(12) STANDARD PATENT (11) Application No. **AU 2011223614 B2**
(19) AUSTRALIAN PATENT OFFICE

(54) Title
Information protection using zones

(51) International Patent Classification(s)
G06F 15/00 (2006.01) ***G06F 21/00*** (2006.01)

(21) Application No: **2011223614** (22) Date of Filing: **2011.03.02**

(87) WIPO No: **WO11/109543**

(30) Priority Data

(31) Number	(32) Date	(33) Country
12/718,843	**2010.03.05**	**US**

(43) Publication Date: **2011.09.09**
(44) Accepted Journal Date: **2014.07.03**

(71) Applicant(s)
Microsoft Corporation

(72) Inventor(s)
Panasyuk, Anatoliy;Bablani, Girish;McColgan, Charles;Parthasarathy, Krishna Kumar

(74) Agent / Attorney
Davies Collison Cave, Level 15 1 Nicholson Street, MELBOURNE, VIC, 3000

(56) Related Art
US 6,366,912 B1
US 6,073,142 A

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
9 September 2011 (09.09.2011)

PCT

(10) International Publication Number
WO 2011/109543 A3

(51) International Patent Classification:
G06F 21/24 (2006.01) *G06F 21/00* (2006.01)
G06F 21/20 (2006.01) *G06F 15/00* (2006.01)

(21) International Application Number:
PCT/US2011/026898

(22) International Filing Date:
2 March 2011 (02.03.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/718,843 5 March 2010 (05.03.2010) US

(71) Applicant *(for all designated States except US)*: **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **PANASYUK, Anatoliy**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **BABLANI, Girish**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **MCCOLGAN, Charles**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **PARTHASARATHY, Krishna Kumar**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *with international search report (Art. 21(3))*

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

(88) Date of publication of the international search report:
12 January 2012

(54) Title: INFORMATION PROTECTION USING ZONES

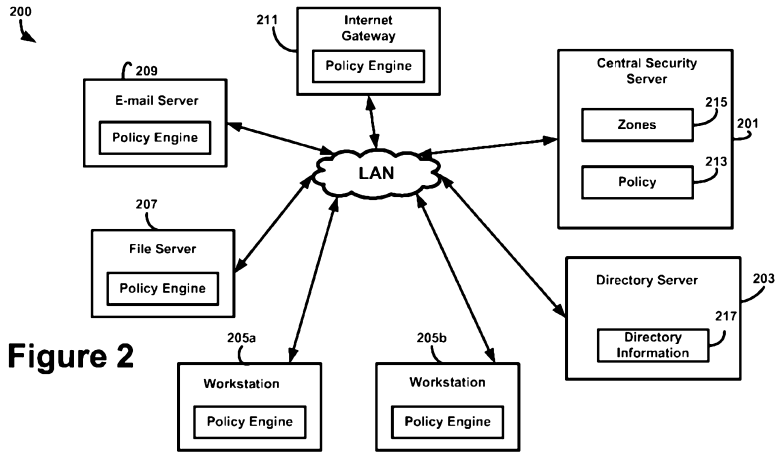

Figure 2

(57) Abstract: Some embodiments are directed to an information protection scheme in which devices, users, and domains in an information space may be grouped into zones. When information is transferred across a zone boundary, information protection rules may be applied to determine whether the transfer should be permitted or blocked, and/or whether any other policy actions should be taken (e.g., requiring encryption, prompting the user for confirmation of the intended transfer, or some other action).

WO 2011/109543 A3

INFORMATION PROTECTION USING ZONES

BACKGROUND

[0001] Within an organization, information is frequently created and shared. For example, workers create and send e-mails both to other workers in the organization and people outside of the organization. In addition, workers create documents, upload these documents to internal file servers, transfer them to portable storage media (e.g., removable flash memory drives), and send them to other users outside of the organization.

[0002] Some of the information created by workers in an organization may be confidential or sensitive. Thus, it may be desired to allow workers in possession of such information to only share it with those authorized to access it and/or to reduce the risk of workers accidentally transferring such information to someone who is not authorized to access it.

SUMMARY

[0003] The inventors have recognized that when information is shared, it may sometimes be sent to someone not authorized or not intended to have access to it or may be maliciously intercepted by someone not authorized to access it.

[0004] Thus, some embodiments are directed to an information protection scheme in which devices, users, and domains in an information space may be grouped into zones. When information is transferred across a zone boundary, information protection rules may be applied to determine whether the transfer should be permitted or blocked, and/or whether any other policy actions should be taken (e.g., requiring encryption, prompting the user for confirmation of the intended transfer, or some other action).

[0005] One embodiment is directed to a method for information protection performed by a computer comprising at least one processor and at least one tangible memory, the computer operating in an information space comprising a plurality of zones of users, devices, and/or domains, wherein each of the plurality of zones is a logical grouping of users, devices, and/or domains, and wherein the method comprises: in response to initiation of a transfer of information, determining whether the transfer of information would cause the information to cross a zone boundary between two of the plurality of zones; when it is determined that the transfer would not cause the information to cross the zone boundary, permitting the transfer; when it is determined that the transfer would cause the information to cross the zone boundary: accessing information protection rules; applying the information protection rules to the transfer to determine whether a

policy action is to be performed; and when it is determined the policy action is to be performed, performing the policy action.

[0006] Another embodiment is directed to at least one computer readable medium encoded with instructions that when executed on a computer comprising at least one processor and at least one tangible memory, perform a method in an information space comprising a plurality of zones of users, device, and/or domains, wherein each of the plurality of zones is a logical grouping of users, devices, and/or domains, wherein the computer is grouped into one of the plurality of zones, the method comprising: creating a document at the computer; automatically determining a first classification for the document; embedding information identifying the determined first classification into the document; receiving user input identifying a second classification for the document; in response to the user input, overriding the first classification with the second classification by removing the information identifying the first classification from the document and embedding information identifying the second classification into the document.

[0007] A further embodiment is directed to a computer in a computer system comprising: at least one tangible memory; and at least one hardware processor that executes processor-executable instructions to: in response to user input of first information that groups users, devices, and/or domains into logical zones, storing the first information in the at least one tangible memory; in response to user input of second information specifying information protection rules to be applied in response to initiation of a transfer of information that would cause the information to cross a boundary between logical zones, storing the second information in the at least one tangible memory.

[0007a] In a first broad form the present invention seeks to provide a method for information protection performed by a computer of an organization, the method comprising:

operating the computer in an information space in which users and devices within the organization are logically grouped into a plurality of zones of the organization;

automatically selecting, by an application program on the computer, a first classification for a document created on the computer using the application program, wherein the first classification is based on classification rules defined by an administrator of the organization and criteria comprising at least one of a zone of the organization into which the computer is grouped or a zone of the organization into which a user of the computer is grouped, and wherein the first classification is indicative of confidential information;

embedding, by the application program, the first classification into the document;

receiving, from the user of the computer, user input selecting a second classification for the document, wherein the second classification is manually selected by the user from available classifications configured by the administrator of the organization and is indicative of non-confidential information;

upon determining that the user is permitted to override the first classification, removing the first classification from the document and embedding the second classification into the document;

in response to the user of the computer initiating a transfer of the document from the computer, determining whether the transfer would cause the document to cross a zone boundary between zones of the organization; and

when it is determined that the transfer would cause the document to cross the zone boundary:

accessing information protection rules defined by the administrator of the organization;

applying one or more of the information protection rules to the transfer based on a zone of the organization from which the document is being transferred, a zone of the organization to which the document is being transferred, and the second classification embedded in the document;

determining a policy action to be performed on the document based on the one or more information rules applied to the transfer; and

performing the determined policy action on the document before permitting the transfer of the document from the computer.

[0007b] Typically the method further comprises receiving zone information from a security server, the zone information indicating a first zone of the organization from which the document is being transferred and a second zone of the organization to which the document is being transferred.

[0007c] Typically the security server is a separate device from the computer.

[0007d] Typically the method further comprises:

determining whether the first zone of the organization and the second zone of the organization are the same zone;

when it is determined that the first zone and the second zone are the same zone, determining that the transfer would not cause the document to cross the zone boundary; and

2011223614 28 Mar 2014

when it is determined that the first zone and the second zone are not the same zone, determining that the transfer would cause the information to cross the zone boundary.

[0007e] Typically accessing the information protection rules comprises accessing the information protection rules from a security server that is a separate device from the computer.

[0007f] Typically the second classification assigned by the user of the computer can be overridden by a manager of the user of the computer.

[0007g] Typically the first classification is further based on at least one of:

content of the document; and

keywords or patterns of text in the document.

[0007h] Typically the second classification indicates that the document contains personal information.

[0007i] Typically the determined policy action comprises at least one of:

audit logging the transfer of the document; and

encrypting the document.

[0007j] Typically the determined policy action comprises at least one of:

sending an alert of the transfer of the document to the administrator; or

creating a copy of the document.

[0007k] Typically the determined policy action comprises prompting the user of the computer for confirmation of the transfer of the document.

[0007l] Typically the plurality of zones includes different departments within the organization.

[0007m] Typically the application program classifies the document each time the document is saved.

[0007n] Typically the document is treated as being transferred from the zone into which the user is grouped if the user and the computer are grouped into different zones of the organization.

[0007o] Typically:

initiating the transfer of the document comprises uploading the document from the computer to a file server internal to the organization; and

the computer and the file server are grouped into different zones of the organization.

2011223614 28 Mar 2014

[0007p] Typically the classification rules defined by the administrator of the organization are applied to e-mail messages that are sent from the computer and files that are uploaded from the computer to a file server internal to the organization.

[0007q] In a second broad form the present invention seeks to provide at least one computer readable storage medium encoded with instructions that, when executed, cause a computer to perform a method of information protection, the method comprising:

operating the computer in an information space in which users and devices within an organization are logically grouped into a plurality of zones of the organization;

automatically selecting, by an application program on the computer, a first classification for a document created on the computer using the application program, wherein the first classification is based on classification rules defined by an administrator of the organization and criteria comprising at least one of a zone of the organization into which the computer is grouped or a zone of the organization into which a user of the computer is grouped, and wherein the first classification is indicative of confidential information;

embedding, by the application program, the first classification into the document;

receiving, from the user of the computer, user input selecting a second classification for the document, wherein the second classification is manually selected by the user from available classifications configured by the administrator of the organization and is indicative of non-confidential information;

upon determining that the user is permitted to override the first classification, removing the first classification from the document and embedding the second classification into the document;

in response to the user of the computer initiating a transfer of the document from the computer, determining whether the transfer would cause the document to cross a zone boundary between zones of the organization; and

when it is determined that the transfer would cause the document to cross the zone boundary:

accessing information protection rules defined by the administrator of the organization;

applying one or more of the information protection rules to the transfer based on a zone of the organization from which the document is being transferred, a zone of the organization to which the document is being transferred, and the second classification embedded in the document;

2011223614 28 Mar 2014

determining a policy action to be performed on the document based on the one or more information rules applied to the transfer; and

performing the determined policy action on the document before permitting the transfer of the document from the computer.

[0007r] Typically the first classification is further based, at least in part, on content of the document.

[0007s] Typically the document is created from a template, and wherein the first classification is further based, at least in part, on the template.

[0007t] In a third broad form the present invention seeks to provide a computer comprising:

at least one hardware processor that executes processor-executable instructions; and

memory storing processor-executable instructions including instructions for:

operating the computer in an information space in which users and devices within an organization are logically grouped into a plurality of zones of the organization;

automatically selecting, by an application program on the computer, a first classification for a document created on the computer using the application program, wherein the first classification is based on classification rules defined by an administrator of the organization and criteria comprising at least one of a zone of the organization into which the computer is grouped or a zone of the organization into which a user of the computer is grouped, and wherein the first classification is indicative of confidential information;

embedding, by the application program, the first classification into the document;

receiving, from the user of the computer, user input selecting a second classification for the document, wherein the second classification is manually selected by the user from available classifications configured by the administrator of the organization and is indicative of non-confidential information;

upon determining that the user is permitted to override the first classification, removing the first classification from the document and embedding the second classification into the document;

2011223614 28 Mar 2014

in response to the user of the computer initiating a transfer of the document from the computer, determining whether the transfer would cause the document to cross a zone boundary between zones of the organization; and

when it is determined that the transfer would cause the document to cross the zone boundary:

accessing information protection rules defined by the administrator of the organization;

applying one or more of the information protection rules to the transfer based on a zone of the organization from which the document is being transferred, a zone of the organization to which the document is being transferred, and the second classification embedded in the document;

determining a policy action to be performed on the document based on the one or more information rules applied to the transfer; and

performing the determined policy action on the document before permitting the transfer of the document from the computer.

BRIEF DESCRIPTION OF DRAWINGS

[0008] The accompanying drawings are not intended to be drawn to scale. In the drawings, each identical or nearly identical component that is illustrated in various figures is represented by a like numeral. For purposes of clarity, not every component may be labeled in every drawing. In the drawings:

[0009] Figure 1 is a block diagram of an information space logically divided into a plurality of zones, in accordance with some embodiments;

[0010] Figure 2 is a block diagram of computer system in which information protection techniques of embodiments of the invention may be implemented;

[0011] Figure 3 is a flow chart of a process for providing information protection in an information space logically divided into zones, in accordance with some embodiments; and

[0012] Figure 4 is a block diagram of a computer system on which aspects of some embodiments may be implemented.

DETAILED DESCRIPTION

[0013] The inventors have recognized that when workers in an organization create and/or access confidential or sensitive electronic information, situations may arise in which workers unwittingly or maliciously jeopardize the security of that information. For example, a worker may unintentionally send electronic information to someone who is not authorized to access that information or may store the electronic information in an insecure place (e.g., a file server which is accessible to someone unauthorized to access the information). As another example, a worker may share confidential electronic information in plain text (rather than encrypting it), thereby putting it at greater risk of being intercepted by someone not authorized to access it, or may take other actions that jeopardize the security of the information.

[0014] Thus, some embodiments are directed to a computer system in which users and devices are divided into logical groups called "zones." When electronic information is transferred from a user or device in one zone to a user or device in another zone, the information is considered to have crossed a zone boundary. When a transfer of information is initiated that would cause the information to cross a zone boundary, information control rules may be applied to determine whether the transfer is permitted or whether some action is to be taken before the transfer is permitted (e.g., prompting the worker initiating the transfer, audit logging the transfer, requiring encryption of the information before allowing the transfer, or some other action).

[0015] In some embodiments, the information control rules may take into account the type of information that is being transferred. For example, different information control rules may be applied when attempting to transfer confidential information from a first zone to a second zone than when attempting to transfer non-confidential information from the first zone to the second zone. Thus, in some embodiments, when electronic information is generated, it may be tagged (e.g., automatically, semi-automatically, or manually) with a classification indicative of the sensitivity of the information and/or other properties of the information. The classification rules may take into account the classification of electronic information and the zone to which and from which the information is being transferred when the information is attempted to be transferred across a zone boundary.

[0016] This technique may provide a number of benefits. First, it allows one uniform security policy to be defined and applied across multiple different channels. That is, the same set of classification rules may be applied to transfer of e-mails, transfer of content through the world wide web, file transfer to a file server internal to the organization, and/or to any other type of electronic information or information channel. Second, it allows the information control rules to be customized based on the type of information to which the rules are applied so that a restrictive set of rules that might be warranted for sensitive or confidential information need not be applied to information for which such a restrictive set of rules is not warranted.

[0017] A number of problems with the prior art and a number of benefits provided by the above-discussed techniques are identified above. However, the invention is not limited to addressing any or all of these problems or providing any or all of these benefits. That is, while some embodiments may address some or all of these problems and provide some or all of these benefits, some embodiments may not address any of these problems or provide any of these benefits.

[0018] Figure 1 shows an example of an information space that may be classified into zones. As shown in Figure 1, an organization 100 may have a computer system comprising a number of devices. Some of these devices may be used by an engineering department of the organization and some may be used by a public relations department. Because documents or other pieces of content from the engineering department likely include a significant amount of confidential and/or sensitive information, while documents or other pieces of content generated in the public relations department are less likely to include such information, the devices used by the engineering department may be grouped into one zone and the devices used by the public relations department may be grouped into another zone. Thus, as shown in Figure 1, though all the devices in the organization are physically connected via local area network (LAN) 125, engineering file server 103, engineering e-mail server 105, and workstations 107a, 107b, and 107c may be logically grouped in Engineering Department zone 101, while PR file server 109, PR e-mail server 111, and workstations 113a, 113b, and 113d are logically grouped together in PR Department Zone 115.

[0019] In addition, in the example of Figure 1, an organization 121 that is external to organization 100 may be logically grouped into a zone. For example, if organization 121 is a trusted partner of organization 100, it may be desired to apply different information control rules to organization 121, such that information sent to and received

from organization 121 (e.g., via Internet 117) is treated differently from that of other entities external to organization 100. Thus, organization 121 may be logically grouped into Trusted Partner zone 119, while information sent to and received from other entities external to organization 100 (e.g., via Internet 117) may be treated as being sent to and received from general Internet zone 123. As discussed above, when information is sent from one zone to another zone, information protection rules may be applied and action may be taken based on the information protection rules, if warranted.

[0020] In the example of Figure 1, devices within organization 100 are logically grouped into two zones. It should be appreciated that this is merely illustrative as an organization may comprise any suitable number of zones. For example, all devices and users within an organization may be grouped into a single zone or these devices and users may be grouped into three or more different zones. In addition, in the example of Figure 1, only devices are shown as being logically grouped into zones. However, users (e.g., employees of organization 100, other workers, or other persons) or domains may also be logically grouped into zones. For example, employees of organization 100 who work in the engineering department may be grouped into Engineering Department zone 101 and employees who work in the PR department may be grouped into PR Department zone 115.

[0021] As such, the inventors have recognized that a situation may arise where a user that is grouped into one zone is using a device that is grouped into a different zone. Thus, when the user sends information from or receives information at that device, the information may be treated as having been sent from or received at either the zone of the user or the zone of the device. Thus, for example, if an employee of the engineering department who is grouped into the Engineering Department zone logs in and works from workstation 113a, which is grouped in the PR Department zone, the employee may attempt to upload a document to engineering file server 103. This document may be treated as either being sent from the Engineering Department zone or the PR Department zone.

[0022] In some embodiments, the zone of the user may take precedence over the zone of the device which the user is using. Thus, in the example above, when the engineering department employee uploads a document to engineering file server 103 from workstation 113a, the document may be treated as being sent from the Engineering Department zone to the Engineering Department zone (i.e., not crossing a zone boundary). However, the invention is not limited in this respect as, in some embodiments, the zone of the device may take precedence over the zone of the user using the device, and in some

embodiments whether the user's zone or the device's zone takes precedence may be configured by an administrator of the organization.

[0023] As discussed above, the information protection rules may define whether and what actions are to be performed when information is transferred across a zone boundary based on the zone to which the information is being transferred, the zone from which the information is being transferred, and the classification of the information being transferred. Information may be classified in any of a variety of ways and classification of information may be performed at any of a variety of points in the information creation and sharing process. For example, classification may be performed, automatically, semi-automatically, or manually, and may be performed when the information is created, when the information is stored, when the information is transferred, and/or at any other suitable time.

[0024] For example, in some embodiments, when an application program is used to create a document (e.g., an e-mail or other document), the application program may automatically classify the document. The application program may classify the document based on any suitable criteria or criterion. For example, the application program may automatically classify the document based on the zone into which the user and/or device has been grouped or based on keywords or patterns in the document. Thus, for example, documents that include certain keywords or patterns of text may be assigned certain classifications. In some embodiments, documents may be classified by hashing the document using a hash function (e.g., SHA1 or any other suitable hash function), comparing the hash value to a set of stored hash values, and assigning a classification to the documents based on the comparison. In some embodiments, documents may be classified using fuzzy matching the employs shingling techniques to represent the fuzzy hashing of documents (or portions of documents) for similarity detection. In some embodiments, a document may be classified based on a template from which the document was created, or may be assigned a default classification associated with that application program used to create or edit the document or some other default classification. The application program may classify the document upon initial creation of the document, each time the document is saved, when the document is completed, and/or any other suitable time.

[0025] In some embodiments, instead of or in addition to the application program used to create a document performing classification, classification may be performed by an information protection agent or other software program executing on the computer used

to create the document. Such a software program may perform classification of a document based on any of the criteria (or any combination of the criteria) discussed above, and may perform classification of the document at any suitable time after initial creation of the document For example, such an agent or other software program may classify documents stored on the computer as background process, may classify documents upon initiation a transfer of the documents outside of the computer, or at any other suitable point in time.

[0026] In the examples above, documents are classified on the computer on which they are created. However, the invention is not limited in this respect as, in some embodiments, a document may be classified by an entity that receives the document. For example, if a document is transferred, the device that receives the document may perform classification of the document before applying information control rules to determine, for example, whether the transfer is permitted and should be completed or is not permitted and should be dropped. For example, an e-mail client executing on a workstation may send an e-mail to an e-mail server in the organization for transmission to the intended recipients. In some embodiments, the e-mail server may perform classification of the e-mail. In addition, e-mails or other documents received from an entity external to the organization may not be classified until they are received by a device within the organization, as the external entities may not use the same information protection model to classify documents. Thus, classification may be performed on these documents after they are received within the organization. For example, an e-mail server may perform classification of e-mails received from external senders, or an internal file server may perform classification of documents uploaded from external senders.

[0027] Once the appropriate classification for a document has been determined, the classification may be stored in any of a variety of ways. In some embodiments, the classification may be embedded (e.g., as a tag or label) in the document itself. For example, the classification of an e-mail may be embedded in the e-mail header, and the classification of other types of document may be embedded in metadata included in the document.

[0028] In the examples discussed above, classification of documents is performed automatically. However, the invention is not limited in this respect, as in some embodiments, classification of documents may be performed semi-automatically, such that a classification may be assigned to a document automatically, but a user has the ability to override the automatic classification and assign a different classification to the document.

[0029] In some embodiments, policies may be defined that indicate which users are authorized to assign classification to documents and which users are allowed to override a previously-assigned classification. For example, in some embodiments, a subsequent user may be permitted to override a previously-assigned classification by an initial user, if the subsequent user is a manager or boss of the initial user. The determination as to whether the subsequent user is a manager or boss of the initial user may be made, for example, using organizational chart (org chart) information stored in the directory information of a directory server.

[0030] In some embodiments, classification of documents may be performed manually, such that users manually specify the classification that is to be assigned to each document. In such embodiments, if a document for which a classification has not been assigned is transferred across a zone boundary, it may be assigned a default classification so that the information protection rules may be applied.

[0031] Any suitable classification scheme may be used to classify documents. In some embodiments, the classifications that are available to be assigned to a document may be configured by an administrator of the organization. Examples of classifications that may be used include, "Company Confidential," "Personal," "Non-Confidential," "Financial Data," and/or any other suitable classification.

[0032] Figure 2 is a block diagram of a computer system 200 for an organization in which information protection rules based on zones and information classification may be employed. Computer system 200 comprises a central security server 201, which stores zone information 215 and policy information 213. Zone information 215 indicates the zones that have been defined (e.g., by a network administrator) and the devices, users, and/or domains that are grouped into each of the defined zones. Policy information 213 specifies the information protection rules (e.g., that have been defined by an administrator) that are to be applied when information is transmitted across a zone boundary.

[0033] Computer system 200 may also include a directory server 203 that stores directory information 217. Directory information 217 includes information about users of and devices in the computer system. In addition, directory information may define organizational units or groups of users and devices. For example, directory information 217 may define an "Engineering Group" that includes users and/or devices in the engineering department and may define a "PR Group" that includes users and/or devices in the PR department.

[0034] In some embodiments, directory information 217 may be used to group users, devices, and/or domains into zones. For example, zone information 215 may be configured to indicate that every user or device in the "Engineering Group" is grouped into the "Engineering Department" zone and every user or device in the "PR Group" is grouped into the "PR Department" zone.

[0035] The inventors have recognized that when an entity (e.g., an organization) is an external to the organization operating computer system 200, an administrator of computer system 200 may not have access to directory information identifying the users and devices of the external organization. Thus, if it is desired to group the external organization into a zone, the domain name of the organization may be used. For example, if an external organization named "Contoso, Inc." uses the domain name "contoso.com," and it is desired to group this organization into a zone (e.g., a "Trusted Partner" zone), then the zone information may identify the domain name "contoso.com" as belonging to this zone. In some embodiments, directory information 217 may define a group of Trusted Partners that includes the domain names of external entities, and the zone information may indicate that all of the domain names in that group are grouped into a particular zone (e.g., the "Trusted Partner" zone).

[0036] Computer system 200 may also include a number of other devices. For example, in Figure 2, computer system 200 includes an e-mail server 209, a file server 207, workstations 205a and 205b, and Internet gateway 211. Internet gateway 211 may serve as a gateway to the Internet for the devices in computer system 200, and the devices in computer system 200 may communicate with each other via local area network (LAN) 218.

[0037] Devices 205a, 205b, 207, 209, and 211 each include a policy engine. The policy engine on each of these devices may operate when information is received from another device or is being sent to another device to determine when the information has crossed or would, if transmitted, cross a zone boundary. If so, the policy engine may determine based on the information protection rules, whether any policy action is warranted, and may perform the policy action.

[0038] In the example of Figure 2, each of devices 205a, 205b, 207, 209, and 211 executes a policy engine. However, the invention is not limited in this respect. That is, in some embodiments, only those devices that are at a zone boundary (e.g., devices that are capable of directly transmitting information to or receiving information from another zone) may execute a policy engine. Thus, if such embodiments were employed in the

example of Figure 2, and if all of the devices and users in computer system 200 were grouped into a single zone, then only Internet gateway 211 need execute a policy engine.

[0039] Figure 3 shows an illustrative information protection process that may be used in a computer system such as computer system 200 to implement information protection rules. The process begins at act 301, where a piece of content (e.g., a document) is created or received. The process next continues to act 303, where the piece of content is classified and the classification for the piece of content is stored.

[0040] After act 303, the process continues to act 305, where transfer of the piece of content to another device is initiated. The process next continues to act 307, where it is determined if the transfer causes or would cause the piece of content to cross a zone boundary. Act 307 may be performed, for example, by a policy engine on the device which is initiating sending the piece of content or on another device that receives the piece of content after it has been transmitted from the device which initiated the transfer.

[0041] The policy engine may determine whether the transfer causes or would cause the information to cross a zone boundary in any of a variety of ways. For example, in some embodiments, the policy engine may communicate with the central security server 201 (which, as discussed above, stores zone information 215) to determine the zone of the device or user that initiated the transfer and the zone of the device or user that is the intended recipient of the transfer. Alternatively, in some embodiments, all or portions of this zone information may be cached locally on the device, and the policy engine may use the locally cached information to determine the zone of the device or user that initiated the transfer and the zone of the device or user that is the intended recipient. If the zone of the device or user that initiated the transfer and the zone of the device or user that is the intended recipient of the piece of content are the same, it may be determined that the transfer does not cause the piece of content to cross a zone boundary, and the process may end.

[0042] If the zone of the device or user that initiated the transfer and the zone of the device or user that is the intended recipient of the piece of content are different, it may be determined that the transfer causes or would cause the piece of content to cross a zone boundary, and the process may continue to act 309. At act 309, the policy engine may determine whether any policy actions are to be taken as a result of the intended transfer and perform the policy actions. The policy engine may determine whether any policy actions are to be taken in any suitable way. For example, the policy engine may communicate with the central security server 201 to determine the information protection

rules stored in policy information 213, and may apply these rules to the transfer in question. Alternatively, in some embodiments, all or some of the rules stored in policy information 213 may be cached locally on the device, and the policy engine may use the locally cached information to determine the classification rules.

[0043] The classification rules may specify any suitable policy action based on the classification rules. For example, the policy engine may block the transfer, require encryption of the content to complete the transfer, create an audit log entry of the transfer, prompt the user for confirmation before completing the transfer, create a copy of the information desired to be transferred, send an alert to a user or an administrator notifying him or her of the transfer, and/or take any other suitable action.

[0044] Figure 4 shows a schematic block diagram of an illustrative computer 400 on which aspects of the invention may be implemented. Only illustrative portions of the computer 400 are identified for purposes of clarity and not to limit aspects of the invention in any way. For example, the computer 400 may include one or more additional volatile or non-volatile memories (which may also be referred to as storage media), one or more additional processors, any other user input devices, and any suitable software or other instructions that may be executed by the computer 400 so as to perform the function described herein.

[0045] In the illustrative embodiment, the computer 400 includes a system bus 410, to allow communication between a central processing unit 402 (which may include one or more hardware general purpose programmable computer processors), a tangible memory 404, a video interface 406, a user input interface 408, and a network interface 412. The network interface 412 may be connected via network connection 420 to at least one remote computing device 418. Peripherals such as a monitor 422, a keyboard 414, and a mouse 416, in addition to other user input/output devices may also be included in the computer system, as the invention is not limited in this respect.

[0046] In some embodiments, the devices illustrated and described above may be implemented as computers, such as computer 400. For example, in some embodiments, devices 201, 203, 205a, 205b, 207, 209, and 211 may each be implemented as a computer, such as computer 400. In this respect, it should be appreciated that the above-described functionality of these devices may be implemented by central processing unit 402 executing software instructions to perform this functionality, and that information described above as being stored on these devices may be stored in memory 404.

[0047] Having thus described several aspects of at least one embodiment of this invention, it is to be appreciated that various alterations, modifications, and improvements will readily occur to those skilled in the art.

[0048] Such alterations, modifications, and improvements are intended to be part of this disclosure, and are intended to be within the spirit and scope of the invention. Accordingly, the foregoing description and drawings are by way of example only.

[0049] The above-described embodiments of the present invention can be implemented in any of numerous ways. For example, the embodiments may be implemented using hardware, software or a combination thereof. When implemented in software, the software code can be executed on any suitable processor or collection of processors, whether provided in a single computer or distributed among multiple computers.

[0050] Further, it should be appreciated that a computer may be embodied in any of a number of forms, such as a rack-mounted computer, a desktop computer, a laptop computer, or a tablet computer. Additionally, a computer may be embedded in a device not generally regarded as a computer but with suitable processing capabilities, including a Personal Digital Assistant (PDA), a smart phone or any other suitable portable or fixed electronic device.

[0051] Also, a computer may have one or more input and output devices. These devices can be used, among other things, to present a user interface. Examples of output devices that can be used to provide a user interface include printers or display screens for visual presentation of output and speakers or other sound generating devices for audible presentation of output. Examples of input devices that can be used for a user interface include keyboards, and pointing devices, such as mice, touch pads, and digitizing tablets. As another example, a computer may receive input information through speech recognition or in other audible format.

[0052] Such computers may be interconnected by one or more networks in any suitable form, including as a local area network or a wide area network, such as an enterprise network or the Internet. Such networks may be based on any suitable technology and may operate according to any suitable protocol and may include wireless networks, wired networks or fiber optic networks.

[0053] Also, the various methods or processes outlined herein may be coded as software that is executable on one or more processors that employ any one of a variety of operating systems or platforms. Additionally, such software may be written using any of a

number of suitable programming languages and/or programming or scripting tools, and also may be compiled as executable machine language code or intermediate code that is executed on a framework or virtual machine.

[0054] In this respect, the invention may be embodied as a computer readable medium (or multiple computer readable media) (e.g., a computer memory, one or more floppy discs, compact discs (CD), optical discs, digital video disks (DVD), magnetic tapes, flash memories, circuit configurations in Field Programmable Gate Arrays or other semiconductor devices, or other non-transitory, tangible computer storage medium) encoded with one or more programs that, when executed on one or more computers or other processors, perform methods that implement the various embodiments of the invention discussed above. The computer readable medium or media can be transportable, such that the program or programs stored thereon can be loaded onto one or more different computers or other processors to implement various aspects of the present invention as discussed above.

[0055] The terms "program" or "software" are used herein in a generic sense to refer to any type of computer code or set of computer-executable instructions that can be employed to program a computer or other processor to implement various aspects of the present invention as discussed above. Additionally, it should be appreciated that according to one aspect of this embodiment, one or more computer programs that when executed perform methods of the present invention need not reside on a single computer or processor, but may be distributed in a modular fashion amongst a number of different computers or processors to implement various aspects of the present invention.

[0056] Computer-executable instructions may be in many forms, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically the functionality of the program modules may be combined or distributed as desired in various embodiments.

[0057] Also, data structures may be stored in computer-readable media in any suitable form. For simplicity of illustration, data structures may be shown to have fields that are related through location in the data structure. Such relationships may likewise be achieved by assigning storage for the fields with locations in a computer-readable medium that conveys relationship between the fields. However, any suitable mechanism may be used to establish a relationship between information in fields of a data structure, including

through the use of pointers, tags or other mechanisms that establish relationship between data elements.

[0058] Various aspects of the present invention may be used alone, in combination, or in a variety of arrangements not specifically discussed in the embodiments described in the foregoing and is therefore not limited in its application to the details and arrangement of components set forth in the foregoing description or illustrated in the drawings. For example, aspects described in one embodiment may be combined in any manner with aspects described in other embodiments.

[0059] Also, the invention may be embodied as a method, of which an example has been provided. The acts performed as part of the method may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

[0060] Use of ordinal terms such as “first,” “second,” “third,” etc., in the claims to modify a claim element does not by itself connote any priority, precedence, or order of one claim element over another or the temporal order in which acts of a method are performed, but are used merely as labels to distinguish one claim element having a certain name from another element having a same name (but for use of the ordinal term) to distinguish the claim elements.

[0061] Also, the phraseology and terminology used herein is for the purpose of description and should not be regarded as limiting. The use of "including," "comprising," or "having," “containing,” “involving,” and variations thereof herein, is meant to encompass the items listed thereafter and equivalents thereof as well as additional items

[0062] The reference in this specification to any prior publication (or information derived from it), or to any matter which is known, is not, and should not be taken as an acknowledgment or admission or any form of suggestion that the prior publication (or information derived from it) or known matter forms part of the common general knowledge in the field of endeavour to which this specification relates.

[0063] Throughout this specification and claims which follow, unless the context requires otherwise, the word “comprise”, and variations such as “comprises” or “comprising”, will be understood to imply the inclusion of a stated integer or group of integers or steps but not the exclusion of any other integer or group of integers.

2011223614 28 Mar 2014

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method for information protection performed by a computer of an organization, the method comprising:

operating the computer in an information space in which users and devices within the organization are logically grouped into a plurality of zones of the organization;

automatically selecting, by an application program on the computer, a first classification for a document created on the computer using the application program, wherein the first classification is based on classification rules defined by an administrator of the organization and criteria comprising at least one of a zone of the organization into which the computer is grouped or a zone of the organization into which a user of the computer is grouped, and wherein the first classification is indicative of confidential information;

embedding, by the application program, the first classification into the document;

receiving, from the user of the computer, user input selecting a second classification for the document, wherein the second classification is manually selected by the user from available classifications configured by the administrator of the organization and is indicative of non-confidential information;

upon determining that the user is permitted to override the first classification, removing the first classification from the document and embedding the second classification into the document;

in response to the user of the computer initiating a transfer of the document from the computer, determining whether the transfer would cause the document to cross a zone boundary between zones of the organization; and

when it is determined that the transfer would cause the document to cross the zone boundary:

accessing information protection rules defined by the administrator of the organization;

applying one or more of the information protection rules to the transfer based on a zone of the organization from which the document is being transferred, a zone of the organization to which the document is being transferred, and the second classification embedded in the document;

determining a policy action to be performed on the document based on the one or more information rules applied to the transfer; and

performing the determined policy action on the document before permitting the transfer of the document from the computer.

2. The method of claim 1, further comprising:

receiving zone information from a security server, the zone information indicating a first zone of the organization from which the document is being transferred and a second zone of the organization to which the document is being transferred.

3. The method of claim 2, wherein the security server is a separate device from the computer.

4. The method of claim 2 or claim 3, further comprising:

determining whether the first zone of the organization and the second zone of the organization are the same zone;

when it is determined that the first zone and the second zone are the same zone, determining that the transfer would not cause the document to cross the zone boundary; and

when it is determined that the first zone and the second zone are not the same zone, determining that the transfer would cause the information to cross the zone boundary.

5. The method of any one of claims 1 to 4, wherein accessing the information protection rules comprises:

accessing the information protection rules from a security server that is a separate device from the computer.

6. The method of any one of claims 1 to 5, wherein:

the second classification assigned by the user of the computer can be overridden by a manager of the user of the computer.

7. The method of any one of claims 1 to 6, wherein the first classification is further based on at least one of:

content of the document; and

keywords or patterns of text in the document.

8. The method of any one of claims 1 to 7, wherein the second classification indicates that the document contains personal information.

9. The method of any one of claims 1 to 8, wherein the determined policy action comprises at least one of:

audit logging the transfer of the document; and

encrypting the document.

10. The method of any one of claims 1 to 9, wherein the determined policy action comprises at least one of:

sending an alert of the transfer of the document to the administrator; and

creating a copy of the document.

11. The method of any one of claims 1 to 10, wherein the determined policy action comprises:

prompting the user of the computer for confirmation of the transfer of the document.

12. The method of any one of claims 1 to 11, wherein the plurality of zones includes different departments within the organization.

13. The method of any one of claims 1 to 12, wherein the application program classifies the document each time the document is saved.

14. The method of any one of claims 1 to 13, wherein the document is treated as being transferred from the zone into which the user is grouped if the user and the computer are grouped into different zones of the organization.

15. The method of any one of claims 1 to 14, wherein:

initiating the transfer of the document comprises uploading the document from the computer to a file server internal to the organization; and

the computer and the file server are grouped into different zones of the organization.

16. The method of any one of claims 1 to 15, wherein the classification rules defined by the administrator of the organization are applied to e-mail messages that are sent from the computer and files that are uploaded from the computer to a file server internal to the organization.

17. At least one computer readable storage medium encoded with instructions that, when executed, cause a computer to perform a method of information protection, the method comprising:

operating the computer in an information space in which users and devices within an organization are logically grouped into a plurality of zones of the organization;

automatically selecting, by an application program on the computer, a first classification for a document created on the computer using the application program, wherein the first classification is based on classification rules defined by an administrator of the organization and criteria comprising at least one of a zone of the organization into which the computer is grouped or a zone of the organization into which a user of the

computer is grouped, and wherein the first classification is indicative of confidential information;

embedding, by the application program, the first classification into the document;

receiving, from the user of the computer, user input selecting a second classification for the document, wherein the second classification is manually selected by the user from available classifications configured by the administrator of the organization and is indicative of non-confidential information;

upon determining that the user is permitted to override the first classification, removing the first classification from the document and embedding the second classification into the document;

in response to the user of the computer initiating a transfer of the document from the computer, determining whether the transfer would cause the document to cross a zone boundary between zones of the organization; and

when it is determined that the transfer would cause the document to cross the zone boundary:

accessing information protection rules defined by the administrator of the organization;

applying one or more of the information protection rules to the transfer based on a zone of the organization from which the document is being transferred, a zone of the organization to which the document is being transferred, and the second classification embedded in the document;

determining a policy action to be performed on the document based on the one or more information rules applied to the transfer; and

performing the determined policy action on the document before permitting the transfer of the document from the computer.

18. The at least one computer-readable storage medium of claim 17, wherein the first classification is further based, at least in part, on content of the document.

19. The at least one computer-readable storage medium of claim 17 or claim 18, wherein the document is created from a template, and wherein the first classification is further based, at least in part, on the template.

20. A computer comprising:

at least one hardware processor that executes processor-executable instructions; and

memory storing processor-executable instructions including instructions for:

operating the computer in an information space in which users and devices within an organization are logically grouped into a plurality of zones of the organization;

automatically selecting, by an application program on the computer, a first classification for a document created on the computer using the application program, wherein the first classification is based on classification rules defined by an administrator of the organization and criteria comprising at least one of a zone of the organization into which the computer is grouped or a zone of the organization into which a user of the computer is grouped, and wherein the first classification is indicative of confidential information;

embedding, by the application program, the first classification into the document;

receiving, from the user of the computer, user input selecting a second classification for the document, wherein the second classification is manually selected by the user from available classifications configured by the administrator of the organization and is indicative of non-confidential information;

upon determining that the user is permitted to override the first classification, removing the first classification from the document and embedding the second classification into the document;

in response to the user of the computer initiating a transfer of the document from the computer, determining whether the transfer would cause the document to cross a zone boundary between zones of the organization; and

when it is determined that the transfer would cause the document to cross the zone boundary:

accessing information protection rules defined by the administrator of the organization;

applying one or more of the information protection rules to the transfer based on a zone of the organization from which the document is being transferred, a zone of the organization to which the document is being transferred, and the second classification embedded in the document;

determining a policy action to be performed on the document based on the one or more information rules applied to the transfer; and

performing the determined policy action on the document before permitting the transfer of the document from the computer.

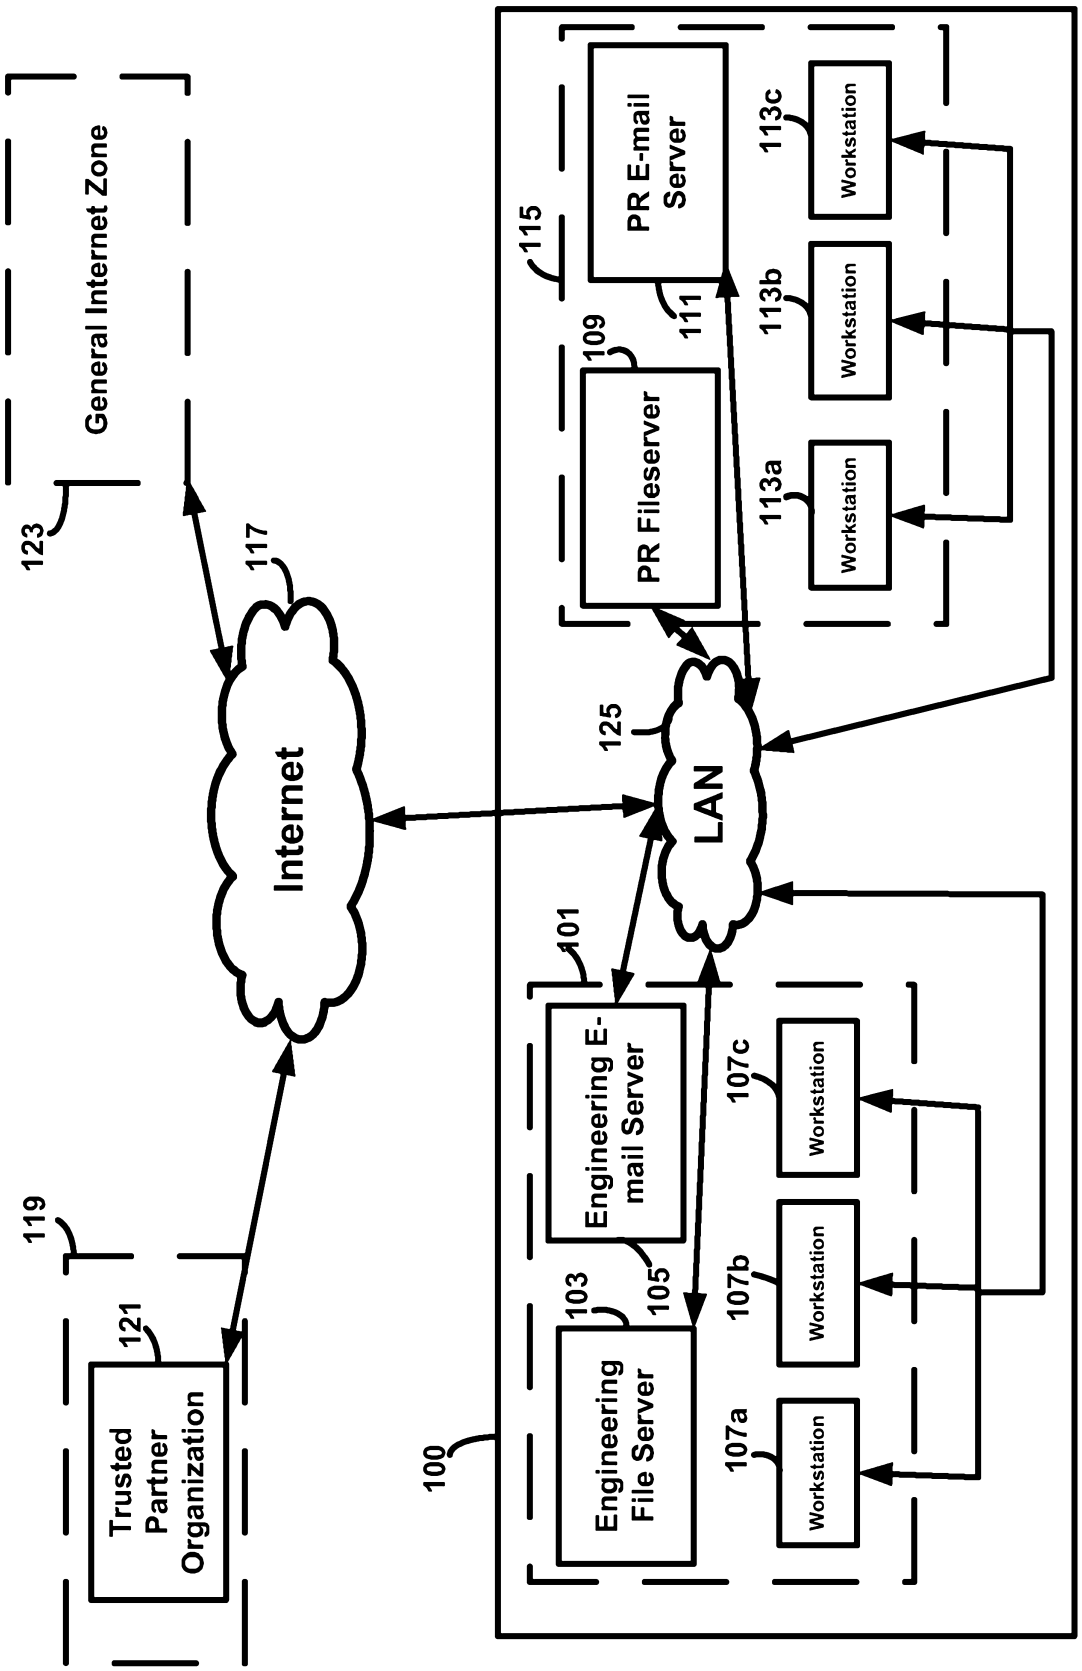
General Internet Zone
123
117
Internet
119
121
Trusted Partner Organization
100
125
LAN
115
PR E-mail Server
111
109
PR Fileserver
101
Engineering E-mail Server
105
103
Engineering File Server
113c
113b
113a
107c
107b
107a
Workstation

Figure 1

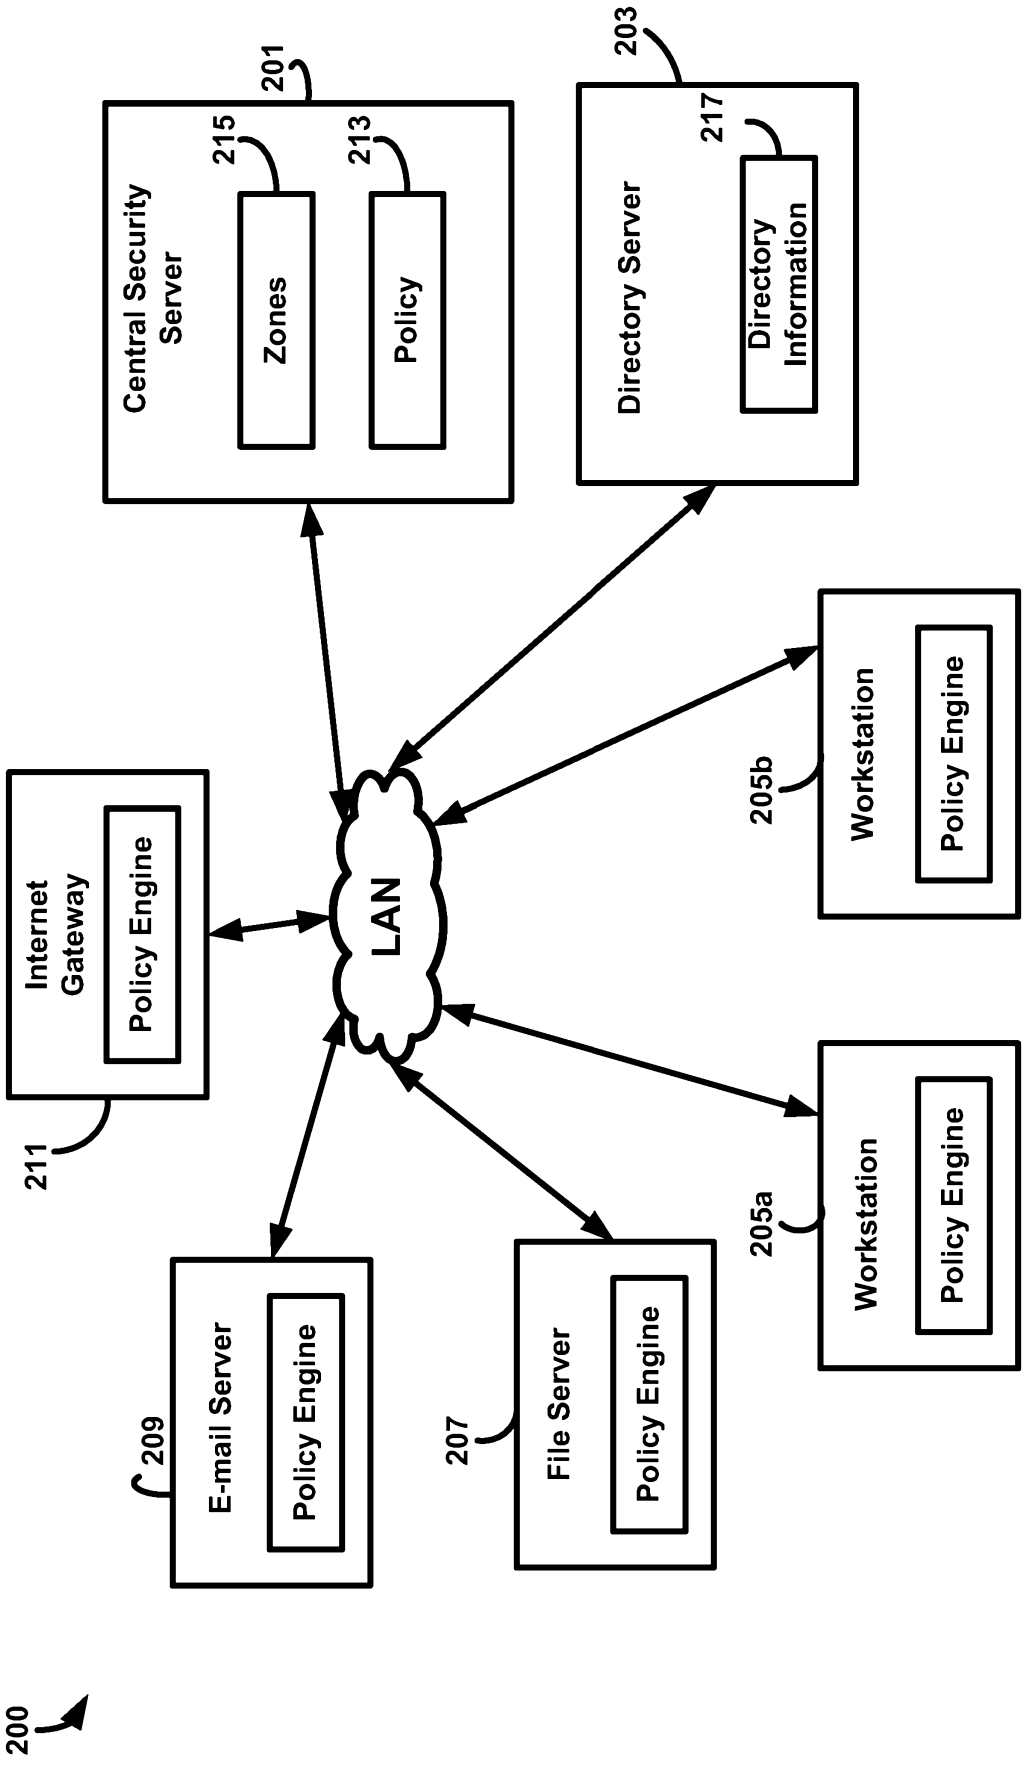
200
Central Security Server
201
Zones
215
Policy
213
Directory Server
203
Directory Information
217
LAN
Internet Gateway
Policy Engine
211
Workstation
Policy Engine
205b
Workstation
Policy Engine
205a
E-mail Server
Policy Engine
209
File Server
Policy Engine
207

Figure 2

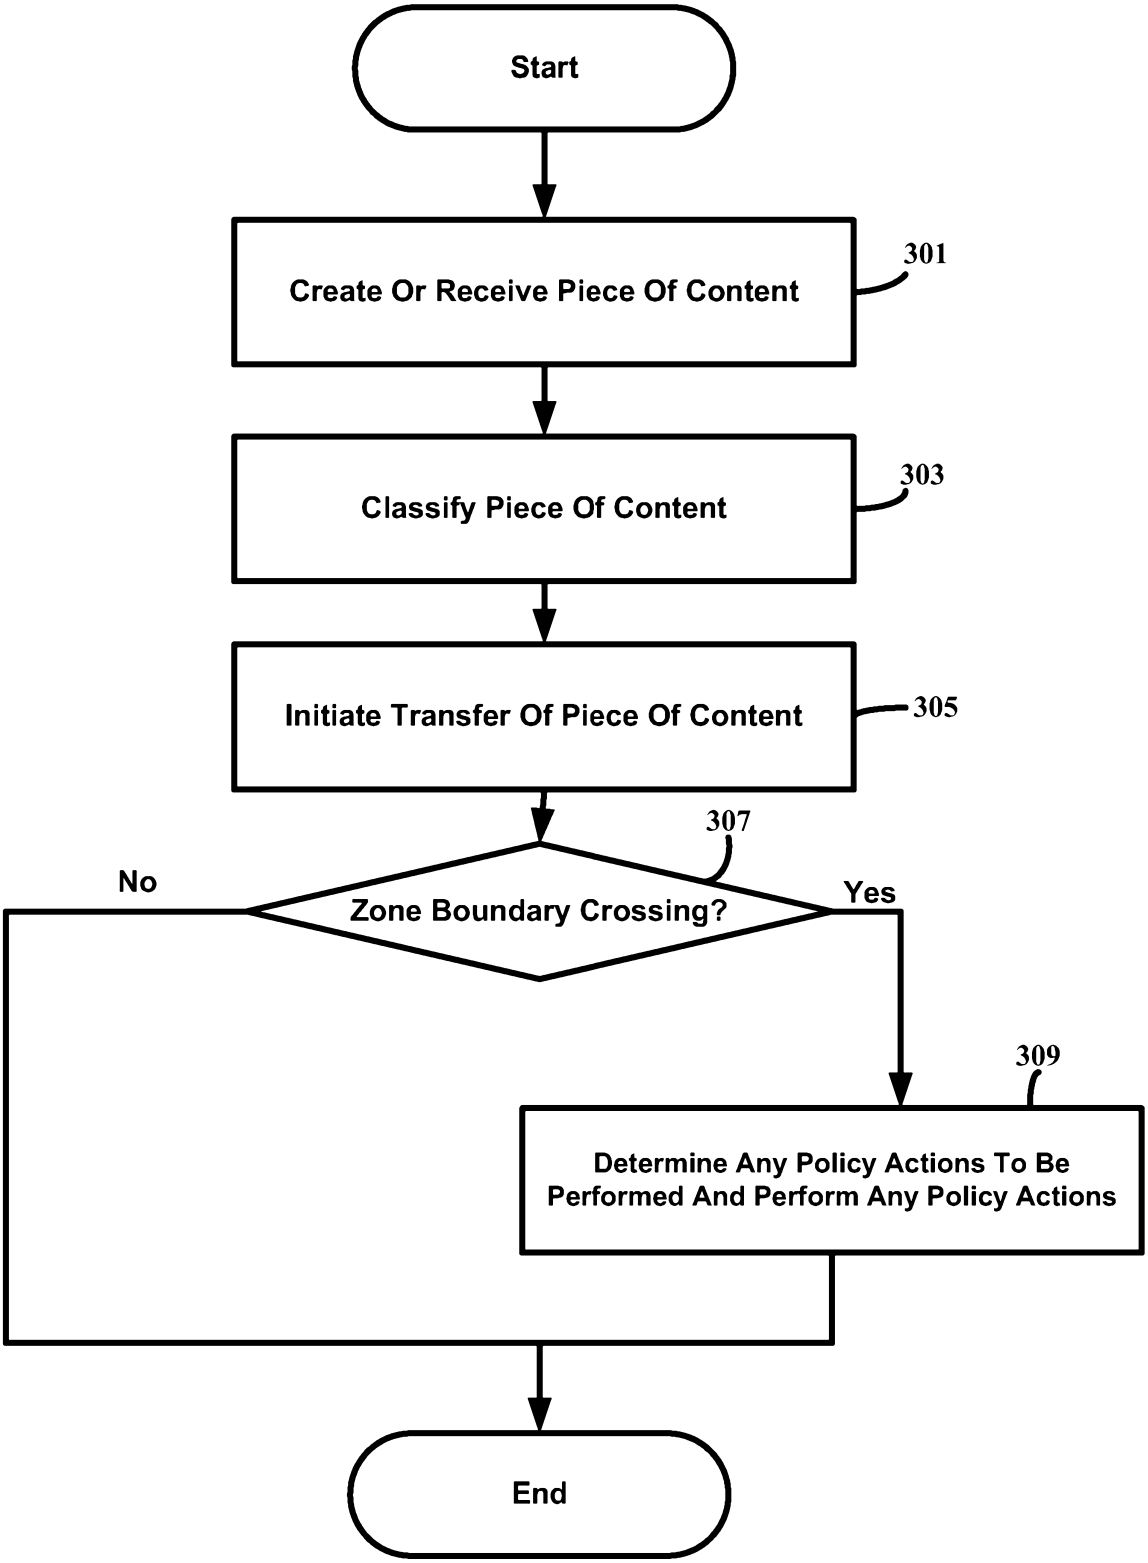
Start
Create Or Receive Piece Of Content
301
Classify Piece Of Content
303
Initiate Transfer Of Piece Of Content
305
307
Zone Boundary Crossing?
No
Yes
309
Determine Any Policy Actions To Be Performed And Perform Any Policy Actions
End

Figure 3

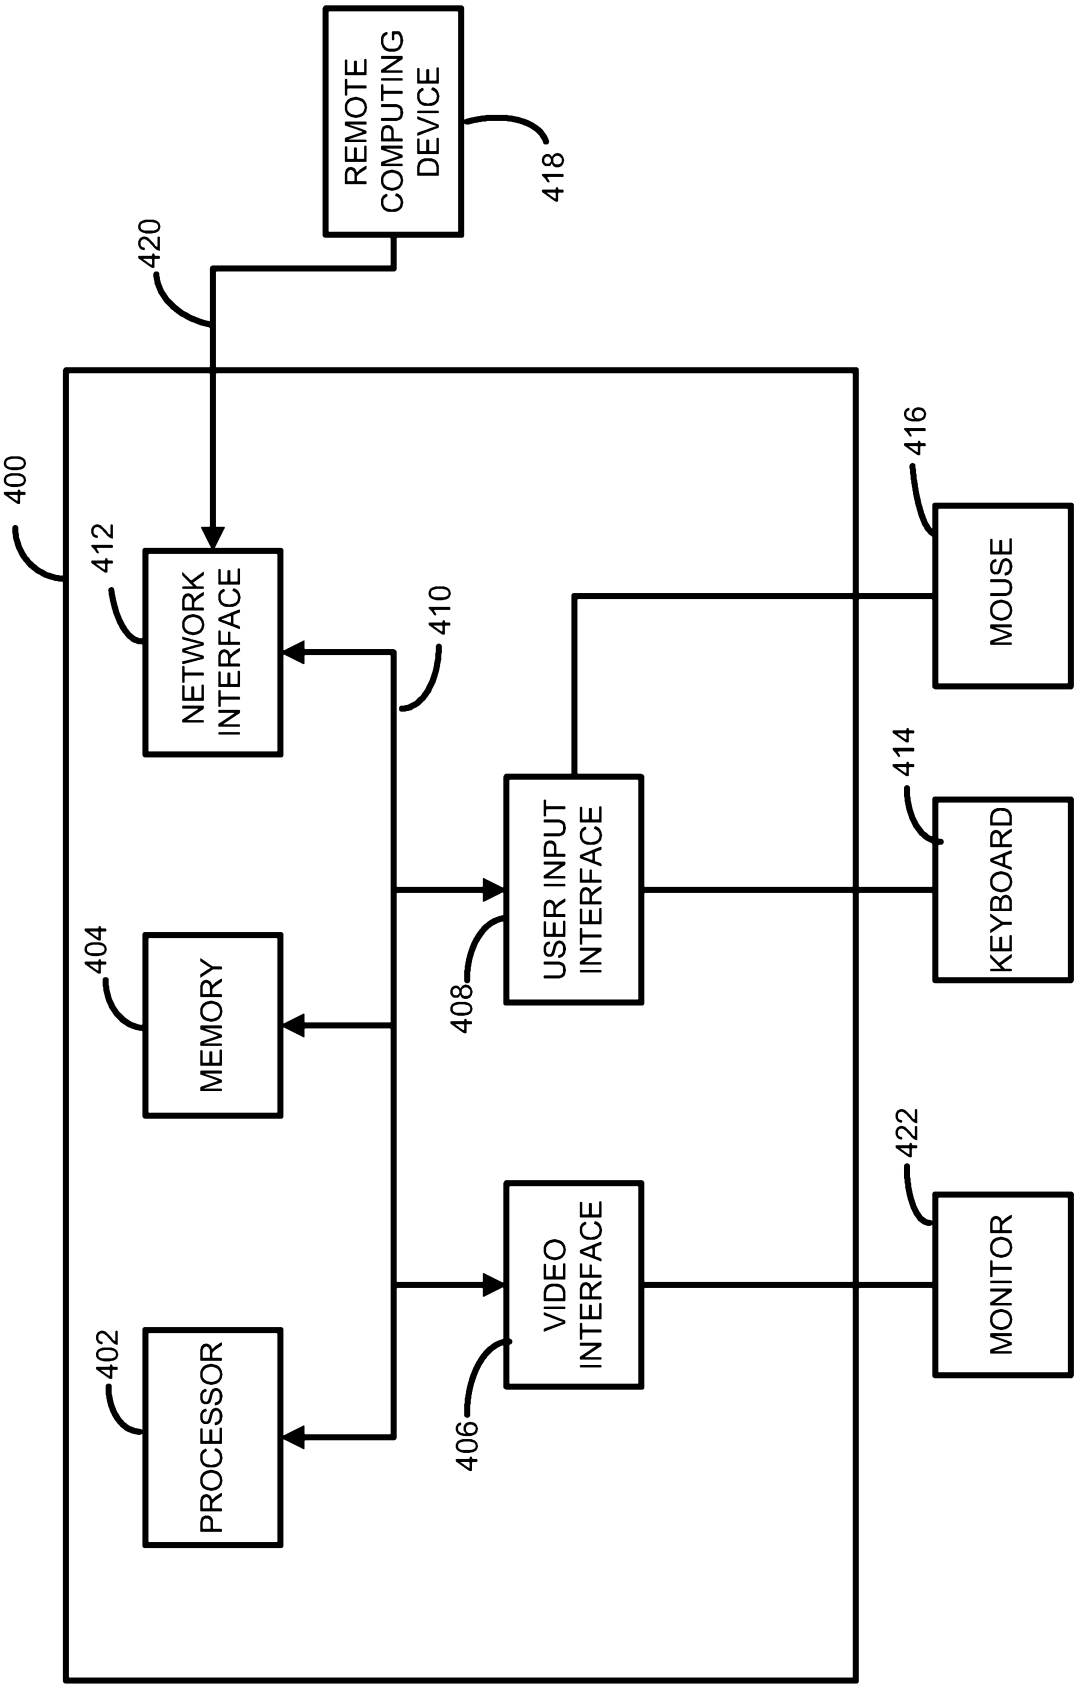
REMOTE COMPUTING DEVICE
418
420
400
412
NETWORK INTERFACE
410
404
MEMORY
408
USER INPUT INTERFACE
402
PROCESSOR
406
VIDEO INTERFACE
416
MOUSE
414
KEYBOARD
422
MONITOR

Figure 4