



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

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

(45) Date de délivrance/Issue Date: 2009/09/15

(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)


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

(73) Propriétaire/Owner:
ORCHESTRIA LIMITED, 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 transmit to the network or has just received from the
10 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
15 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
20 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.

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

10 Electronic commerce ("eCommerce"), particularly 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 personal meetings. Sellers advertise products and
15 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 services of a particular market sector. Buyers can find suppliers, select goods, obtain
20 quotations, place and track orders, and even make payments entirely electronically and at any time. eCommerce brings the promise of increased flexibility, choice and efficiency, with dramatically reduced procurement costs.

25 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 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
30 web browser is Microsoft Corporation's "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
5 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
10 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
15 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
20 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
25 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
30 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
35 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 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 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 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 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 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 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 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 one or more workstations adapted for connection to a computer network, each workstation having a memory; an
5 application stored in the memory of each workstation for transmitting outbound data to the network and receiving inbound data from the network; policy data containing rules specifying an appropriate encryption strength for outbound data, the encryption strength
10 depending on the content of the data; and an analyser, the analyser being operable in conjunction with the policy data to monitor the outbound data and to determine, in accordance with the rules in the policy data, an appropriate encryption strength for the
15 outbound data; wherein the analyser controls transmission of the outbound data from the application in dependence upon the determination of an appropriate encryption strength.

According to another aspect of the invention, a
20 method of managing information comprises the steps of providing one or more workstations adapted for connection to a computer network, each workstation having a memory; providing an application stored in the memory of each workstation for transmitting outbound
25 data to the network and receiving inbound data from the network; providing policy data containing rules specifying an appropriate encryption strength for outbound data, the encryption strength depending on the content of the data; and analysing the outbound data to
30 determine, in accordance with the rules in the policy data, an appropriate encryption strength for the outbound data; and controlling transmission of the outbound data from the application in dependence upon the determination of an appropriate encryption strength
35 in the analysing step.

- 11 -

According to a further aspect of the invention, there is provided a computer program product for controlling a computer connected to a public network to manage information, the computer having access to

5 policy data containing rules specifying an appropriate encryption strength for outbound data transmitted to the public network, the encryption strength depending on the content of the data, comprising a recording medium readable by the computer, having program code

10 recorded thereon which when executed on the computer, configures the computer to determine, in conjunction with an application running on the computer that is operable at least to transmit outbound data to the public network, with reference to the rules in the

15 policy data, an appropriate encryption strength for the outbound data; and control the transmission of the outbound data by the application in dependence upon the determination of an appropriate encryption strength.

The information management system provides many

20 advantages in the eCommerce environment to on-line trading companies, who may benefit by being able to regulate the transactions made by their staff according to their instructions encoded in the policy data, automatically maintain records of passwords and

25 business conducted on-line, avoid paying for 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

30 management system will now be described in more detail, 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

35 the Internet according to the prior art;

- 12 -

Figure 2 is schematic illustration of the preferred embodiment of the information management system implemented in a corporate environment;

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

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 preferred embodiment;

Figure 6 is a flowchart illustrating the operation of a plug-in module, according to a preferred embodiment, for capturing username and password values transmitted by a user to a remote web 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, according to a preferred embodiment, for recognizing credit card numbers contained in data transmitted to or from a web server or e-mail client;

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, for identifying transmissions from a user or to a user that comprise part of an eCommerce transaction;

- 13 -

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
5 operation of a plug-in module, 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
10 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
15 requires approval, and for identifying an appropriate approver;

Figure 17 is a flowchart illustrating the operation of a plug-in module, according to a preferred embodiment of the invention, for determining an
20 appropriate level of encryption for a transmission and 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
25 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
30 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 up-
35 load of information to a site external to the company, utilizing the policy data shown in Figure 19;

- 14 -

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
5 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

10 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
15 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
20 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 intelligently manage the transmission of e-mails.

25 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 implementation and use of the system by a company of
30 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 Internet may benefit from the functionality provided by
35 the preferred system.

- 15 -

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 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 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 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 1998. The use of the 'Browser Helper Object' and '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 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 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 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

- 16 -

Internet and retrieve web page information from web servers 14 and 18. Alternately users with e-mail 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 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 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 can be both a buyer and a seller, and as buyers abc.com 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 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 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 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

- 17 -

is also connected to Back Office Application Plug-ins 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 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 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 preferred system.

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

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.

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.

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.

- 18 -

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 determines that a particular transaction requires approval, a request is sent to the Central Management Server and queued pending approval by an authorized operator.

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 a policy file on the central Management Server for access by any of the Operator Consoles 44, Back Office 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 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 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 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.

- 19 -

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
5 all purchases to require his approval if company cash 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.

10 Each user of the network will have his or her own 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
15 is preferably stored in file form on the Central Management Server, it is not intended that storage of the policy data be restricted to file form only. Any other representation or encoding of policy settings may be employed within the preferred system.

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

Use of the preferred system in a web browser

Figure 3 shows the simplified operation of a web
25 browser. The web browser is launched at step S100 in response to a start request from either the user or 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
30 is booted-up. After the web browser has been started it typically requests a 'home page', the default web 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
35 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.

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

10 The browser then typically waits for user input at step S112. Such input may include filling displayed 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
15 be another web page address, or it may contain additional data such as that typed in to displayed fields by the user.

Figure 4 shows a sample web page display, in which a GUI is presented to the user in order to receive the
20 users name and e-mail address. It will be seen from reference to Figure 4 that the user has entered his 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.

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

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

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

- 21 -

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
5 via e-mail, or to gain access to other web pages.

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
10 is implemented to respond to a number of significant 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
15 requests and data submitted to the web server from the browser and identify data that is unique to the user. It may do this simply by searching the outbound data stream for the presence of pre-determined words or phrases. In the above case shown in Figure 4, the two
20 variable definitions 'UserID' and 'email' may be searched for, and the data that follows them extracted 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
25 data. The BHO may also monitor the inbound data stream received from the web site being connected to.

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-
30 dependent software modules or objects that something significant has just occurred or that an action has 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
35 is normally available. The BHO is implemented to trap these events and to take action in dependence on them.

- 22 -

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 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 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 the web site into memory. The page is encoded in the form of a Document Object, conforming to Microsoft's™ 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 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 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 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

- 23 -

asterisks. This too may be determined from analysis of 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
5 by the user into an HTML form. Usually, the potentially sensitive user data is transmitted to the 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
10 the DOM to extract the user data, and, if necessary, prevent the data from being transmitted.

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
15 before it is encrypted or after it is decrypted. This is advantageous since there is no need for the BHO to perform any encoding or decoding of data itself. This does not affect the ability to determine if the link is secure or not, since a secure link can be identified by
20 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 after decryption occurs.

25 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 implementation of the preferred embodiment in an e-mail client will now be described.

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

35 An e-mail client's 'receive message' operation is initiated at step S120. This may be done automatically

- 24 -

at predetermined intervals in order to keep the user 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
5 the e-mail server 95 and download any new messages to the user's 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.
10 The user then clicks on 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 response, the e-mail
15 client provides an interface comprising a text editor in which the user can enter 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
20 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.

In the preferred embodiment, additional functionality is
25 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 the plug-in module is encoded as an Exchange client extension. These are described in the
30 document "Microsoft Outlook and Exchange Client Extensions" by Sharon Lloyd mentioned earlier.

- 25 -

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

In the preferred embodiment, the Exchange Client
20 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 actions. The 'events' in question are provided by the COM interfaces mentioned above. The monitoring of the
25 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.

The e-mail client plug-in module is implemented to respond to the 'OnDelivery' event for example which is
30 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 information to access the different parts of the e-mail message that have been downloaded and which are held in
35 memory. The message header, the message body and any message attachments are encoded in memory as properties

- 26 -

of the message object which may be separately accessed through Mail Application Program Interface (MAPI) calls.

Via the information supplied as part of the
5 'OnDelivery' event, the plug-in module may access the message header and extract the identity of the sender 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
10 pertinent data. It may search for evidence of an eCommerce transaction, by identifying significant words 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
15 message may be deleted unseen.

The analysis of a received e-mail occurs therefore at point A in Figure 5 prior to being viewed by the user. Preferably the e-mail is examined before the e-mail is even placed in the Inbox. Where a message is
20 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-
25 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
30 '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
35 encryption takes place. The new message which is to be transmitted is similarly stored in memory as an object

- 27 -

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
5 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
10 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.

15 Identifying and Recording Usernames, Passwords and other Information

The preferred system provides means to automatically identify, collect and store data contained in transmissions to and from a user's
20 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
25 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
30 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
35 drawback of the password file being stored locally is that if the user moves to another machine he does not

- 28 -

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 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 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 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 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,

- 29 -

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
5 preferred system, the password information may be 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
10 implemented to extract user name and password 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
15 the browser. This occurs at point 'C' in the process illustrated in Figure 3. Control then passes to step S152 where the plug-in module determines whether or not the data which is to be transmitted contains username or password information.

20 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
25 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.

30 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
35 themselves meaningful, the nature of data being transmitted, may be inferred from the field type of the

- 30 -

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 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 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, 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 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

- 31 -

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
5 may then transmit the data to the web server.

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
10 data, in particular those relating to eCommerce 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
15 applied to e-mail systems.

The information may be extracted in the manner described above, via the DOM or via MAPI calls to the COM representation of e-mail content, or may be extracted from the language in which a web page is
20 encoded. Web pages are typically encoded in Hyper-Text Mark-Up Language (HTML), a human readable text based 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
25 involve recording just password and username information, recording the URL of a web page being 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
30 retrieved later and viewed.

The plug-in modules provided by the preferred 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
35 preferred system to instruct the operation of each of

- 32 -

the plug-in modules thereby controlling their functionality.

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

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

15 The browser branch contains three sub-branches called 'DataToRecord', 'WhenToStartRecording' and 'WhenToStopRecording'. The DataToRecord branch specifies the type of data that is to be extracted from transmissions to or from the user's workstation and a
20 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 that is entered by a user into fields provided on the web page and submitted to a web site, and any passwords
25 and usernames that are entered by the user. These are referred to on four distinct sub-branches of the DataToRecord branch, entitled 'URL', 'HTML', 'SubmittedFields' and 'Passwords'. A Yes/No option on each of those sub-branches specifies whether or not the
30 indicated data is to be recorded.

The WhenToStartRecording branch sets out a number of conditions which specify the point at which the data specified in the DataToRecord branch is to be recorded. Five conditions are illustrated in this example each
35 one set out on a separate branch and respectively entitled 'WhenBrowserIsOpened', 'IfCreditCard

- 33 -

NumberSubmitted', 'IfPasswordSubmitted',
'IfKeywordsReceived' and 'IfKeywordsSent'. Whether or
not these conditions are to be used in order to
determine when to start recording is indicated by
5 Yes/No markers on each branch.

Similarly, the WhenToStopRecording branch lists
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
10 conditions are 'WhenUserClosesBrowser',
'WhenUserChangesSite' and 'WhenUserChangesPage'. The
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
15 operation of the plug-in module.

The e-mail branch of the policy data is divided
into a DataToRecord branch and a WhenToRecord branch.
Each of these branches are sub-divided into branches
that deal with Sent Mail and Received Mail. The type
20 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
the case of messages signed with a digital signature, a
copy of the digital certificate accompanying the
25 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
number, a keyword or a digital certificate in the mail
that is sent or received.

30 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
different users to be assigned to different branches of
the tree in order to receive different policies.
35 Although, the tree like structure is preferred, other
arrangements may also be possible. The branches shown

- 34 -

in this diagram are only intended to be illustrative.

Identification of Credit Card Numbers

The preferred system also searches for Credit Card
5 numbers or other account information in the data to be
transmitted to the web server or e-mail client by
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
10 used by credit card companies to validate card numbers.

If a credit card number is found in the transmission,
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,
15 renegotiate a higher level of encryption to safe-guard
the transmission if the credit card number belongs to
the company, or prevent the transmission from taking
place altogether.

The most common test to identify a credit card
20 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
25 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
30 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',
35 '18' counts as '1 + 8 = 9'. If the total of the sum is

- 35 -

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 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

- 36 -

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
5 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
10 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
15 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.

However, if the checksum indicates that the
20 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 action to take. The action may be determined from such factors such as the number itself, the identity of the
25 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 that a higher encryption strength is required for the transmission before it may proceed.

30 This policy checking step allows credit card transactions to be controlled at a higher level than the user making the transaction. Thus financial decisions may be swiftly and easily implemented, and may be enforced automatically without the need for
35 monitoring. An organisation may, for example, wish to restrict the capability to make credit card

- 37 -

transactions on the organisation's account to particular authorised people or it may wish to restrict transactions on a particular account altogether.

In step S170, the credit card number, and other
5 details of the transaction are compared to the settings in the policy file and it is determined whether or not 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,
10 control passes to step S172 where transmission of the data is halted, and then to step S174 where the module 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
15 3, in the case of a web browser, or to step S132, 'compose mail' in Figure 5, in the case of an e-mail client.

If in step S172, it is determined that the credit card number may be transmitted, control passes to step
20 S176 where transmission of the data occurs, and then to step S180 where the module exits. In this case control 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.

25 Credit card numbers need not be identified in step S162 solely by scanning the content of the 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
30 to be transmitted or also from the representation of the web page in memory. The discussion above about the identification of passwords explains this in more detail.

The preferred system may also be configured to
35 search outgoing transmissions for other pertinent financial details, such as account numbers. Company

- 38 -

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 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 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 extract credit card numbers from transmissions that are being received.

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.

- 39 -

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
5 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,
10 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
15 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 issuing authority. A private key is also issued to the
20 certificate holder who should not disclose it to anyone else.

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
25 his private key has been compromised.

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
30 after decrypting the message with his private key.

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
35 encrypted hash may then be used as a digital signature for the message that is being transmitted. The original

- 40 -

message, the digital signature and the transmitter's 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
5 produce a hash for the received message. If, the received hash, having been decrypted with the holder's 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,
10 and that the message has not been altered en route from its original form. Digital certificates are therefore of considerable and increasing importance to companies conducting business on the Internet.

In cases where an on-line trader makes use of
15 Digital Certificates to ensure the identity of its customers, it is necessary to check with the issuer that the certificate is still valid before any transaction is authorised. Such checks can be carried out on-line using an independent verification service
20 such as that provided by Valicert, Inc. A fee is usually charged for such services.

It may be that individual employees of an organisation each receive e-mails from a single client, each signed with his digital certificate, on separate
25 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 themselves manually, and as a result, individual employees might request that the same certificate be
30 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 validation fees spent on an already revoked certificate are unnecessary. Additionally, the recipient may wish
35 to make a business judgement as to whether a previously validated certificate should be re-checked or not.

- 41 -

For example, if a digitally signed order for \$1M worth 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
5 certificate, the organisation may consider a second validation check to be unnecessary, thereby saving the validation charge.

The preferred system provides means to record information about the Digital Certificates that have
10 been received, the status of the certificate at the last check as well as, where applicable, transaction 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
15 preferred system also provides means to use the stored information to decide whether or not a validation check is desirable, and to accept or refuse transmissions depending on the status of the digital certificate. Thus, users of the system may receive and review
20 transmissions without the need to establish their authenticity themselves.

Figure 9 illustrates the operation of a plug-in module of the preferred system implemented to extract digital certificates from transmissions received by
25 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 so on. The module first checks to determine if the certificate is obviously invalid, and that the message
30 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 a thumbprint may be a checksum for the certificate itself for example. The message has not been correctly
35 signed if the signature cannot be verified from the information contained within the certificate. Details

- 42 -

of certificate validation and message signing are more 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
5 in the database and only records those that have not. Where a copy of the certificate is already stored, the module checks the database record to determine whether it has been previously identified as revoked in which case the transmission is immediately rejected.

10 Otherwise, the module then determines, in accordance with policy defining business rules, whether or not to validate the certificate. Taking in to account the results of any such validation, it then determines whether the certificate should be relied upon, and
15 therefore whether or not the transmission signed by the digital certificate should be rejected or accepted. The module is initiated at step S190, following receipt of data containing a digital certificate. Digital certificates are typically transmitted as attachments
20 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 indicate whether an attachment is a digital certificate or not.

25 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 implemented in an e-mail client. Following initialisation, the module proceeds to step S191 in
30 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 message is incorrectly signed, the module proceeds to step S198 and rejects the transmission. Otherwise the
35 module proceeds to step S192 in which the database is searched for a previously received copy of the digital

- 43 -

certificate. Control then passes to decision step S194. 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
5 previously been marked as revoked. This will have occurred if a previous validity check revealed that a 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
10 certificate and the date on which it was received are stored in the database, together with additional 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
15 etc. If the certificate has already been marked as revoked in step S196 then control passes directly to step S198 in which the transmission comprising the digital certificate is automatically rejected. This may involve for example, transmitting an automatically
20 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 certificate from taking any further steps in connection with the refused transmission. The module then exits at
25 step S200.

If, however, the certificate has not been previously marked as revoked in step S196, then control passes to step S204 in which the history of transmissions signed by the certificate, by other
30 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 of the certificate is required. Control also passes to step S204 after a new digital certificate has been
35 added to the database in step S202.

- 44 -

The policy data contains instructions which when considered in conjunction with the history of signed transmissions previously received and revocation checks previously made indicate whether or not the certificate
5 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.

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

35 Issuers of digital certificates may provide many different classes of certificate suited for different

- 45 -

purposes. These may be separately dealt with by the policy data by corresponding sub-branches of the policy datatree.

In the example policy the first branch entitled
5 'IdentityCertificates' deals with transmissions that do not involve a monetary transaction. The branch comprises four separate sub-branches. The first of these, entitled 'AlwaysAcceptFrom' contains a reference to a table, 'table a', which lists the names of
10 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, it is not considered necessary to determine whether or not a digital certificate has been revoked by its
15 issuer.

The second branch entitled 'AlwaysCheckFrom' contains a reference to a separate table, table b, in which the names of organisations and individuals for whom digital certificates should always be checked are
20 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.

The third branch entitled 'CheckIfDaysSince
25 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 are not considered necessary. In this case, the period
30 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
certificate. In the example shown the policy data
35 specifies that validation checks on any given digital certificate need only be made every 30 days. Again, the

- 46 -

number of days specified on both of these branches is left to the user of the preferred system to decide. The amount of time that has passed since a valid digital certificate has been received may be determined
5 by referring to digital certificates and associated data stored in the database. By checking digital certificates periodically, rather than each time they are received allows money spent on making checks to be saved.

10 The MonetaryCertificates branch also contains an AlwaysAcceptFrom and an AlwaysCheckFrom branch which refer to tables x and y respectively. Table x lists all those organisations and individuals for whom no status check of a digital certificate is required;
15 table y lists all those for whom a check is always required. The MonetaryCertificate branch also contains a CheckIfAmountExceeds branch which specifies a threshold transaction amount over which all digital certificates must be checked and lastly an
20 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 that digital certificates received for transactions for
25 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 which interacts with the policy data shown in figure 9.
30 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 Digital Certificate it has received. The sub-process
35 begins in step S220 from which control passes to decision step S222 in which it is determined if the

- 47 -

transmission is monetary from the class of the Digital Certificate used to sign the message. If the 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.

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 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 results of that check control flows either to step S242 in which confidence in the Digital Certificate is 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 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 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 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

- 48 -

'AlwaysCheckFrom' table. Similarly, if the sender is listed in this table control flows to step S244 at which an on-line check of the Digital Certificate's status is deemed necessary. Control returns from step
5 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 S226 to the next decision step in the chain which represents the next condition listed as a sub-branch
10 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 involve looking up the Digital Certificate in the database of stored Digital Certificates and extracting
15 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 days, control flows to step S242 where confidence is established. If the information in the database of
20 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
25 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.

If the other Digital Certificate has been checked in
30 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
35 through in decision steps S232 to decision step S240. If in decision step S232 the sender of the Digital

- 49 -

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. 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 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

- 50 -

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
5 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
10 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
15 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
20 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 more general policies which are to be used in
25 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 before to see if the amount and the goods and services
30 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 has been revoked, then a third party may have acquired
35 the original holder's private key and be attempting to make fraudulent transactions.

- 51 -

Having checked the policy data in step S204, confidence in the digital certificate will either have been established or will not have been established. In decision step S206, if confidence has been established
5 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 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
10 case of an e-mail client.

If in step S206 confidence in the digital certificate is not established, then control passes to step S210, where an on-line validation check is performed on the digital certificate. This may involve
15 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 confirm a warranty for the amount promised in the transaction. Control next passes to step S212, in which
20 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 transmission is rejected, or to step S208 where the
25 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 some other indicator. Following either of steps S198
30 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 amount, in order that the information may be used in
35 determining the need for future validation checks.

- 52 -

Recording of information.

The preferred system also provides an automatic 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 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 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 the system for indications that a transaction has begun or is taking place. Such indications are numerous. The 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 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 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 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

- 53 -

configured to record all data transmitted to a web page which has been identified as that of an on-line trading company. Other indications might be an identified credit card number, an electronic receipt, an e-mail
5 acknowledging the sale, use of a digital certificate, in particular a digital warranty certificate, or a purchase code.

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

15 Recording of transaction data may be stopped when the end of the transaction is identified or after a predefined number of transmissions between the purchaser and the trader have taken place. Similarly, once a transaction has been identified, the preferred
20 system may record in the database a predefined number of cached transmissions that took place immediately before the first recognised transmission of the transaction.

This will be useful when, for example, the first
25 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
30 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 that in which the transaction was detected, of a
35 different type, or be a mixture of types. For example a user might visit a web site www.abc.com, obtain

- 54 -

details of goods, and then order them in an email sent to orders@abc.com.

The preferred system records the transaction details in a centralised common database 42.

5 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 techniques so that only a person with the necessary authorisation may access it.

10 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-line. Figure 14 illustrates the process by which the preferred system records an identified transaction in
15 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 approvals policy.

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

The module begins operation at step S250 in response to receiving data or in response to a user initiating transmission of data to a remote site. In
25 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 client it will be after point A or B respectively as shown in Figure 5.

30 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 transmitting data and the site receiving data. This may be achieved by querying the URL address that has
35 been connected to, as mentioned above, or interrogating the web browser to see if encryption is being employed.

- 55 -

In the case of an e-mail message this step is omitted 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 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 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 transactions but which do establish secure connections. Browser based e-mail sites such as Microsoft's™ Hotmail 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 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 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 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

- 56 -

the content of the transmission are made in order to 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
5 this has been described with reference to Figure 8. If a credit card number is found in the transmission then 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
10 instead to decision step S262 where the transmission is scanned to see if it contains an account code. Account 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
15 from descriptive data in the transmission such as a field name like "Account Number" or similar characters appearing in the text of a message.

If in decision step S262 an account code is found, then it is assumed that the transmission constitutes
20 part of a transaction and control passes to step S258 where the module exits. If no account code is found 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
25 decision step S266, a determination on that comparison is made. If the URL is found to be at a known eCommerce 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
30 exits. Similarly, in the case of an e-mail, the destination address may be compared against a list of 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
35 place and control passes to step S258 where the module exits.

- 57 -

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 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.

10 In step S268, a general check is illustrated which 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
15 placed in the data for example. It is preferred that the web browser or e-mail client being used in the preferred system allows the user to mark transmissions with an embedded code to indicate that the transmission is part of a transaction and should be recorded. Also,
20 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
25 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
30 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
35 recording transmissions and possible transaction details if there is just a suspicion that a transaction

- 58 -

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
5 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
10 '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
15 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-branch corresponds to decision step S252 shown in Figure 12. It will be appreciated with reference to
20 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
25 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
30 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
35 PreviousPages sub-branch of the IfCreditCardNumberPresent branch lists the number of

- 59 -

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 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

- 60 -

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 '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 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.

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 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

- 61 -

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
5 card numbers, account codes or other identifiable 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
10 along with means for recording transmissions between a 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
15 making paper copies of each transmission received or sent. Thus, a company's record keeping is made considerably easier and more accurate.

Figure 14 illustrates the operation of a module for recording transmissions which comprise a
20 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.
25 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.

30 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
35 then passes to decision step S274 where, if it is

- 62 -

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
5 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
10 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
15 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 containing an electronic receipt for transactions for smaller amounts.

20 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
25 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
30 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
35 depend on the trader's identity, geographical location,

- 63 -

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
5 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
10 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
15 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 services; the amount of data to record depending on these keywords and names as well.

20 Also, the tables might be expanded to refer to the amount of different types of data that should be 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
25 from the DOM representation of the web page and stored in the database.

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
30 the page initially identified as being part of a transaction.

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
35 identified transaction.

- 64 -

Following determination of how much data to record 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 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 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 the identified transaction and any retrieved prior transmissions are stored in the system database 42.

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 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 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 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

- 65 -

party may be recorded until the system detects that the 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
5 identifying whether a transaction is taking place. The most simple implementation is to record transmission information until an electronic receipt or shipping order is received. Alternatively, recording of
10 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 of time has passed since the transaction was identified.

Transmissions may be made simpler if each time the
15 user changes web site the cache is emptied. This keeps the memory required for the cache memory low, as well as reducing the number of previous transmissions that need to be searched if searching techniques are to be employed.

20 It will be appreciated that the methods described above can also be used to record associated 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
25 confirmatory e-mail sent from the seller to the buyer.

This e-mail can be detected as forming part of the transaction, since it will contain common characteristics, such as order number, account number, goods description, price etc. It may also be sent from
30 an address similar to the web site address, for example 'customerservices@abc.com' when the original web site 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
35 are considered as being associated with the original transaction.

- 66 -

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 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 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 keystrokes and mouse clicks required to complete a purchase, the amount of 'idle' time while the user 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 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 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.

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

- 67 -

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.

5 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 between each received transmission and associates these totals with the received transmission. Keystrokes and
10 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 determined that the transaction is complete, the totals of keystrokes and mouse clicks which occurred between
15 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
20 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
25 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.
30 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

- 68 -

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.

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 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 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 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 typically contained within a single e-mail and can be

- 69 -

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
5 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
10 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
15 once the required details are known.

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
20 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
25 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
30 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
35 Transactions Approval branch of the policy data tree.

- 70 -

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
5 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
10 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.

15 The third branch entitled 'ExcludedSites' refers to a table containing web site and e-mail addresses of 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
20 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
25 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
30 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.

35 If the conditions on the first three sub-branches of the transaction branch indicate that approval is

- 71 -

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
5 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
10 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.

15 The operation of the approvals process could for example be extended to any kind of transmission, not 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
20 usernames, addresses or keywords for example that are 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
25 via keyword identification.

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
30 transaction may be proceed. Control returns from step 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
35 which the particulars of the transaction are used to determine an appropriate approver for the transaction.

- 72 -

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.

5 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

10 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

15 designated approver via systems approval queue 100. Following step S298 control passes to decision step 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

20 response has not been received in step S300 control 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

25 to wait for a response from the approver. Thus decision 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

30 is taken depending upon whether the transaction was 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

35 the transmission is not approved then control passes from step S300 to step S306 at which the module exits.

- 73 -

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 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.

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. Control then passes to decision step S314. In step S314 if no request is pending control passes back to 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.

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 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 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 be exceeded, or if they are simply over a maximum

- 74 -

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
5 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
10 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
15 S284. The approver receives the location details in Step S316 and the system retrieves the transmissions which constitute the transaction from the database, suitably displaying them such that the approver can consider them in making his/her determination of
20 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
25 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
30 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.

35 Security

- 75 -

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
5 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,
10 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
15 security levels in dependence on the determined nature of the data may also be employed within the scope of the invention.

Figure 17 shows an example illustration of policy data defining appropriate encryption strengths for
20 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
25 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
30 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
35 both have a desired encryption strength of 128 bits, submitted key words have a desired encryption strength

- 76 -

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
5 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
10 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
15 corresponding encryption strengths specific to each password or number.

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
20 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
25 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
30 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
35 bit). The first part of the notification describes the type of the encryption used while the second part

- 77 -

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
5 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
10 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
15 module allows the information to be transmitted;

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
20 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.

c) If the security of the link is not sufficient
25 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
30 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
35 still wish for transmission to take place.

- 78 -

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
5 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
10 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
15 begins operation at step S320 at point C in Figure 3, just prior to the transmission of the data to a web 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
20 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
25 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
30 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
35 for indications of other sensitive data in the data to be transmitted. Such indications will need to be

- 79 -

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.

10 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.

15 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 comparing the detected data with pre-determined entries in the policy data.

25 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. 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, 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.

35 Once the appropriate security level has been determined in step S332, control passes to step S334 in

- 80 -

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 e-mail application prior to transmitting the message.

5 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 which will occur prior to transmission.

10 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 lower than or equal to that determined in step S334,
15 then there is deemed to be enough protection for the 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
20 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
25 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
30 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
35 server and the web browser, or in the case of an e-mail by the sending and receiving e-mail applications.

- 81 -

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. Furthermore, certain types of data may be given a
5 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 changed the e-mail encryption settings, to a higher
10 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 encryption strength at step S340 has not been
15 successful, then it is deemed unsafe to transmit the 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
20 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
25 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
30 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
35 requirement, to allow or prevent transmission based on such security requirements, and on the determined

- 82 -

security level of the link (encryption strength). It will be appreciated that the preferred system provides a significantly improved system for secure transmission which reduces the possibility of human error.

5

Monitoring Out-bound E-mails for sensitive information

In addition to the problem of sensitive data being intercepted by a third party between sender and
10 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 documents, such as customer lists, before leaving the
15 employ of an organisation is easily achieved, virtually 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
20 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
25 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
30 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
35 attachments to the message, the means by which the message is intended to be transmitted, and whether or

- 83 -

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 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 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 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.

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-directed to him/her, and to whom it was originally sent.

- 84 -

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 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.

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 example by using special icons, or by putting up a message box to notify him/her. Means may also be 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 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 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 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 is then sent, either automatically, or when the user presses the ``Send`` button. In this way, with the

- 85 -

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 redirection had never taken place.

5 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 20. Figure 19 is a policy data tree having a number of
10 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
15 plug-in module steps through six steps determining 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
20 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
25 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
30 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
35 e-mail message, and any attachments attached to the e-mail message, is compared with entries in a look-up

- 86 -

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 company. In the next step S356, the e-mail message is
5 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 next decision step S358, it is determined whether or
10 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.

In any of the decision steps S352, S354 or S362
15 where a look-up table is consulted, a match between an 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
20 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
25 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.
30 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.

35 In decision step S356 if encryption of the message is detected then that is deemed enough to warrant re-

- 87 -

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 modified for re-direction. In the case of a message
5 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 the message, preferably also includes the encryption
10 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 digital certificate could be included with the re-
15 directed message. The public key, or certificate, as 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
20 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
25 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.

30 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
35 step S370, and a check of sensitive keywords in the information being uploaded in step S374. Providing the

- 88 -

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
5 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.

10 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.

15

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
20 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.
25 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
30 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
35 only the subject line, particularly when they have large amounts of e-mails to process. As a consequence,

- 89 -

sensitive information is often unintentionally disclosed, both within the organisation, and more dangerously, outside it. There have been several high profile cases where considerable sums have been lost as
5 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 forwarded e-mail is transmitted and preventing the e-
10 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.

Preferably, forwarding occurs in accordance with
15 the content of the e-mail, and the addresses of the 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
20 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>''
25 (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
30 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.
35 Similarly, if the original e-mail was only sent to a single recipient, then it may be determined that the

- 90 -

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 described can be determined in accordance with policy.

5 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 added at the beginning of the original mail by the E-
10 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 immediately after these keystings. Internal and
15 external recipients can be readily distinguished by 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.

20 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-
25 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
30 give a warning to the e-mail client user whenever an e-mail is about to 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
35 affecting forwarding of e-mails within the company and forwarding of e-mails to people outside the company.

- 91 -

The ``PreventKeywords`` branch specifies a look-up table in which keywords indicative of sensitive information are stored. Such keywords may be predefined key strings such as <NOFORWARD> or <NOFORWARDEXTERNAL>, or
5 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 allowed to be forwarded. The last two sub-branches
10 PreventIfNotSentExternally when sent 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 forwarded e-mail may be transmitted to all of the
15 recipients internal to the company, with the external 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.

20 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.

The operation of the plug-in module is illustrated
25 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 keystore ``----Original message----``, since this key string is automatically
30 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 keystore, then it may be deduced that the e-mail message is an original message and is not being forwarded and the
35 message may be transmitted in step S384. If however in step S382 the message is found to contain the key

- 92 -

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 allow the forwarded message to be transmitted. Control
5 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 module scans the e-mail message to see whether the e-
10 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 e-mail message has been sent out of the company before,
15 then a warning is issued to the user in step S392, 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
20 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
25 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
30 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
35 table. If any matches between entries in the e-mail and those in the table are found then the message is

- 93 -

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.

5

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
10 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
15 care is used when digitally signing a document, just as 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
20 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.

25 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;
- 30 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
35 preferred system takes into account a number of factors, including the nature of the message content

- 94 -

(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
5 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
10 query' techniques. Similarly, the nature of the intended or available digital certificates can be determined by reference to the issuer and type of certificate.

Figure 25 illustrates the operation of a plug-in
15 module for ensuring that an e-mail is appropriately 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
20 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
25 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
30 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
35 passes to decision step S408 where the KeywordsTable on

- 95 -

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.

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

15 If the recipients of the e-mail are not to be found in table f, and the e-mail does not contain 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
20 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
25 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
30 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
35 such sensitive keywords, the user of the e-mail client is either notified in step S412 of the desirability to

- 96 -

digitally sign the message, or alternatively, the e-mail message is transmitted without signature in step S416.

5 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 DenySigning branch of the policy data. Table m is
10 specified on the CertificatesUsed branch under the DenySigning branch, and specifies the issuer, type, certificate number or signing key of digital certificates that are deemed of interest. Should the digital certificate or signing key that is to be used
15 to sign the outbound e-mail be found in table m, then 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-
20 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-
25 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.

30 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
35 SuggestNotSigning branch of the policy data tree. As for the DenySigning branch of the policy data tree, the

- 97 -

three decision steps S426, S428 and S430, correspond to the look-up commands for checking the digital certificate or signing key used with the e-mail, the recipient of the outbound e-mail and the text of the
5 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 S428 or step S430, the recipient or the text of the
10 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 the outbound e-mail message. The user however may still
15 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 step S416.

20

Instant Messaging and Telephony Applications

In addition to browser and email activity, additional applications such as Instant Messaging (also known as 'chat') and digital telephony applications
25 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 ITU-T Recommendation H.323 (1998). Many aspects of
30 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 'conversation' is carried out 'live' with both parties
35 present throughout. For the purposes of the present invention however the procedures are identical.

- 98 -

Figure 5 in the drawings can represent Instant Messaging by replacing the word ``E-mail'' in steps S122, S124, S132 and S134 with the word ``Message''. The ``E-mail Server'' description 95 is replaced with

5 ``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 application, or by developing an Instant Messaging application which contains the plug-in functionality.

10 Alternatively, it will be appreciated that the interception could occur at the protocol level, intercepting network packets before they leave the user machine, or as they arrive at the user machine, or at some intermediate point on the network.

15 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 content analysis is impractical, however means to

20 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 within the user machine, or at some intermediate point

25 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 implemented by providing web browsers, e-mail clients,

30 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.

35

Claims

1. An information management system comprising:
one or more workstations adapted for
connection to a computer network, each workstation
having a memory;

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 containing rules specifying an
appropriate encryption strength for outbound data, the
encryption strength depending on content of the data;
and

an analyser, said analyser being operable in
conjunction with said policy data to monitor said
outbound data and to determine, in accordance with said
rules in said policy data, an appropriate encryption
strength for the outbound data;

wherein said analyser controls transmission
of said outbound data from said application in
dependence upon said determination of the appropriate
encryption strength.

2. The system of claim 1, wherein said rules in
said policy data define confidential data which should
not be transmitted, said analyser being operable in
conjunction with said policy data to prevent said
confidential data being transmitted from said
application.

- 100 -

3. The system of claim 1 or 2, wherein said analyser is further operable to determine present encryption strength in use for transmitting said outbound data; and

wherein said analyser controls transmission of said outbound data from said application both in dependence upon said determination of the appropriate encryption strength and in dependence upon said determination of the present encryption strength in use.

4. The system of claim 3, wherein if the analyser determines that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, then said analyser prevents transmission of said outbound data from said application.

5. The system of claim 3, wherein if the analyser determines that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, then said analyser prevents transmission of said outbound data from said application and controls said application to renegotiate an encryption strength for transmission that is appropriate.

6. The system of claim 3, wherein if the analyser determines that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, then said

- 101 -

analyser modifies the outbound data such that the present encryption strength is the appropriate encryption strength for the transmission of the modified outbound data.

7. The system of claim 3, wherein if the analyser determines that the present encryption strength in use for transmitting the outbound data is less than said appropriate encryption strength, then said analyser controls said application to notify a user of said application that the encryption strength in use is not sufficient.

8. The system of any one of claims 1 to 7, wherein the analyser is further operable to identify credit card numbers in said outbound data.

9. The system of claim 8, wherein the analyser is further operable to distinguish a predetermined set of credit card numbers from other credit card numbers, wherein said rules of said policy data define different appropriate encryption strengths for outbound data containing credit card numbers in the predetermined set than for other credit card numbers.

10. The system of claim 9, wherein said rules of said policy data specify that there is no appropriate encryption strength for a selected set of one or more credit card numbers.

- 102 -

11. The system of any one of claims 1 to 10, wherein said analyser is further operable to identify at least one or more of, credit card numbers, account codes, usernames, passwords, names and addresses and other predetermined keywords in the content of said outbound data.

12. The system of any one of claims 1 to 11, wherein said rules in said policy data specify an appropriate encryption strength for said outbound data, that is dependent on an address to which said outbound data is to be transmitted.

13. The system of any one of claims 1 to 12, wherein said analyser is located on each of said one or more workstations.

14. The system of any one of claims 1 to 13, wherein said application is a web browser.

15. The system of claim 14, wherein said analyser is a plug-in module of said web browser.

16. The system of claim 15, wherein said web browser is Microsoft's™ Internet Explorer™ and said analyser is a Browser Helper Object.

17. The system of any one of claims 1 to 13, wherein said application is an e-mail client.

- 103 -

18. The system of claim 17, wherein said analyser is a plug-in module of said e-mail client.

19. The system of claim 18, wherein said e-mail client is Microsoft's™ Outlook e-mail client and said analyser is a Microsoft™ client extension.

20. The system of any of claims 1 to 13, wherein said application is an Instant messaging application.

21. The system of claim 20, wherein said analyser is a plug-in module of said Instant messaging application.

22. The system of any one of claims 1 to 12, wherein said network comprises a server and said analyser is located at a point on said network intermediate said one or more workstations and said server, or said analyser is located at said server.

23. The system of any one of claims 1 to 22, 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.

24. The system of any one of claims 1 to 23, 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.

- 104 -

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

providing one or more workstations adapted for connection to a computer network, each workstation having a memory;

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 containing rules specifying an appropriate encryption strength for outbound data, the encryption strength depending on content of the data; and

analysing said outbound data to determine, in accordance with said rules in said policy data, an appropriate encryption strength for the outbound data;

controlling transmission of said outbound data from said application in dependence upon the determination of the appropriate encryption strength in said analysing step.

26. The method of claim 25, wherein said rules in said policy data define confidential data which should not be transmitted, and wherein in said controlling step transmission of said confidential data is prevented.

27. The method of claim 25 or 26, wherein said analysing step further comprising the step of determining present encryption strength in use for transmitting said outbound data; and

- 105 -

wherein in said controlling step the transmission of said outbound data from said application is dependent upon both the determination of the appropriate encryption strength and the determination of the present encryption strength in use.

28. The method of claim 27, wherein if it is determined that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, then in said controlling step transmission of said outbound data from said application is prevented.

29. The method of claim 27, wherein if in said analysing step it is determined that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, then in said controlling step an encryption strength appropriate for transmission of said outbound data is negotiated before transmission.

30. The method of claim 27, wherein if in said analysing step it is determined that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, then in said controlling step the outbound data is modified such that the present encryption strength is the appropriate encryption strength.

- 106 -

31. The method of claim 27, wherein in said analysing step if it is determined that the present encryption strength in use for transmitting the outbound data is less than said appropriate encryption strength, then in said controlling step a user of said application is notified that the encryption strength in use is not sufficient.

32. The method of any of claims 25 to 31, wherein said analysing step includes identifying credit card numbers in said outbound data.

33. The method of claim 32, wherein said analysing step includes distinguishing a pre-determined set of one or more credit card numbers from other credit card numbers, and wherein said rules of said policy data define a different appropriate encryption strength for outbound data containing credit card numbers in that pre-determined set than for other credit card numbers.

34. The method of claim 33, wherein said rules of said policy data specify that there is no appropriate encryption strength for said pre-determined set of one or more credit card numbers.

35. The method of any of claims 25 to 34, wherein said analysing step includes identifying at least one or more of, credit card numbers, account codes, usernames, passwords, names and addresses and other

- 107 -

predetermined keywords in the content of said outbound data.

36. The method of any of claims 25 to 35, wherein said rules in said policy data specify an appropriate encryption strength for said outbound data, that is dependent on an address to which said outbound data is to be transmitted.

37. The method of any of claims 25 to 36, wherein said analysing step is performed at said one or more workstations.

38. The method of any of claims 25 to 37, wherein said application is a web browser.

39. The method of claim 38 wherein said analysing step is performed by a plug-in module of said web browser.

40. The method of claim 39 wherein said web browser is Microsoft's™ Internet Explorer™ and said plug-in module is a Browser Helper Object.

41. The method of any of claims 25 to 37, wherein said application is an e-mail client.

42. The method of claim 41, wherein said analysing step is performed by a plug-in module of said e-mail client.

- 108 -

43. The method of claim 42, wherein said e-mail client is Microsoft's™ Outlook e-mail client and said plug-in module is a Microsoft™ Exchange client extension.

44. The method of any of claims 25 to 37, wherein said application is an Instant messaging application.

45. The method of claim 44, wherein said analyser is a plug-in module of said Instant messaging application.

46. The method of any of claims 25 to 36, wherein said network comprises a server and said analysing step is performed at a point on said network intermediate said one or more workstations and said server, or said analysing step is performed at said server.

47. The method of any of claims 25 to 46, 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.

48. The method of any of claims 25 to 47, 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.

- 109 -

49. A computer program product for controlling a computer connected to a public network to manage information, the computer having access to policy data containing rules specifying an appropriate encryption strength for outbound data transmitted to the public network, the encryption strength depending on content of the data, comprising:

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

determine, in conjunction with an application running on the computer that is operable at least to transmit outbound data to said public network, with reference to said rules in said policy data, an appropriate encryption strength for the outbound data; and

control the transmission of said outbound data by said application in dependence upon the determination of the appropriate encryption strength.

50. The computer program product of claim 49, wherein said rules in said policy data define confidential data which cannot be transmitted, wherein said program code when executed on said computer is operable to prevent transmission of said confidential data from said application.

51. The computer program product of claim 49 or 50, wherein said program code when executed on said computer is further operable to determine present

- 110 -

encryption strength in use for transmitting said outbound data; and

to control the transmission of said outbound data from said application in dependence upon both the determination of an appropriate encryption strength and the determination of the present encryption strength in use.

52. The computer program product of claim 51, wherein said program code when executed on said computer is further operable, if it is determined that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, to prevent the transmission of said outbound data from said application.

53. The computer program product of claim 51, wherein said program code when executed on said computer is further operable, if it is determined that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, to negotiate the appropriate encryption strength for transmission of said outbound data before transmission.

54. The computer program product of claim 51, wherein said program code when executed on said computer is further operable, if it is determined that the present encryption strength in use for transmitting outbound data is less than said appropriate encryption strength, to modify the outbound data such that the

- 111 -

present encryption strength is the appropriate encryption strength.

55. The computer program product of claim 51, wherein said program code when executed on said computer is further operable, if it is determined that the present encryption strength in use for transmitting the outbound data is less than said appropriate encryption strength, to provide notification that the encryption strength in use is not sufficient.

56. The computer program product of any of claims 49 to 55, wherein said program code when executed on said computer is further operable to identify credit card numbers in said outbound data.

57. The computer program product of claim 56, wherein said program code when executed on said computer is further operable to identify a pre-determined set of one or more credit card numbers from other credit card numbers, and wherein said rules of said policy data define a different appropriate encryption strength for outbound data containing credit card numbers in that pre-determined set than for other credit card numbers.

58. The computer program product of claim 57, wherein said rules of said policy data specify that there is no appropriate encryption strength for said pre-determined set of one or more credit card numbers.

- 112 -

59. The computer program product of any of claims 49 to 58, wherein said program code when executed on said computer is further operable, to identify at least one or more of, credit card numbers, account codes, usernames, passwords, names and addresses and other predetermined keywords in the content of said outbound data.

60. The computer program product of any of claims 49 to 59, wherein said rules in said policy data specify an appropriate encryption strength for said outbound data, that is dependent on an address to which said outbound data is to be transmitted.

61. The computer program product of any of claims 49 to 60, wherein said program code is executable on said computer.

62. The computer program product any of claims 49 to 61, wherein said application is a web browser.

63. The computer program product of claim 62, wherein said program code when executed on said computer is a plug-in module of said web browser.

64. The computer program product of claim 63, wherein said web browser is Microsoft's Internet Explorer™ and said plug-in module is a Browser Helper Object.

- 113 -

65. The computer program product of any of claims 49 to 61, wherein said application is an e-mail client.

66. The computer program product of claim 65, wherein said program code when executed on said computer is a plug-in module of said e-mail client.

67. The computer program product of claim 66, wherein said e-mail client is Microsoft's™ Outlook e-mail client and said plug-in module is a Microsoft Exchange client extension.

68. The computer software product of any of claims 49 to 61, wherein said application is an Instant messaging application.

69. The computer software product of claim 68, wherein said program code when executed on said computer is a plug-in module of said Instant messaging application.

70. The computer program product of any of claims 49 to 60, 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 program code is executable at said server.




FIG. 1




FIG. 2

3 / 24




FIG. 3

4 / 24




FIG. 4




FIG. 5




FIG. 6

7 / 24




FIG. 7




FIG. 8

9 / 24




FIG. 9




FIG. 10

11 / 24




FIG. 11

12 / 24




FIG. 12




FIG. 13




FIG. 14




FIG. 15




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




FIG. 17




FIG. 18




FIG. 19

19 / 24




FIG. 20





FIG. 21



Example of Table j
confidential
secret
contract
quotation
order
project x

FIG. 22




FIG. 23




FIG. 24




FIG. 25

Back-Office Application Plug-Ins

