314. In step S314 if no request is pending control passes back to

- 74 -

step S312 where the system queue is polled once again.

These steps are repeated until an approval request is received or until the approver deactivates the approval process.

5 In step S314 if an approval request is received, control passes to step S316 in which the approval request is downloaded from the systems queue and the approver himself decides whether to approve the request or refuse it. Control then passes to step S318 in
10 which the approver's response is transmitted back to the system approvals queue and from there back to the users work station.

Control passes from step S318 back to step S312 in which the system approvals queue is queried for new
15 approval requests. It will be appreciated that the approvals process could be entirely automated in some circumstances. For example, transactions may be automatically rejected if the company does not have sufficient funds, if they would cause budget amounts to
20 be exceeded, or if they are simply over a maximum amount. Such automation could alternatively be provided as part of the user process, such that an approvals request is not even made.

In determining whether to approve a given
25 transaction, the approver should ideally be able to have a complete view of it, for example such that he/she can see exactly what is being purchased, rather than simply summary information such as total price and supplier. The preferred system provides for this by
30 combining the features of recording transmissions described above, with the approvals feature. The approvals request submitted at Step S298 is supplemented with a reference to the location in the database of the transaction information stored in Step
35 S284. The approver receives the location details in Step S316 and the system retrieves the transmissions

- 75 -

which constitute the transaction from the database, suitably displaying them such that the approver can consider them in making his/her determination of approval. Operation then continues normally at step S318. Clearly it is important that the recording step S284 takes place before the approval request is made at step S298, otherwise the recorded information will not yet be available. Since at step S284 the transaction will have been identified, but not yet completed (since it has not yet been approved), it is necessary for the database record made in step S284 to contain a flag which identifies the transaction as ``pending''. This flag can then be updated in Step S316 to show that the transaction has been approved or denied, or alternatively, if approval is denied the database record can be deleted since the transaction did not take place.

Security

The preferred system provides means to assign an appropriate security rating to the transmission in dependence on the identified nature of the data being transmitted. The security rating assigned may be set by the user of the system using the policy data to reflect his needs.

The simplest implementation of the policy data in this case is a list containing in a first column possible types of data, such as employee passwords, employer passwords, credit card numbers, banking details and so on, and containing in a second column, the desired encryption strength (in key bits for example) deemed appropriate for each data type. It will be appreciated that other ways of assigning security levels in dependence on the determined nature of the data may also be employed within the scope of the invention.

- 76 -

Figure 17 shows an example illustration of policy data defining appropriate encryption strengths for various types of data. The policy data takes the form of a number of key-value pairs arranged on separate branches of the policy data tree. The key specifies the type of data that is being transmitted such as passwords, credit card numbers, submitted key words and a general key for any other submitted data. The values that correspond to these keys are the encryption strength in bits that is deemed appropriate for the transmission of the data specified in the key. The key value pairs are arranged on several branches of the RequiredEncryptionLevel branch of the TransmittedDataSecurity branch of the policy data tree. Thus, in the example, it may be seen that passwords have a desired encryption strength of 40 bits, company credit card numbers and personal credit card numbers both have a desired encryption strength of 128 bits, submitted key words have a desired encryption strength of 40 bits and other submitted data requires no encryption.

The SubmittedKeywords branch refers to particular words or strings or text that have been designated as sensitive and requiring some form of encryption. These may be usernames, address information, financial information or pre-selected words such as 'Confidential' or Secret'. The submitted keywords may be detected by referring to a table or file in which they are stored.

Furthermore, each branch of the policy data may, instead of giving a general encryption strength, refer to a table in which different passwords or credit card numbers, for example, are listed along with corresponding encryption strengths specific to each password or number.

- 77 -

Once a security rating has been assigned, the plug-in module interrogates either the web browser to determine the security of the link that has been established by the web browser with the web server for transmission of that information, or in the case of an e-mail transmission, the encryption settings that the user or application have determined will be applied to the message. Typically, this will be the cryptographic strength of the encryption algorithm used to encode the data for transmission. Such transmission details are received by the web browser as part of the 'electronic handshake' from the web service provider.

A secure link is usually indicated in a Browser window by the presence of a closed padlock icon in the bottom right corner. A user can click on the icon to interrogate the level of security that has been provided by the handshake. In doing so they may receive a notification of the form ASSL secured. (128 bit). The first part of the notification describes the type of the encryption used while the second part describes the encryption strength. The plug-in module is implemented to automatically obtain this data from the browser so that it may be used to determine whether or not the security level is adequate for a proposed transmission. Similarly, in the case of an e-mail message, the plug-in module determines the encryption settings that the user or application has specified are to be used prior to transmission of the message.

The module compares the specified encryption strength with that of the link or message and depending on the result of the comparison performs one of the following actions:

a) If the security of the link is appropriate for the nature of the information being transmitted, the module allows the information to be transmitted;

- 78 -

b) If the security of the link is more than that required for the transmission of the information then, the module may either allow the information to be transmitted at that level of security, automatically
5 renegotiate with the web server and the web browser a new appropriate level of security and transmit the information at that level, or prompt the user that the present level of security is unnecessary and invite them to take action.

10 c) If the security of the link is not sufficient for the nature of the information being transmitted then, the module may either prevent the transmission from taking place and warn the user, automatically renegotiate with the web server and the web browser a
15 new appropriate level of security and then transmit the information at that level, or in the case of an e-mail automatically increase the encryption strength setting, or prompt the user that the present level of security is not sufficient and invite them to confirm that they
20 still wish for transmission to take place.

It will be appreciated that the plug-in module could be configured to respond to a difference in the determined desirable level of security and that being provided in a number of ways and that the actions
25 outlined above are only illustrative.

Further actions that may be taken by the system might include requesting a different web page to be down loaded to the user's machine or modifying the submitted field data such that sensitive information is
30 not transmitted.

The operation of a browser or e-mail plug in module for monitoring the data being transmitted by a user of the preferred system is illustrated in Figure 18, to which reference should be made. The module
35 begins operation at step S320 at point C in Figure 3, just prior to the transmission of the data to a web

- 79 -

server or at point B in Figure 5 just prior to transmission of an e-mail. Control then passes to step S322, in which the module parses the data about to be transmitted and searches for credit card numbers. A possible method for doing this was described earlier with reference to Figure 8. If no credit card number is detected in the data then control passes to step S314 in which the module searches for passwords in the data about to be transmitted. A method for doing this has been described above, with reference to Figure 6. If no password is found in the data, then control passes to step S316 in which the module searches for company account or purchase codes in the data. Recognising account or purchase codes may be achieved by storing the codes of the company in a file and attempting to match these codes with any strings of digits found in the outgoing data. If no account code is found then control passes to step S318, where the module searches for indications of other sensitive data in the data to be transmitted. Such indications will need to be defined in advance preferably in a separate file used for detection, and will be dependent on the requirements of the users of the preferred system. Examples might be specified keywords relating to projects the company is undertaking, project titles themselves, personal details address of the recipient of the data, or of the sender, or even the word 'confidential' or 'private' included in the data itself.

30 If no such indications are found that the data is sensitive and requires stronger protection before it is transmitted, then the transmission is allowed to proceed at the current level of encryption. This may mean that transmission takes place without any encryption being applied.

- 80 -

If however any of the checks in steps S322 to S328 reveal data that is deemed sensitive then control passes to step S332 in which a security rating is assigned to the detected data. This is achieved by
5 comparing the detected data with pre-determined entries in the policy data.

Each entry on the branch of the policy data has a pre-assigned level of encryption which is the minimum level that may be used for transmission of that data.
10 The entries in the table and the assigned level of encryption, as with all the policy settings, are decided by the company using the preferred system in dependence on their requirements. Assigning a security rating is then simply a matter of looking up password,
15 credit card number or other data in the policy data and reading off the corresponding rating. References to tables on a sub branch of the policy data may be used to allocate different encryption strengths to different passwords, credit card numbers and so on.

20 Once the appropriate security level has been determined in step S332, control passes to step S334 in which the module determines the level of encryption that has been negotiated with the web server to which the data is being transmitted, or is to be used by the
25 e-mail application prior to transmitting the message. This may be achieved by interrogating the web browser or e-mail application, or by setting encryption strength variables at the time the link is established or e-mail encryption requirements determined, both of
30 which will occur prior to transmission.

Control then passes to decision step S336 in which the desired level of security, i.e. the encryption strength, is compared with that determined in the previous step. If the desired level of encryption is
35 lower than or equal to that determined in step S334, then there is deemed to be enough protection for the

- 81 -

data to be transmitted and control passes to end step S330, where the module exits. Following step S330, control returns to either point C in Figure 3 or point B in Figure 5 depending on whether the module is
5 implemented in a web browser or an e-mail client. Transmission of the data may then proceed in the usual way.

If however in step S336, the desired level of encryption is greater than that currently set, then the
10 module does not allow transmission to go ahead until the proper level of encryption has been negotiated. Control passes to decision step S338 in which the module determines if it is able to increase the encryption strength, and if so passes control to step
15 S340 where a new stronger encrypted link is negotiated, or in the case of an e-mail a higher encryption strength set.

The highest level of encryption that is available depends on the software being used by both the web
20 server and the web browser, or in the case of an e-mail by the sending and receiving e-mail applications. There may then be cases in which the appropriate level of encryption is not available by one party and transmission of the data is never allowed to proceed.
25 Furthermore, certain types of data may be given a security rating that indicates that no level of encryption will ever be high enough to protect it, ie: preventing that data from ever being transmitted.

Having attempted to re-establish the link, or
30 changed the e-mail encryption settings, to a higher encryption strength, control passes back to step S334 to ensure the link or settings are now at a suitable strength. If the appropriate encryption level cannot be renegotiated in step S338 or an attempt to increase the
35 encryption strength at step S340 has not been successful, then it is deemed unsafe to transmit the

- 82 -

data, and control passes to end step S342 where the module exits. Following exit at step S342, control returns to point A in Figure 3, or to step S132 in Figure 5 'compose e-mail', for the user to reconsider
5 and edit or abort the transmission. A suitable message may also be displayed to the user explaining the reasons for the transmission being prevented.

The preferred system therefore provides a way of ensuring that transmission of data is as secure as
10 possible. It precludes the possibility of a user forgetting to secure a transmission, and negotiates a more appropriate level of security if that being used is not sufficient.

Web Browsers may provide similar facilities to
15 warn the user that user entered data is about to be sent over an insecure link or provide facilities to encrypt all messages by default. The preferred system however provides the ability to examine the content of data to be transmitted to determine its security
20 requirement, to allow or prevent transmission based on such security requirements, and on the determined security level of the link (encryption strength). It will be appreciated that the preferred system provides a significantly improved system for secure transmission
25 which reduces the possibility of human error.

Monitoring Out-bound E-mails for sensitive information

In addition to the problem of sensitive data being
30 intercepted by a third party between sender and recipient, organisations are at considerable risk from the deliberate release of sensitive information by their users. For example, the practice of
"electronically" stealing copies of confidential
35 documents, such as customer lists, before leaving the employ of an organisation is easily achieved, virtually

- 83 -

untraceable, and consequently widespread. All that is required is for the user to send the document to his own private email address for later retrieval. The document need not even be sent via the organisation's own e-mail system, since an internet mail service, such as "Hotmail" may be used, making tracing the unauthorised 'leak' virtually impossible by current means.

In addition to providing means to ensure that an appropriate level of encryption is applied to messages, the preferred system allows messages identified as potentially sensitive to be automatically redirected or copied to another destination without the user's knowledge. In determining whether to re-direct such messages the preferred system takes into account a number of factors including the identity of the sender, the identity of the intended recipient, the nature of the intended recipients' address, the nature of the message content, the nature and existence of any attachments to the message, the means by which the message is intended to be transmitted, and whether or not the message, and/or any attachments, are encrypted.

The nature of the message can be determined by scanning it for one or more keywords, or combinations of keywords, or by using standard 'natural language query' techniques. The nature of the intended recipients' address may be determined by reference to a list of known internet mail service domains. For example "hotmail.com", "yahoo.com" and "aol.com" are all predominantly used by individuals, not corporations. Similarly, the address may be examined for similarities to the sender's name, for example an email known to be sent by Fred Smith to the address "smith900@aol.com" could be considered suspicious by the inclusion of both "smith" and "aol.com" in the recipient address. Further examination of the message

- 84 -

may support the likelihood of it being an unauthorised release of confidential data, for example if the message consists only of file attachments with the message text and subject blank, an important clue since
5 the sender may be less likely to type text only he will read. The means by which the message is being transmitted is an important factor, for example a transmission sent using an internet mail service such as hotmail may be considered much more suspect than one
10 sent through the usual corporate e-mail system.

Indeed the 'uploading' of files to an internet mail service is potentially so suspicious that the preferred embodiment provides means to prohibit the uploading of files to such services altogether.

15 When re-directing mail, the preferred system adds additional text to the beginning of the mail, for example ``----Redirected Mail----``, together with the addresses of the originally intended recipients, such that the new recipient can determine that it was re-
20 directed to him/her, and to whom it was originally sent.

If the forwarded message is to be encrypted, the preferred system may re-direct the message encrypted or unencrypted to the third party. Preferably, the
25 sender's encryption key is transmitted with the message and means are provided for the third party to, unencrypt the message, if it has already been encrypted, and encrypt the message with the original sender's encryption key for transmission.

30 The preferred system also identifies incoming mail which has been re-directed (ie sent to a user who is not the originally intended recipient), by searching for the ``----Redirected Mail----`` text. Such mail can be flagged for attention by the new recipient for
35 example by using special icons, or by putting up a message box to notify him/her. Means may also be

- 85 -

provided to allow the new recipient to easily 'approve' the message and have it sent to the originally intended recipient(s). This can be achieved, for example by providing an 'Approve' button. If this button is
5 pressed, then the preferred system creates a new message in the same manner as if the user had pressed the normal 'Forward' button. However, instead of adding text to the message to note that it has been forwarded, in the case of the 'Approve' button, the
10 system extracts the originally intended recipient list from the message, and then strips out the redirection details to leave the message in its original state. The destination fields 'To', 'Cc' and 'Bcc' are then automatically filled with the extracted original
15 recipient addresses, and the 'From' field (which exists for every message even if it is not normally displayed) is filled with the original sender's name/address. The date/time field can also be adjusted to the date/time of the original message. The message
20 is then sent, either automatically, or when the user presses the 'Send' button. In this way, with the press of just one or two buttons, the re-directed mail can be approved and sent, and when delivered, will appear to have come from the original recipient as if
25 redirection had never taken place.

Sample policy data for controlling the operation of a plug-in module implemented to re-direct mail is shown in Figure 19 and a schematic illustration of the operation of such a plug-in module is shown in Figure
30 20. Figure 19 is a policy data tree having a number of branches which correspond to decision steps in Figure 20.

The plug-in module is initiated at step S350 which corresponds to point B in the illustration of the e-mail client operation in Figure 5. Once initiated, the
35 plug-in module steps through six steps determining

- 86 -

different particulars of the outbound e-mail message. Firstly, in step S351 the identity of the person sending the e-mail is checked against entries in a pre-determined list of names or addresses. E-mails from

5 users on this list are deemed to have the authority to transmit e-mails directly to the intended recipient's regardless of the contents of the message and regardless of the recipient. Anybody, not included on the list may have their e-mail re-directed or not

10 depending on its content. Thus, in decision step S351, if the user's name or address is found on the list, then control passes to step S364 where the e-mail is allowed to be transmitted without further interaction. If however, the user is not on the list, control passes

15 to step S352 for further checks. In step S352 the recipient of the outbound e-mail message is checked against look-up table s, specified on the Recipients branch of the policy data shown in Figure 19. In the following decision step S354, the text comprising the

20 e-mail message, and any attachments attached to the e-mail message, is compared with entries in a look-up table t. Table t is referred to on the Keywords branch of the policy data and contains words that indicate the nature of the e-mail message might be sensitive to the

25 company. In the next step S356, the e-mail message is checked to see whether or not it is to be encrypted. It will be remembered that encryption does not occur until the moment the e-mail is transmitted, so at this stage the e-mail will only be flagged for encryption. In the

30 next decision step S358, it is determined whether or not the e-mail message contains attachments, and in the following decision step S360 whether the e-mail message contains text to accompany the attachments, that is whether the body text of the e-mail message is blank.

35 In any of the decision steps S352, S354 or S362 where a look-up table is consulted, a match between an

- 87 -

entry in the look-up table and an entry in the e-mail indicates that the e-mail should be re-directed to a third party for checking before being sent out of the company. For example, if in step S354, the e-mail is
5 found to contain the words 'confidential' or 'secret', that will be enough to warrant the e-mail being checked by a third party before it is delivered to its intended recipient. This ensures that sensitive information is not sent out of the company without the company being
10 aware. Control flows therefore from these steps to step S364 where text indicating that the e-mail has been redirected is added to the message, and the recipient address is changed to that of the verifying party to whom the re-directed message should be transmitted.
15 Control then passes to step S364 where the e-mail is transmitted. If the e-mail message has been marked for re-direction, in step S364 it will of course be transmitted to the verifying party for review rather than the original recipient of the message.

20 In decision step S356 if encryption of the message is detected then that is deemed enough to warrant re-directing the message to a third party for review. Accordingly, if the message is to be encrypted, control passes from step S356 to step S364 where the message is
25 modified for re-direction. In the case of a message flagged for encryption, the encryption flag is preferably cleared, such that the message is re-directed without it being encrypted, allowing the new recipient to read it. The re-direction text, added to
30 the message, preferably also includes the encryption key (which is generally the intended recipient's public key, and therefore not sensitive) such that the message can be re-encrypted prior to transmission.

Alternatively, the intended recipient's entire
35 digital certificate could be included with the re-directed message. The public key, or certificate, as

- 88 -

the case may be, would then be removed by the automated approval process described above.

If in step S358 the e-mail does not contain attachments then the e-mail is deemed to be unlikely to contain any documents or files of potentially sensitive information, and the e-mail is thus allowed to be transmitted without any further intervention in step S364. If the e-mail does contain attachments, and in step S360, the e-mail is determined not to contain any text entered into the body or the subject of the message, then the message is recognised as one that is likely being sent to a different account of the sender. The e-mail is then forwarded in step S362 and S364 to a third party for checking.

Figure 21 shows the operation of the plug-in module to block the up-load of information from the company's computer system to an external site. A simple two stage check is used, comprising a check of the external site address in S372, following initiation at step S370, and a check of sensitive keywords in the information being uploaded in step S374. Providing the external site address is not found to be a prohibited site in step S372, and no sensitive keywords are detected in step S374, then up-load is allowed to proceed in step S376. Otherwise, the up-load is blocked in step S378.

The policy data for controlling the operation of the plug-in module for selectively blocking upload of information is straightforward, and is shown at the bottom of Figure 19.

In this way, outbound transmissions containing sensitive information that are being transmitted for reasons that are not in the company's interests can be verified before transmission, and transmission can be prevented if necessary.

- 89 -

Managing Forwarding of E-mails

E-mail applications provide means to "Forward" incoming mail to one or more users. The user typically clicks a "Forward" button, which causes a
5 copy of the incoming mail to be automatically entered into the mail composition window, as if the user had keyed it himself. All that is then required is for the user to enter the names of the intended recipients of the forwarded mail and to click the "Send" button.
10 This is a useful feature allowing a user to easily share the contents of a received e-mail with others.

A problem can arise however in the case of mail containing sensitive information, particularly if the sensitive nature of the mail is not immediately
15 apparent, for example if the sensitive information appears lower down in the mail, such that the user is required to scroll down the viewing window in order to read it. Users frequently forward e-mails after scanning only the first few lines, or in some cases
20 only the subject line, particularly when they have large amounts of e-mails to process. As a consequence, sensitive information is often unintentionally disclosed, both within the organisation, and more dangerously, outside it. There have been several high
25 profile cases where considerable sums have been lost as a result.

The preferred system therefore provides means of controlling the forwarding of e-mails. Such controls include providing warnings to the user before a
30 forwarded e-mail is transmitted and preventing the e-mail from being forwarded altogether. Means to approve the e-mail prior to transmission, or to re-direct it to another user, as described above could also be provided.

35 Preferably, forwarding occurs in accordance with the content of the e-mail, and the addresses of the

- 90 -

recipients to whom it has already been sent. For example, the sensitive nature of the e-mail can be determined by a number of methods including scanning it for keywords such as ``confidential'', or by checking if
5 the sensitivity attribute has been set to ``Personal'', ``Private'' or ``Confidential''. Means for the original composer of the message to mark as unsuitable for future forwarding can also be provided by inserting pre-defined keystings, such as ``<NOFORWARD>''
10 (preventing any forwarding), or ``<NOFORWARDEXTTERNAL>'' (preventing forwarding outside the organisation). Such means could also be provided in the form of an additional attribute to the message.

In addition to content based factors, the
15 preferred system queries the list of previous recipients of the message. If the e-mail has already been sent out of the organisation by the original composer, then it may for example be considered safe to allow it to be further forwarded externally.
20 Similarly, if the original e-mail was only sent to a single recipient, then it may be determined that the original composer did not intend it to be widely circulated and an appropriate warning could be issued. The course of action in response to the factors
25 described can be determined in accordance with policy..

The fact that an e-mail that is about to be transmitted is a forwarded e-mail can be readily determined by scanning the e-mail for keystings such as ``-----Original Message-----'' which is automatically
30 added at the beginning of the original mail by the E-Mail program. Similarly, the list of previous recipients can be determined by scanning for the keystings ``To:'' and ``Cc:'' which follow the original message string. The list of recipients is found
35 immediately after these keystings. Internal and external recipients can be readily distinguished by

- 91 -

reference to the domain name (if any). For example, a recipient name shown as 'Fred Smith' is typically internal, whereas 'fsmith@xyz.com' is typically external.

5 Policy data for instructing the operation of a plug-in module implemented to control forwarding of e-mail messages is shown in Figure 22, and the operation of such a module in Figure 23.

The policy data tree contains a number of sub-
10 branches specifying parameters against which command values can be set. For example, the first sub-branch 'PreventAll' when set to 'YES' instructs the module to prevent all e-mails from being forwarded. The next sub-branch WarnAll when set to YES, requires the module to
15 give a warning to the e-mail client user whenever an e-mail is about to be forwarded. The next two sub-branches PreventExternal and WarnExternal contain corresponding parameters for external e-mails only and allow the user of the e-mail client to distinguish between rules
20 affecting forwarding of e-mails within the company and forwarding of e-mails to people outside the company. The 'PreventKeywords' branch specifies a look-up table in which keywords indicative of sensitive information are stored. Such keywords may be predefined key
25 strings such as <NOFORWARD> or <NOFORWARDEXTERNAL>, or one or more pre-determined words.

The e-mail is scanned before transmission and if such a keyword is found in the text of the e-mail or in any attachments of the e-mail, the e-mail will not be
30 allowed to be forwarded. The last two sub-branches PreventIfNotSentExternally when set to YES will inhibit the transmission of the forwarded e-mail outside of the company if it has not been transmitted outside of the company before. In practice the
35 forwarded e-mail may be transmitted to all of the recipients internal to the company, with the external

- 92 -

recipients simply being deleted from the recipient list or alternatively, before transmission, the user may be required to modify the address list such that it does not contain external recipients.

5 Finally, the parameter set on the PreventIfSingleRecipient branch when sent to YES prevents forwarding of any e-mail messages if the original recipient of the message was a single person.

10 The operation of the plug-in module is illustrated in Figure 23. The plug-in module is initiated at step S380, again at point B in Figure 5. In decision step S382, the e-mail is given a preliminary scan to see whether it contains the keystring ``---Original message---'', since this key string is automatically
15 added at the beginning of the original mail by the e-mail program when generating a message for forwarding. If the e-mail message does not contain this keystring, then it may be deduced that the e-mail message is an original message and is not being forwarded and the
20 message may be transmitted in step S384. If however in step S382 the message is found to contain the key string ``---Original message---'', then the e-mail message is clearly a forwarded message, and the module takes further steps to determine whether or not to
25 allow the forwarded message to be transmitted. Control then flows to step S386 in which the recipients of the forwarded message are checked to see if any are external to the on-line company. If there is an external recipient, then in step S388, the plug-in
30 module scans the e-mail message to see whether the e-mail has been forwarded to a recipient outside of the company before. If it has not then the e-mail message is prevented from being forwarded in step S390 and the user of the e-mail client is notified. However, if the
35 e-mail message has been sent out of the company before, then a warning is issued to the user in step S392,

- 93 -

following which the e-mail message may either be transmitted by the user or may be returned to the user for him to revise the intended recipients.

If in step S386 the forwarded message is not to be sent to an address outside of the company, then control passes to step S394, in which it is determined whether the user was the only recipient of the original message. If he was, then it may be the case that the original composer of the message did not intend for it to be transmitted to a wide audience. Accordingly, control flows to step S390 where the transmission of the forwarded e-mail message is blocked. This is in accordance with the policy data shown in Figure 22, which specifies taking such action. Alternatively, a warning may be issued to the user attempting to forward the message.

The last check made is in decision step S396 in which the contents of the message and any attachments attached to it are compared with entries in a keywords table. If any matches between entries in the e-mail and those in the table are found then the message is considered to contain sensitive information and is not forwarded. The module then ends at step S390. If no sensitive keywords are found then the e-mail is allowed to be forwarded in step S384.

Managing Signing of Out-bound Transmissions

The ability to use a Digital Certificate to sign a message is clearly valuable to the recipient of the message in establishing the identity of the sender, and to both parties in ensuring the message is not tampered with. The sender of the message should however be aware that a digitally signed message potentially constitutes a binding contract, which cannot be denied or revoked once sent. It is therefore imperative that care is used when digitally signing a document, just as

- 94 -

care should be exercised when appending a written signature to a paper document. E-mail applications, such as Microsoft Corporation's 'Outlook' provide means to digitally sign messages automatically, and while this may be valuable in confirming the identity of the sender to the recipient, for the reasons described above, it is also potentially dangerous, requiring the user to take extreme care with the message content.

10 The preferred system provides means to control the signing of outgoing messages, according to policy data. Sample policy data is shown in Figure 24. Such controls include:

- enforcing that a message be signed;
- 15 suggesting to the user that a message be signed;
- suggesting to the user that a message should NOT be signed; and
- preventing a message being signed.

In determining which course of action to take, the preferred system takes into account a number of factors, including the nature of the message content (including any attachments), the identity of the intended recipient and/or his organisation, the identity of the sender, the nature of the digital certificate being used (if the message has already been flagged for signature), and the nature of the digital certificate(s) available to sign the message.

The nature of the message can be determined by scanning it for one or more keywords, or combinations of keywords, or by using standard 'natural language query' techniques. Similarly, the nature of the intended or available digital certificates can be determined by reference to the issuer and type of certificate.

35 Figure 25 illustrates the operation of a plug-in module for ensuring that an e-mail is appropriately

- 95 -

digitally signed or not. The module is initiated in the step S400, at point B in the operation of the e-mail client illustrated in Figure 5. Control then passes to decision step S402 in which the outbound e-mail is scanned to see whether it has been marked for signature already. Actual 'signing' of the message will not occur until directly before transmission. If it is not marked for signature control passes to S404 where the module consults a Recipients look-up table (table f) to determine whether the Recipient of the outbound e-mail has been identified as one to whom e-mails should always be digitally signed. If the Recipient is in table f, then control passes to step S406 where the user of the e-mail client is notified that the e-mail will not be sent unless it is digitally signed. Alternatively, the plug-in module may automatically digitally sign the e-mail using the e-mail author's digital certificate.

If in step S404, the recipient of the outbound e-mail is not found in the look-up table, then control passes to decision step S408 where the KeywordsTable on the EnforceSigning branch of the policy tree is then consulted. Should any of the keywords in table g be found in the text of the e-mail or in any of the attachments of the e-mail, digitally signing of the e-mail will be required, and control will pass to step S406 as before. It will be appreciated that decision steps S404 and S408 correspond to look-up commands for consulting look-up Recipients and Keywords tables on the EnforceSigning branch of the policy data.

Such Keywords may be predetermined words like 'Confidential', 'Secret', 'Contract', 'Quotation', 'Order' and so on as illustrated in Figure 24.

If the recipients of the e-mail are not to be found in table f, and the e-mail does not contain

- 96 -

Keywords specified in table g, then control passes to decision step S410, corresponding to a look-up command on the SuggestSigning branch of the sample policy data. In decision step S410 the recipient's address is

5 checked against those in the look-up table h to determine whether signing of the e-mail is advised. If the recipient's name is found in table h, then control passes to step S412, where the user of the e-mail client is notified of the desirability to digitally

10 sign the outbound e-mail message. However, the user of the e-mail client is not required to digitally sign the e-mail message and the e-mail may therefore be transmitted without signature if the user so chooses. Following decision step S410, if the recipient's name

15 is not found in table h, control passes to decision step S414, where as before, the e-mail text is searched for a number of keywords which might indicate that it contains sensitive data and requires a digital signature. Depending on whether the e-mail contains

20 such sensitive keywords, the user of the e-mail client is either notified in step S412 of the desirability to digitally sign the message, or alternatively, the e-mail message is transmitted without signature in step S416.

25 If in step S402 following initiation of the plug-in module, the e-mail is found to have been marked for signature then control passes to decision step S418. In decision step S418 the plug-in module consults look-up table m in the DenySigning branch specified under the

30 DenySigning branch of the policy data. Table m is specified on the CertificatesUsed branch under the DenySigning branch, and specifies the issuer, type, certificate number or signing key of digital

35 digital certificate or signing key that is to be used to sign the outbound e-mail be found in table m, then

- 97 -

further checks as to Recipient and the nature of the outbound e-mail will be made to determine whether it is appropriate to sign the message or not. Control passes to step S420, where the recipient of the outbound e-mail is checked against Recipient table n and then to decision step S422 where the text of the e-mail and any attachments are scanned for various keywords. If in either of decision steps S420 or S422, the recipient or any text in the message match that defined in the look-up tables control passes to step S424 where transmission of the e-mail is blocked. The user of the e-mail client may then be returned to the e-mail text entry stage, and he may be required to re-transmit the message without digitally signing it.

15 If in either of steps S418 or S422 the certificate or signing key is not found to be one of interest, and the text of the e-mail is found not to contain any sensitive words, control passes to step S426 which corresponds to the first sub-branch on the

20 SuggestNotSigning branch of the policy data tree. As for the DenySigning branch of the policy data tree, the three decision steps S426, S428 and S430, correspond to the look-up commands for checking the digital

25 certificate or signing key used with the e-mail, the recipient of the outbound e-mail and the text of the outbound e-mail against predetermined sensitive entries in look-up tables j, k and l respectively. If in step S426 the certificate used to sign the outbound e-mail is found to be one of interest, and in either of steps

30 S428 or step S430, the recipient or the text of the outbound e-mail is found to match entries in the specified look-up tables, control will flow to step S432 in which the user of the e-mail client will be notified of the desirability of not digitally signing

35 the outbound e-mail message. The user however may still

- 98 -

be free to send the signed e-mail message if he so wishes.

If in either of decision steps S426, S428 and S430 no match is found then the e-mail is sent normally in
5 step S416.

Instant Messaging and Telephony Applications

In addition to browser and email activity, additional applications such as Instant Messaging (also
10 known as 'chat') and digital telephony applications such as 'Voice Over IP') are now becoming popular in business situations. Instant Messaging standards are defined in RFC's 2778 and 2779 and by the IETF SIMPLE working group. Voice Over IP standards are defined in
15 ITU-T Recommendation H.323 (1998). Many aspects of the present invention can be applied to the data transmitted and received by such applications. Instant Messaging is conceptually similar to a series of sent and received emails, except that the
20 'conversation' is carried out 'live' with both parties present throughout. For the purposes of the present invention however the procedures are identical. Figure 5 in the drawings can represent Instant Messaging by replacing the word 'E-mail' in steps
25 S122, S124, S132 and S134 with the word 'Message'. The 'E-mail Server' description 95 is replaced with 'Message Relay'. The preferred embodiment is arranged so as to intercept at points A and B as before, by providing a plug-in to the Internet Messaging
30 application, or by developing an Instant Messaging application which contains the plug-in functionality. Alternatively, it will be appreciated that the interception could occur at the protocol level, intercepting network packets before they leave the user
35 machine, or as they arrive at the user machine, or at some intermediate point on the network.

- 99 -

Voice Over Internet Protocol (VOIP) is conceptually similar to Instant Messaging, except that the message content consists of digitised speech, which is encoded and transmitted immediately. Message
5 content analysis is impractical, however means to record the message and to place controls at the 'call' level are viable and are implemented in a similar manner, either as a plug-in to the voice messaging application, or at the network protocol level, either
10 within the user machine, or at some intermediate point on the network.

Although the implementation of the preferred system has been described with reference to plug-in modules for existing applications, the invention may be
15 implemented by providing web browsers, e-mail clients, instant messaging applications or Voice Over IP applications in which the functionality of the plug-in modules described here is already coded from the outset.

20

25

30

35

- 100 -

Claims

1. An information management system comprising:
a plurality of client workstations adapted for connection to a computer network, each workstation having a memory;

a data repository arranged to receive data from each of said client workstations;

an application stored in said memory of each workstation for transmitting outbound data to said network and receiving inbound data from said network;

policy data defining rules for the recording of data that may comprise part of a transaction conducted between a client workstation and a third party across said computer network; and

an analyser, said analyser being operable in conjunction with said policy data to analyse at least one of said outbound data and said inbound data, to identify the existence of a transaction occurring between a client workstation and a third party by analysing said outbound or said inbound data, and to cause transaction data that is all or part of said outbound data or said inbound data related to an identified transaction to be stored in said data repository.

2. The system of claim 1, wherein said analyser is operable to determine whether a secure link has been negotiated between said application and a remote site on said network, and to identify the existence of a transaction if said outbound data or said inbound data is transmitted on a secure link.

3. The system of claim 1 or 2, wherein said network is the Internet, and said rules of said policy data define the addresses of non-eCommerce web sites

- 101 -

and/or non-eCommerce e-mail accounts, said analyser being operable to disregard any transactions that are identified between a client workstation and a non-eCommerce web site and/or e-mail account such that no transaction data related to a transaction made to a non-eCommerce web sites or a non-eCommerce e-mail account is stored in the data repository.

4. The system of any one of claims 1 to 3, wherein said analyser is operable to identify the existence of a transaction by reference to said rules of said policy data, said rules of said policy data defining the addresses of known eCommerce locations.

5. The system of any one of claims 1 to 4, wherein said analyser is operable to identify credit card numbers, and to identify the existence of a transaction by identifying credit card numbers in said outbound data or inbound data.

6. The system of any one of claims 1 to 5, wherein said analyser is operable to identify the existence of a transaction by reference to said rules of said policy data, said rules of said policy data defining one or more of pre-determined digital certificates, account codes, pre-determined keywords, pre-determined names and addresses and embedded codes.

7. The system of any one of claims 1 to 6, wherein said analyser is operable to identify embedded codes in said inbound data, said embedded code having been placed in said inbound data to identify it as transaction data.

8. The system of any one of claims 1 to 7, wherein said analyser is operable to identify

- 102 -

electronic receipts, and to identify the existence of a transaction by identifying an electronic receipt in said outbound or inbound data.

9. The system of any one of claims 1 to 8, wherein said analyser is operable to record a pre-determined number of subsequent transmissions of said outbound data or said inbound data following an identification of the existence of a transaction by said analyser, providing that the address or organisation to which the subsequent transmissions are sent, or from which they are received, is the same as the address or organisation to which the outbound data was sent or from which the inbound data was received and in which the existence of a transaction was identified.

10. The system of claim 9, wherein said analyser is operable to detect one or more indicators of the nature of the transaction, and said rules of said policy data define the number of subsequent transmissions of said outbound data and said inbound data that are to be recorded in said data repository based on the identified nature of the transaction.

11. The system of claim 9, wherein said rules of said policy data define the number of subsequent transmissions of said outbound and said inbound data that are to be stored in said data repository in dependence on the indicator by which the existence of a transaction was identified.

12. The system of any one of claims 1 to 11, wherein said analyser is operable to record all subsequent transmissions of said outbound data or said inbound data that occur within a pre-determined amount

- 103 -

of time following an identification of the existence of a transaction by said analyser, providing that the address or organisation to which the subsequent transmissions are sent, or from which they are received, is the same as the address or organisation to which the outbound data was sent or from which the inbound data was received and in which the existence of a transaction was identified.

13. The system of claim 12, wherein said analyser is operable to detect one or more indicators of the nature of the transaction, and said rules of said policy data define the amount of time for which all subsequent transmissions of said outbound data and said inbound data are to be recorded in said data repository based on the identified nature of the transaction.

14. The system of claim 12, wherein said rules of said policy data define the amount of time for which subsequent transmissions of said outbound and said inbound data are to be stored in said data repository in dependence on the indicator by which the existence of a transaction was identified.

15. The system of any one of claims 1 to 14, wherein said analyser is further operable to identify the completion of a transaction by analysing said outbound data or said inbound data, and to cause all or part of said outbound data transmitted by said application and all or part of said inbound data received by said application after said analyser has identified the existence of a transaction and before said analyser has identified the completion of a transaction to be stored in said data repository.

- 104 -

16. The system of claim 15, wherein said analyser is operable to identify subsequent related data in said outbound data transmitted by said application and said inbound data received by said application after said analyser has identified the completion of a transaction, and to cause said subsequent related data to be stored in said data repository with said transaction data already identified.

17. The system of claim 16, wherein said analyser is operable to identify said subsequent related data by identifying common indicators in both said transaction data already identified and said outbound data transmitted by said application and said inbound data received by said application after said analyser has identified the completion of a transaction, said common indicators being one or more of the address of the location to which said outbound data is transmitted or from which said inbound data is received, part of the data path to the location to which said outbound data is transmitted or from which said inbound data is received, account code or reference numbers.

18. The system of any one of claims 1 to 17, wherein said application is operable such that a user of said application can indicate said outbound and said inbound data that is related to a transaction, said analyser being operable to identify said outbound and said inbound data so indicated.

19. The system of any one of claims 1 to 18, wherein said application is operable to store all of said outbound data and said inbound data in said memory of said workstation as previous data, irrespective of whether it may or may not be part of a transaction and, said analyser is operable, if the existence of a

- 105 -

transaction is identified, to retrieve a pre-determined amount of previous data from said outbound data and said inbound data stored in said memory of said workstation, and to cause said previous data to be stored in said data repository with said transaction data.

20. The system of claim 19, wherein said rules of said policy data specify the amount of previous data that is to be retrieved in dependence on the indicator by which the existence of a transaction is identified.

21. The system of claim 19 or 20, wherein said network is the Internet and said application is a web browser, said web browser being operable to store each web page that is viewed by said web browser in memory as previous data.

22. The system of claim 21, wherein said rules of said policy data specify the number of web pages that are to be retrieved from those previously stored in memory in dependence on the indicator by which the existence of a transaction is identified.

23. The system of any one of claims 1 to 22, wherein said application is operable to store all of said outbound data and said inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction and, said analyser is operable, if the existence of a transaction is identified, to identify, in said previous data, earlier relevant data that is related to said transaction data already identified, and to cause said earlier relevant data to be stored in said data repository with said transaction data.

- 106 -

24. The system of claim 23, wherein said analyser is operable to identify said earlier relevant data in said previous data, by identifying common indicators in both said transaction data and said outbound data and said inbound data previously stored in said memory of said workstation, said common indicators being one or more of the address of the location to which said outbound data is transmitted or from which said inbound data is received, part of the data path to the location to which said outbound data is transmitted or said inbound data is received, account code or reference number.

25. The system of any one of claims 1 to 24, wherein said application is operable to store all of said outbound data and said inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction, and is further operable such that, if said analyser identifies the existence of a transaction, a user of said application can select earlier relevant data from said previous data, said earlier relevant data selected by the user being stored in said common data repository together with said transaction data.

26. The system of any one of claims 1 to 25, wherein said analyser is operable, once it has identified the existence of a transaction, to determine the nature of said transaction by analysing the content of said outbound and inbound data, and said rules of said policy data define how said transaction data is to be stored in said data repository in dependence on the nature of the transaction determined by said analyser, said transaction data being stored in said database according to said determination and said rules of said policy data.

- 107 -

27. The system of claim 26, wherein said analyser is operable to determine the nature of the transaction by identifying in said outbound data and said inbound data one or more indicators, said indicators being defined in said rules of said policy data, and being one or more of: the address of the network location to which said data that may be part of a transaction is transmitted or from which it is received; part of the data path to the network location to which said transaction data is transmitted or from which it is received; account codes; reference numbers; credit card numbers; digital certificates and pre-determined keywords.

28. The system of any one of claims 1 to 27, wherein said analyser is operable to identify, once the existence of a transaction has been identified, one or more indicators of the nature of said transaction, said transaction data being stored in said data repository such that it is organised by said one or more indicators to form a record.

29. The system of claim 28, wherein said rules of said policy data define said one or more indicators of the nature of a transaction, said indicators being one or more of: the address of the location to which said transaction data is transmitted or from which it is received; part of the data path to the location to which said transaction data is transmitted or from which it is received; account codes, reference numbers, credit card numbers, digital certificates and pre-determined keywords.

30. The system of any one of claims 1 to 29, wherein the analyser is operable to determine the

- 108 -

length of time taken to complete a transaction, and cause this to be stored in the data repository with the transaction data.

31. The system of any one of claims 1 to 30, wherein the analyser is operable to determine the number of keystrokes or mouse-clicks made in order to complete the transaction, and cause this to be stored in the data repository with the transaction data.

32. The system of any one of claims 1 to 31, wherein the analyser is operable to determine the length of time spent waiting for data to be downloaded to said application during said transaction, and cause this to be stored in the data repository with the transaction data.

33. The system of any one of claims 1 to 32, wherein the analyser is operable to determine the length of time the web server takes to respond to outbound messages sent from said application, and cause this to be stored in said data repository.

34. The system of any one of claims 1 to 33, wherein the analyser is operable to determine the length of time a user of said application takes to transmit a response to data downloaded to said application during said transaction, and cause this to be stored in the data repository with the transaction data.

35. The system of claim 34, wherein the analyser is operable to determine from the length of time a user of said application takes to transmit a response to data downloaded to said application, the length of time the user spends entering the data that

- 109 -

forms the response, and cause this to be stored in the data repository with the transaction data.

36. The system of any one of claims 1 to 35, wherein said data repository is accessible by one or more of an accounts application, an order processing application or other transaction management application.

37. The system of any one of claims 1 to 36, wherein any data transmitted to said data repository is encrypted before it is transmitted to said data repository.

38. The system of anyone of claims 1 to 37, wherein any data stored in said data repository is encrypted.

39. The system of any one of claims 1 to 38, wherein said analyser is located on each of said one or more workstations.

40. The system of any one of claims 1 to 39, wherein said application is a web browser.

41. The system of claim 40, wherein said analyser is a plug-in module of said web browser.

42. The system of claim 41, wherein said web browser is Microsoft's Internet Explorer and said analyser is a Browser Helper Object.

43. The system of any of claims 1 to 39, wherein said application is an e-mail client.

- 110 -

44. The system of claim 43, wherein said analyser is a plug-in module of said e-mail client.

45. The system of claim 44, wherein said e-mail client is Microsoft's Outlook e-mail client and said analyser is a Microsoft Exchange client extension.

46. The system of any of claims 1 to 39, wherein said application is an Instant messaging application.

47. The system of claim 46, wherein said analyser is a plug-in module of said Instant messaging application.

48. The system of any of claims 1 to 38, wherein said network comprises a server, and said analyser is located at a point on said network intermediate said one or more work stations and said server, or said analyser is located at said server.

49. The system of any one of claims 1 to 48, wherein said computer network to which said one or more workstations are adapted for connection is a public computer network, and wherein said one or more workstations together form a private computer network.

50. The system of any one of claims 1 to 49, further comprising a supervisor workstation, said policy data being accessible by said supervisor workstation, such that a user of said supervisor workstation can edit said policy data.

51. A method of managing information comprising the steps of:

- 111 -

providing a plurality of client workstations adapted for connection to a computer network, each workstation having a memory;

providing a data repository arranged to receive data from each of said client workstations;

providing an application stored in said memory of each workstation for transmitting outbound data to said network and receiving inbound data from said network;

providing policy data defining rules for the recording of data that may comprise part of a transaction conducted between a client workstation and a third party across said computer network; and

analysing, at least one of said outbound data and said inbound data to identify, with reference to said rules of said policy data, the existence of a transaction occurring between a client workstation and a third party; and

storing transaction data that is all or part of said outbound data or said inbound data related to an identified transaction in said data repository.

52. The method of claim 51, wherein in said analysing step the existence of a transaction is identified by determining whether a secure link has been negotiated between said application and a remote site on said network, and by determining whether said outbound data or said inbound data is transmitted on that link.

53. The method of claim 51 or 52, wherein said network is the Internet, and said rules of said policy data define the addresses of non-eCommerce web sites and/or non-eCommerce e-mail accounts, wherein said analysing step includes disregarding any transactions that are identified between a client workstation and a

- 112 -

non-eCommerce web site and/or e-mail account such that no transaction data related to a transaction made to a non-eCommerce web site or a non-eCommerce e-mail account is stored in the data repository.

54. The method of any of claims 51 to 53, wherein said analysing step includes identifying the existence of a transaction by reference to said rules of said policy data, said rules of said policy data defining the addresses of known eCommerce locations.

55. The method of any of claims 51 to 54, wherein said analysing step includes identifying credit card numbers, and the existence of a transaction is identified by identifying credit card numbers in said outbound data or inbound data.

56. The method of claims 51 to 55, wherein in said analysing step the existence of a transaction is identified by reference to said rules of said policy data, said rules of said policy data defining one or more of pre-determined digital certificates, account codes, pre-determined keywords, pre-determined names and addresses and embedded codes.

57. The method of any of claims 51 to 56, wherein said analysing step includes detecting an embedded code in said inbound data, said embedded code having been placed in said inbound data to identify it as transaction data.

58. The method of claims 51 to 57, wherein in said analysing step, the existence of a transaction is identified by identifying an electronic receipt in said outbound or inbound data.

- 113 -

59. The method of any of claims 51 to 58, further comprising the step of recording a pre-determined number of subsequent transmissions of said outbound data or said inbound data following an identification of the existence of a transaction in said analysing step, providing that the address or organisation to which the subsequent transmissions are sent, or from which they are received, is the same as the address or organisation to which the outbound data was sent or from which the inbound data was received and in which the existence of a transaction was identified.

60. The method of claim 59, wherein said analysing step includes detecting one or more indicators of the nature of the transaction, and said rules of said policy data define the number of subsequent transmissions of said outbound data and said inbound data that are to be recorded in said data repository based on the identified nature of the transaction.

61. The method of claim 59, wherein said rules of said policy data define the number of subsequent transmissions of said outbound and said inbound data that are to be stored in said data repository in dependence on the indicator by which the existence of a transaction was identified.

62. The method of any of claims 51 to 61, further comprising the step of recording all subsequent transmissions of said outbound data or said inbound data that occur within a pre-determined amount of time following an identification of the existence of a transaction in said analysing step, providing that the address or organisation to which the subsequent transmissions are sent, or from which they are

- 114 -

received, is the same as the address or organisation to which the outbound data was sent or from which the inbound data was received and in which the existence of a transaction was identified.

63. The method of claim 62, wherein said analysing step includes detecting one or more indicators of the nature of the transaction, and said rules of said policy data define the amount of time for which all subsequent transmissions of said outbound data and said inbound data are to be recorded in said data repository based on the identified nature of the transaction.

64. The method of claim 62, wherein said rules of said policy data define the amount of time for which subsequent transmissions of said outbound and said inbound data are to be stored in said data repository in dependence on the indicator by which the existence of a transaction was identified.

65. The method of any of claims 51 to 64, wherein said analysing step includes identifying the completion of a transaction by analysing said outbound data or said inbound data, and said storing step includes storing all or part of said outbound data transmitted by said application and all or part of said inbound data received by said application, after the existence of a transaction has been identified and before the completion of a transaction has been identified, in said data repository.

66. The method of claim 65, wherein said analysing step includes identifying subsequent related data contained in said outbound data transmitted by said application and said inbound data received by said

- 115 -

application after the completion of a transaction, and said storing step includes storing said subsequent related data in said data repository with said transaction data already identified.

67. The method of claim 66, wherein said analysing step includes identifying said subsequent related data by identifying common indicators in both said transaction data already identified and said outbound data transmitted by said application and said inbound data received by said application after the completion of a transaction has been identified,

said common indicators being one or more of the address of the location to which said outbound data is transmitted or from which said inbound data is received, part of the data path to the location to which said outbound data is transmitted or from which said inbound data is received, account code or reference numbers.

68. The method of any of claims 51 to 67, wherein said application is operable such that a user of said application can indicate said outbound and said inbound data that is related to a transaction, said analysing step including identifying indicated outbound and inbound data.

69. The method of any of claims 51 to 68, further comprising the step of storing all of said outbound data and said inbound data in said memory of said workstation as previous data, irrespective of whether it may or may not be part of a transaction and, said analysing step includes retrieving a pre-determined amount of previous data from said outbound data and said inbound data stored in said memory of said workstation if the existence of a transaction is

- 116 -

identified, and said storing step includes storing said previous data in said data repository with said transaction data.

70. The method of claim 69, wherein said rules of said policy data specify the amount of previous data that is to be retrieved in dependence on the indicator by which the existence of a transaction is identified.

71. The method of claim 69 or 70, wherein said network is the Internet and said application is a web browser, said web browser being operable to store each web page that is viewed by said web browser in memory as previous data.

72. The method of claim 71, wherein said rules of said policy data specify the number of web pages that are to be retrieved from those previously stored in memory in dependence on the indicator by which the existence of a transaction is identified.

73. The method of any of claims 51 to 72, further comprising the step of storing all of said outbound data and said inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction and, said analysing step includes identifying, in said previous data, earlier relevant data that is related to said transaction data already identified, and said storing step includes storing said earlier relevant data in said data repository with said transaction data.

74. The method of claim 73, wherein said analysing step includes identifying said earlier relevant data in said previous data, by identifying common indicators in both said transaction data and

- 117 -

said previous data, said common indicators being one or more of the address of the location to which said outbound data is transmitted or from which said inbound data is received, part of the data path to the location to which said outbound data is transmitted or said inbound data is received, account code or reference number.

75. The method of any of claims 51 to 74, further comprising the steps of storing all of said outbound data and said inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction; and

if the existence of a transaction is identified, providing a user of said application with a selector for selecting earlier relevant data from said previous data, and wherein said storing step includes storing said earlier relevant data selected by the user in said data repository together with said transaction data.

76. The method of any of claims 51 to 75, wherein said analysing step includes, once the existence of a transaction has been identified, determining the nature of said transaction by analysing the content of said outbound and inbound data, said rules of said policy data defining how said transaction data is to be stored in said data repository in dependence on the nature of the transaction determined in said analysing step, said transaction data being stored in said database according to said determination and said rules of said policy data.

77. The method of claim 76, wherein said analysing step includes determining the nature of the transaction by identifying in said outbound data and

- 118 -

said inbound data one or more indicators, said indicators being defined in said rules of said policy data, and being one or more of: the address of the network location to which said data that may be part of a transaction is transmitted or from which it is received; part of the data path to the network location to which said transaction data is transmitted or from which it is received; account codes; reference numbers; credit card numbers; digital certificates and pre-determined keywords.

78. The method of any of claims 51 to 77, wherein said analysing step includes identifying, once the existence of a transaction has been identified, one or more indicators of the nature of said transaction, and said storing step includes organising transaction data stored in said data repository by said one or more indicators such that it forms a record.

79. The method of claim 78, wherein said rules of said policy data define said one or more indicators of the nature of a transaction, said indicators being one or more of: the address of the location to which said transaction data is transmitted or from which it is received; part of the data path to the location to which said transaction data is transmitted or from which it is received; account codes, reference numbers, credit card numbers, digital certificates and pre-determined keywords.

80. The method of any of claims 51 to 79, wherein said analysing step includes determining the length of time taken to complete a transaction, and causing this to be stored in the data repository with the transaction data.

- 119 -

81. The method of any of claims 51 to 80, wherein said analysing step includes determining the number of keystrokes or mouse-clicks made in order to complete the transaction, and causing this to be stored in the data repository with the transaction data.

82. The method of any of claims 51 to 81, wherein said analysing step includes determining the length of time spent waiting for data to be downloaded to said application during said transaction, and causing this to be stored in the data repository with the transaction data.

83. The method of any of claims 51 to 82, wherein said analysing step includes determining the length of time the web server takes to respond to outbound messages sent from said application, and causing this to be stored in said data repository.

84. The method of any of claims 51 to 83, wherein said analysing step includes determining the length of time a user of said application takes to transmit a response to data downloaded to said application during said transaction, and causing this to be stored in the data repository with the transaction data.

85. The method of claim 84, wherein said analysing step includes determining from the length of time a user of said application takes to transmit a response to data downloaded to said application, the length of time the user spends entering the data that forms the response, and causing this to be stored in the data repository with the transaction data.

86. The method of any of claims 51 to 85, wherein said data repository is accessible by one or more of an

- 120 -

accounts application, an order processing application or other transaction management application.

87. The method of any of claims 51 to 86, further comprising the step of encrypting any relevant data identified in said analysing step before it is stored in said data repository.

88. The method of any of claims 51 to 87, further comprising the step of encrypting the data stored in said data repository.

89. The method of any of claims 51 to 88, wherein said analysing step is performed at said one or more workstations.

90. The method of any of claims 51 to 89, wherein said application is a web browser.

91. The method of claim 90 wherein said analysing step is performed by a plug-in module of said web browser.

92. The method of claim 91, wherein said web browser is Microsoft's Internet Explorer and said plug-in module is a Browser Helper Object.

93. The method of any of claims 51 to 89, wherein said application is an e-mail client.

94. The method of claim 93, wherein said analysing step is performed by a plug-in module of said e-mail client.

95. The method of claim 94, wherein said e-mail client is Microsoft's Outlook e-mail client and said

- 121 -

plug-in module is a Microsoft Exchange client extension.

96. The method of any of claims 51 to 89, wherein said application is an Instant messaging application.

97. The method of claim 96, wherein said analyser is a plug-in module of said Instant messaging application.

98. The method of any of claims 51 to 88, wherein said network comprises a server, and said analysing step is performed at a point on said network intermediate said one or more work stations and said server, or said analysing step is performed at said server.

99. The method of any of claims 51 to 98, wherein said computer network to which said one or more workstations are adapted for connection is a public computer network, and wherein said one or more workstations together form a private computer network.

100. The method of any of claims 51 to 99, further comprising the step of providing a supervisor workstation, said policy data being accessible by said supervisor workstation, such that a user of said supervisor workstation can edit said policy data.

101. A computer program product for controlling a plurality of computers in a private network to manage information, the network having a data repository arranged to receive data from the plurality of computers, and policy data defining rules for the recording of data that may comprise part of a transaction conducted between a computer in the private

- 122 -

network and a third party across a public network, comprising:

a recording medium readable by a computer, having program code recorded thereon which when executed on each of said plurality of computers configures said computers to:

analyse, in conjunction with an application running on the computer that is operable to transmit outbound data to said public network and receive inbound data from said public network, at least one of said outbound data and said inbound data to identify, with reference to said rules of said policy data, the existence of a transaction occurring between the computer and a third party; and

to control said computer to store transaction data that is all or part of said outbound data or said inbound data related to an identified transaction in said data repository.

102. The computer program product of claim 101 wherein said program code when executed on said computer is operable to identify the existence of a transaction by determining whether a secure link has been negotiated between said application and a remote site on said public network, and whether the outbound data or said inbound data is transmitted on that link.

103. The computer program product of claim 101 or 102, wherein said public network is the Internet, and said rules of said policy data define the addresses of non-eCommerce web sites and/or non-eCommerce e-mail accounts, wherein said program code when executed on said computer is operable to disregard any transactions that are identified between the computer and a non-eCommerce web site and/or e-mail account such that no transaction data related to a transaction made to a

- 123 -

non-eCommerce web sites or a non-eCommerce e-mail account is stored in the data repository.

104. The computer program product of any of claims 101 to 103, wherein said program code when executed on said computer is operable to identify the existence of a transaction by reference to said rules of said policy data, said rules of said policy data defining the addresses of known eCommerce locations.

105. The computer program product of any of claims 101 to 104, wherein said program code when executed on said computer is operable to identify credit card numbers, and the existence of a transaction is identified by identifying credit card numbers in said outbound data or inbound data.

106. The computer program product of any of claims 101 to 105, wherein said program code when executed on said computer is operable to identify the existence of a transaction by reference to said rules of said policy data, said rules of said policy data defining one or more of pre-determined digital certificates, account codes, pre-determined keywords, pre-determined names and addresses and embedded codes.

107. The computer program product of any of claims 101 to 106, wherein said program code when executed on said computer is operable to identify in said inbound data an embedded code, said embedded code having been placed in said inbound data to identify it as transaction data.

108. The computer program product of any of claims 101 to 107, wherein said program code when executed on said computer is operable to identify the existence of

- 124 -

a transaction by identifying an electronic receipt in said outbound or inbound data.

109. The computer program product of any of claims 101 to 108, wherein said program code when executed on said computer is further operable to control the computer to record a pre-determined number of subsequent transmissions of said outbound data or said inbound data following an identification of the existence of a transaction, providing that the address or organisation to which the subsequent transmissions are transmitted, or from which they are received, is the same as the address or organisation to which the outbound data was sent or from which the inbound data was received and in which the existence of a transaction was identified.

110. The computer program product of claim 109, wherein said program code when executed on said computer is operable to detect one or more indicators of the nature of the transaction, and said rules of said policy data define the number of subsequent transmissions of said outbound data and said inbound data that are to be recorded in the data repository based on the identified nature of the transaction.

111. The computer program product of claim 109, wherein said rules of said policy data define the number of subsequent transmissions of said outbound and said inbound data that are to be stored in said data repository in dependence on the indicator by which the existence of a transaction was identified.

112. The computer program product of any of claims 101 to 111, wherein said program code when executed on said computer is operable to control the computer to

- 125 -

record all subsequent transmissions of said outbound data or said inbound data that occur within a pre-determined amount of time following an identification of the existence of a transaction, providing that the address or organisation to which the subsequent transmissions are transmitted, or from which they are received, is the same as the address or organisation to which the outbound data was transmitted or from which the inbound data was received and in which the existence of a transaction was identified.

113. The computer program product of claim 112, wherein said program code when executed on said computer is operable to detect one or more indicators of the nature of the transaction, and said rules of said policy data define the amount of time for which all subsequent transmissions of said outbound data and said inbound data are to be recorded in said data repository based on the identified nature of the transaction.

114. The computer program product of claim 112, wherein said rules of said policy data define the amount of time for which subsequent transmissions of said outbound and said inbound data are to be stored in said data repository in dependence on the indicator by which the existence of a transaction was identified.

115. The computer program product of any of claims 101 to 114, wherein said program code when executed on said computer is operable to identify the completion of a transaction, and control the computer to store all or part of said outbound data transmitted by said application and all or part of said inbound data received by said application after the existence of a transaction has been identified and before the

- 126 -

completion of a transaction has been identified in said data repository.

116. The computer program product of claim 115, wherein said program code when executed on said computer is operable to identify subsequent related data contained in said outbound data transmitted by said application and said inbound data received by said application after the completion of a transaction, and control the computer to store said subsequent related data in the data repository with said transaction data already identified.

117. The computer program product of claim 116, wherein said program code when executed on said computer is operable to identify said subsequent related data by identifying common indicators in both said transaction data already identified and said outbound data transmitted by said application and said inbound data received by said application after the completion of a transaction has been identified,

said common indicators being one or more of the address of the location to which said outbound data is transmitted or from which said inbound data is received, part of the data path to the location to which said outbound data is transmitted or from which said inbound data is received, account code or reference numbers.

118. The computer program product of any of claims 101 to 117, wherein said application is operable such that a user of said application can indicate said outbound and said inbound data that is related to a transaction, said program code when executed on said computer being operable to identify said outbound and said inbound data so indicated.

- 127 -

119. The computer program product of any of claims 101 to 118, wherein said program code when executed on said computer is operable to control the computer to store all of said outbound data and said inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction and, to retrieve a pre-determined amount of previous data from said outbound data and said inbound data stored in memory if the existence of a transaction is identified, and to control the computer to store said previous data in the data repository with said transaction data.

120. The computer program product of claim 119, wherein said rules of said policy data specify the amount of previous data that is to be retrieved in dependence on the indicator by which the existence of a transaction is identified.

121. The computer program product of claim 119 or 120, wherein said public network is the Internet and said application is a web browser, said web browser being operable to store each web page that is viewed by said web browser in memory as previous data.

122. The computer program product of claim 121, wherein said rules of said policy data specify the number of web pages that are to be retrieved from those previously stored in memory in dependence on the indicator by which the existence of a transaction is identified.

123. The computer program product of any of claims 101 to 122, wherein said program code when executed on said computer is further operable to control the computer to store all of said outbound data and said

- 128 -

inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction and, to identify, in said previous data, earlier relevant data that is related to said transaction data already identified, and control the computer to store the earlier relevant data in the data repository with said transaction data.

124. The computer program product of claim 123, wherein said program code when executed on said computer is further operable to identify said earlier relevant data in said previous data, by identifying common indicators in both said transaction data and said previous data, said common indicators being one or more of the address of the location to which said outbound data is transmitted or from which said inbound data is received, part of the data path to the location to which said outbound data is transmitted or said inbound data is received, account codes or reference numbers.

125. The computer program product of any of claims 101 to 124, wherein said program code when executed on said computer is further operable to control the computer to store all of said outbound data and said inbound data in memory as previous data, irrespective of whether it may or may not be part of a transaction; and

wherein said computer program product further comprises a selector, recorded on said recording medium, said selector being operable to select earlier relevant data from said previous data in response to input from a user,

and wherein said program code when executed on said computer is further operable to control the computer to store said earlier relevant data selected

- 129 -

by the user in said data repository together with said transaction data.

126. The computer program product of any of claims 101 to 125, wherein said program code when executed on said computer is operable, once the existence of a transaction has been identified, to determine the nature of said transaction by analysing the content of said outbound and inbound data,

said rules of said policy data defining how said transaction data is to be stored in said data repository in dependence on the nature of the transaction that has been determined, said transaction data being stored in said database according to said determination and said rules of said policy data.

127. The computer program product of claim 126, wherein said program code when executed on said computer is further operable to determine the nature of the transaction by identifying in said outbound data and said inbound data one or more indicators, said indicators being defined in said rules of said policy data, and being one or more of: the address of the public network location to which said data that may be part of a transaction is transmitted or from which it is received; part of the data path to the public network location to which said transaction data is transmitted or from which it is received; account codes; reference numbers; credit card numbers; digital certificates and pre-determined keywords.

128. The computer program product of any of claims 101 to 127, wherein said program code when executed on said computer is further operable, once the existence of a transaction has been identified, to identify one or more indicators of the nature of said transaction,

- 130 -

and to control the computer to organise the storage of said transaction data by said one or more indicators such that it forms a record.

129. The computer program product of claim 128, wherein said rules of said policy data define said one or more indicators of the nature of a transaction, said indicators being one or more of: the address of the public location to which said transaction data is transmitted or from which it is received; part of the data path to the public location to which said transaction data is transmitted or from which it is received; account codes, reference numbers, credit card numbers, digital certificates and pre-determined keywords.

130. The computer program product of any of claims 101 to 129, wherein said program code when executed on said computer is further operable to determine the length of time taken to complete a transaction, and cause this to be stored in the data repository with the transaction data.

131. The computer program product of any of claims 101 to 130, wherein said program code when executed on said computer is further operable to determine the number of keystrokes or mouse-clicks made in order to complete the transaction, and cause this to be stored in the data repository with the transaction data.

132. The computer program product of any of claims 101 to 131, wherein said program code when executed on said computer is further operable to determine the length of time spent waiting for data to be downloaded to said application during said transaction, and cause

- 131 -

this to be stored in the data repository with the transaction data.

133. The computer program product of any of claims 101 to 132, wherein said program code when executed on said computer is further operable to determine the length of time the web server takes to respond to outbound messages sent from said application, and cause this to be stored in said data repository.

134. The computer program product of any of claims 101 to 133, wherein said program code when executed on said computer is further operable to determine the length of time a user of said application takes to transmit a response to data downloaded to said application during said transaction, and cause this to be stored in the data repository with the transaction data.

135. The computer program product of claim 134, wherein said program code when executed on said computer is further operable to determine from the length of time a user of said application takes to transmit a response to data downloaded to said application, the length of time the user spends entering the data that forms the response, and cause this to be stored in the data repository with the transaction data.

136. The computer program product of any of claims 101 to 135, wherein the data repository is accessible by one or more of an accounts application, an order processing application or other transaction management application.

- 132 -

137. The computer program product of any of claims 101 to 136, wherein said program code when executed on said computer is further operable to cause any identified relevant data to be encrypted before it is stored in said data repository.

138. The computer program product of any of claims 101 to 137, wherein said program code when executed on said computer is further operable to cause any relevant data stored in the data repository to be encrypted.

139. The computer program product of any of claims 101 to 138, wherein said program code is executable at each of said computers.

140. The computer program product of any of claims 101 to 139, wherein said application is a web browser.

141. The computer program product of claim 140 wherein said program code when executed on said computer is a plug-in module of said web browser.

142. The computer program product of claim 141, wherein said web browser is Microsoft's Internet Explorer and said plug-in module is a Browser Helper Object.

143. The computer program product of any of claims 101 to 139, wherein said application is an e-mail client.

144. The computer program product of claim 143, wherein said program code when executed on said computer is a plug-in module of said e-mail client.

- 133 -

145. The computer program product of claim 144, wherein said e-mail client is Microsoft's Outlook e-mail client and said plug-in module is a Microsoft Exchange client extension.

146. The computer software product of any of claims 101 to 139, wherein said application is an Instant messaging application.

147. The computer software product of claim 146, wherein said program code when executed on said computer is a plug-in module of said Instant messaging application.

148. Computer program product of any of claims 101 to 138, wherein said network includes a server and said program code is executable at a point on said network intermediate said one or more workstations and said server, or said program code is executable at said server.

149. The computer program product of any of claims 101 to 148, further comprising program code recorded on the recording medium which when executed on a computer in the plurality of computers enable that computer as a supervisor workstation, said supervisor workstation having access to said data repository and being operable to view said relevant data stored in said data repository.

150. The computer program product of any of claims 101 to 149, wherein said policy data is accessible by said supervisor workstation, such that a user of said supervisor workstation can edit said policy data.

1/25





Fig. 1

2/25

Fig. 2



3/25




Fig. 3

4/25




Fig. 4

5/25






Fig. 5

6/25


Fig. 6



7/25




8/25



9/25

Fig. 9



10/25





Fig. 10


11/25

Fig. 11



12/25

Fig. 12



13/25




Table q - Excluded Sites
www.hotmail.com
www.passport.com
ibankon.barclays.co.uk
www.nwolb.co.uk

Table s - Known eCommerce Sites
ecomm.us.dell.com/dellstore
buy.supersaver.co.uk
www.booksforall.com/basket


Table r - Account Codes	
Account Code	Previous Pages To Record
21321234	2
ORCH01	6
58734	1
PETER304	0


Table t - Keywords
"receipt"
"thank you for your order"
"order confirmation"

Fig. 13

14/25

Fig. 14








Table f - Approvers		
Username	Limit	Excluded Sites
F Smith	\$500	www.dell.com
R Jones	\$1000	www.dell.com; www.officemax.com
F Healy	Unlimited	none

Fig. 16

17/25






Fig. 17

Fig. 18





Example of Table s
*@microsoft.com
fred.smith@xyz.com
jjones@hotmail.com

Example of Tables x
* @ hotmail.com
*@ aol.com
username

Example of Tables t, y
confidential
secret
contract
quotation
order
project x

Example of Table u
A.N. Authoriseduser
T.H.E. Boss

Fig. 19

20/25




Fig. 20

21/25





Fig. 21



Example of Tables j
confidential
secret
contract
quotation
order
project x

Fig. 22

23/25




Fig. 23




Fig. 24

25/25




Fig. 25

Back-Office Application Plug-Ins

