



US012519804B2

(12) **United States Patent**
Greevy

(10) **Patent No.:** US 12,519,804 B2
(45) **Date of Patent:** *Jan. 6, 2026

(54) **SYSTEM AND METHOD FOR VERIFYING THE IDENTITY OF EMAIL SENDERS TO IMPROVE EMAIL SECURITY WITHIN AN ORGANIZATION**

(71) Applicant: **Paibox, Inc.**, San Francisco, CA (US)

(72) Inventor: **Hoala Greevy**, San Francisco, CA (US)

(73) Assignee: **Paibox, Inc.**, San Francisco, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 167 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **18/374,270**

(22) Filed: **Sep. 28, 2023**

(65) **Prior Publication Data**

US 2024/0106835 A1 Mar. 28, 2024

Related U.S. Application Data

(63) Continuation-in-part of application No. 18/230,019, filed on Aug. 3, 2023, now Pat. No. 12,137,104, which is a continuation of application No. 17/127,930, filed on Dec. 18, 2020, now Pat. No. 11,765,185, which is a continuation of application

(Continued)

(51) **Int. Cl.**
H04L 9/40 (2022.01)

(52) **U.S. Cl.**
CPC **H04L 63/126** (2013.01); **H04L 63/08** (2013.01); **H04L 63/105** (2013.01)

(58) **Field of Classification Search**
CPC H04L 63/08; H04L 63/126; H04L 63/105;
H04L 63/1483; H04L 63/1433; H04L
63/101; H04L 67/306

USPC 726/4; 713/170
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,206,814 B2 * 4/2007 Kirsch H04L 63/145
709/217
7,366,761 B2 * 4/2008 Murray H04L 51/00
709/224

(Continued)

OTHER PUBLICATIONS

NPL Search Terms (Year: 2025).*


Primary Examiner — Syed A Zaidi

(74) *Attorney, Agent, or Firm* — Run8 Patent Group, LLC; Peter Miller; Leah Raddatz

(57) **ABSTRACT**

One variation of a method includes, intercepting an email addressed to a target recipient within an organization, the email received from a sender at an inbound email address and including a first inbound display name; accessing a whitelist including a set of contact information corresponding to an entity associated with the organization and including a verified display name associated with the entity and a set of verified email addresses associated with the entity. The method further includes, in response to the set of verified email addresses omitting the inbound email address; characterizing a display name difference between the inbound display name and the verified display name associated with the entity; and, in response to the display name difference falling below a threshold difference, withholding transmission of the email to the target recipient and flagging the email for authentication.

20 Claims, 12 Drawing Sheets



Related U.S. Application Data

No. 16/944,091, filed on Jul. 30, 2020, now Pat. No. 10,904,266.

(60) Provisional application No. 63/410,933, filed on Sep. 28, 2022, provisional application No. 62/880,511, filed on Jul. 30, 2019.

(56) References Cited**U.S. PATENT DOCUMENTS**

7,590,698	B1 *	9/2009	Cooley	G06Q 10/107 709/206
8,205,264	B1 *	6/2012	Kailash	H04L 67/02 709/224
8,510,388	B2 *	8/2013	Taylor	H04L 51/234 709/225
8,516,061	B2 *	8/2013	Seon	H04L 51/212 709/206
8,719,350	B2 *	5/2014	Sharma	G06Q 10/107 707/899
10,880,322	B1 *	12/2020	Jakobsson	H04L 51/08
11,500,853	B1 *	11/2022	Hamilton	G06F 16/245
11,729,211	B2 *	8/2023	Jakobsson	H04L 51/212 726/23
11,757,914	B1 *	9/2023	Jakobsson	H04L 51/42 726/25
2007/0208868	A1 *	9/2007	Kidd	G06Q 30/02 709/229
2016/0308809	A1 *	10/2016	Wood	H04L 51/212
2018/0089606	A1 *	3/2018	McNamara	G06Q 30/0201
2021/0218698	A1 *	7/2021	Murillo	H04L 51/48
2021/0311162	A1 *	10/2021	Mai	G01S 7/415

* cited by examiner




FIGURE 1




FIGURE 2






FIGURE 3

**FIGURE 4**

**FIGURE 5A**




FIGURE 5B




FIGURE 5C







FIGURE 6A

**FIGURE 6B**

**FIGURE 6C**






FIGURE 7B

**SYSTEM AND METHOD FOR VERIFYING
THE IDENTITY OF EMAIL SENDERS TO
IMPROVE EMAIL SECURITY WITHIN AN
ORGANIZATION**

**CROSS-REFERENCE TO RELATED
APPLICATIONS**

This Application claims the benefit of U.S. Provisional Application No. 63/410,933, filed on 28 Sep. 2022, which is incorporated in its entirety by this reference.

This Application is also a continuation-in-part of U.S. patent application Ser. No. 18/230,019, filed on 3 Aug. 2023, which is a continuation application of U.S. patent application Ser. No. 17/127,930, filed on 18 Dec. 2020, which is a continuation application of U.S. patent application Ser. No. 16/944,091, filed on 30 Jul. 2020, which claims the benefit of U.S. Provisional Application No. 62/880,511, filed on 30 Jul. 2019, each of which are incorporated in their entireties by this reference.

TECHNICAL FIELD

This invention relates generally to the field of email communications and more specifically to a new and useful method for verifying the identity of email senders in the field of email communications.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 is a flowchart representation of a method; FIG. 2 is a schematic representation of the method; FIG. 3 is a schematic representation of the method; FIG. 4 is a schematic representation of the method; FIGS. 5A, 5B, and 5C are schematic representations of the method; FIGS. 6A, 6B, and 6C are schematic representations of the method; and FIGS. 7A and 7B are schematic representations of the method.

DESCRIPTION OF THE EMBODIMENTS

The following description of embodiments of the invention is not intended to limit the invention to these embodiments but rather to enable a person skilled in the art to make and use this invention. Variations, configurations, implementations, example implementations, and examples described herein are optional and are not exclusive to the variations, configurations, implementations, example implementations, and examples they describe. The invention described herein can include any and all permutations of these variations, configurations, implementations, example implementations, and examples.

1. Method

As shown in FIGS. 1-4, a method S100 includes: intercepting a first email addressed to a first target recipient within an organization in Block S110, the first email received from a first sender at a first inbound email address and including a first inbound display name; accessing a whitelist including a first set of contact information corresponding to a first employee within the organization in Block S120, the first set of contact information including a first verified display name associated with the first employee and a first set of verified email addresses associated with the

first employee; characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first employee in Block S130; and, in response to the first display name difference falling below a threshold difference, comparing the first inbound email address to the first set of verified email addresses associated with the first employee in Block S150. The method S100 further includes: in response to identifying the first inbound email address in the first set of verified email addresses, authorizing transmission of the first email to the first target recipient in Block S160; and, in response to the first set of verified email addresses omitting the first inbound email address, withholding transmission of the first email to the first target recipient and flagging the first email for authentication in Block S170.

One variation of the method S100 includes: intercepting a first email addressed to a first target recipient within an organization in Block S110, the first email received from a first sender at a first inbound email address and including a first inbound display name; accessing a first set of contact information corresponding to a first employee within the organization in Block S120, the first set of contact information including a first set of verified email addresses associated with the first employee and a first verified display name associated with the first employee; comparing the first inbound email address to the first set of verified email addresses associated with the first employee in Block S130; and, in response to the first set of verified email addresses excluding the first inbound email address, characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first employee in Block S140. In this variation, the method S100 further includes: in response to the first display name difference exceeding the threshold difference, authorizing transmission of the first email to the first target recipient in Block S160; and, in response to the first display name difference falling below a threshold difference, withholding transmission of the first email to the first target recipient and flagging the first email for authentication in Block 170.

One variation of the method S100 includes: receiving an email addressed to a target recipient within an organization in Block S110, the email received from a sender at an inbound email address and including an inbound display name and an inbound domain; accessing a whitelist containing contact information of a set of employees within an organization in Block S120; comparing the inbound display name contained in the email to verified display names of employees contained in the whitelist in Block S130; and, in response to the inbound display name differing from a display name of each employee in the set of employees identified in the whitelist, sending the email to the target recipient in Block S140. The method S100 also includes, in response to the inbound display name matching a particular name (or display name) of a particular employee in the set of employees identified in the whitelist: comparing the inbound email address with a set of verified email addresses associated with the particular employee in the whitelist in Block S150; in response to the set of verified email addresses containing the inbound email address of the email, sending the email to the target recipient in Block S160; and, in response to the set of verified email addresses omitting the first inbound address, quarantining the email in Block S170. The method S100 can further include, in response to the set of verified email addresses omitting the first inbound address, prompting further investigation of the email by an email administrator.

Generally, Blocks of the method S100 can be executed by a computer system to verify the identity of a sender before passing an email to its designated recipient in order to detect and suppress email spoofing attempts. In particular, email clients (e.g., mobile email clients) may hide sender email addresses and display only display names of these senders in an inbox, email window, and/or reply window. A phisher may leverage hidden sender addresses in such email clients to spoof email recipients within an organization. For example, a spammer may: leverage an organization directory or social network to identify a keyman within the organization (e.g., a CEO, a CTO, a VP, a board member); create a false email account with a display name identical to the name of this keyman; and deliver an email to an employee or associate of the organization—from this false email account—with an urgent request, such as to provide organization login information or complete a purchase on behalf of the phisher. Because the sender address of this false email account may be hidden in the recipient's email client and because this email contains the keyman's display name, the recipient may perceive the email to be authentic; because the email includes an urgent request, the recipient may also be prompted to act quickly and thus allocate less time to considering authenticity of the request, all of which may result in the recipient completing the action requested in this email on behalf of the phisher.

Therefore, the computer system (e.g., an email server) can execute Blocks of the method to: generate and maintain an approved sender database (hereinafter a "whitelist") containing a keyman profile—including all verified displays names and verified email addresses—for each keyman within an organization; scan this whitelist for a display name of a keyman that matches an inbound display name attached to an inbound email into the organization; and pass this inbound email on to its designated recipient within the organization if no keyman profile for keymen in the organization contains a display name that matches the inbound display name of the inbound email. However, if the computer system identifies a particular keyman profile that contains a name (or display name) that matches the display name attached to the inbound email, the computer system can: deliver this inbound email to its recipient if the email address of the inbound email matches the email address associated with the matched display name in this particular keyman profile; and otherwise quarantine the email for additional investigation (e.g., by an email administrator within the organization) before delivery to the designated recipient. For example, the computer system can notify the email administrator of the flagged email—which represents a possible spoofing attempt on the organization—by: inserting a link to the email into an security alert feed at an information security operations center (or "ISOC"); or inserting hyperlinks to discard or deliver the email and contents of the email into an email notification and then forwarding this email notification to the email administrator's email inbox. The computer system can then selectively deliver the inbound email to its designated recipient or discard the email based on the email administrator's response.

The computer system can thus execute Blocks of the method to: protect the personal security of a select group of keymen within an organization by blocking emails impersonating these keymen; reduce email spoofing within the organization; and increase confidence and trust of employees working under the select group of keymen by verifying

authenticity of emails containing display names indicating these keymen to these employees.

3. Example

5 In one example, during a setup period, a particular employee at an organization may load her contact information into a keyman profile and upload this contact card to the computer system via a web portal. For example, the keyman profile can include work, personal, and/or education email addresses and related display names and work and personal phone numbers. Alternatively, the employee can send test emails from her work, personal, education, and/or other email accounts to the computer system, and the computer system can strip email addresses and related display names from these test emails and load this contact information into a keyman profile for the particular employee. The computer system can similarly interface with other employees, associates, or keymen of the organization to generate a corpus of 10 keyman profiles for the organization. The computer system can then compile these keyman profiles into a whitelist for the organization, such as including one keyman profile for each executive, board member, vice president, and upper-level manager of the organization.

15 Later, upon receiving an email from a sender (e.g., outside of the organization), the computer system can: determine whether a display name of the inbound email matches an employee name or display name contained in any one 20 keyman profile within the whitelist. If the display name of the inbound email matches the employee name or display name within a particular keyman profile, the computer system can compare a sender email address of the inbound email with a verified email address of the particular employee contained in the particular keyman profile. Then, if the sender email address matches one verified email address in this particular keyman profile, the computer system can deliver (or "release") the email to a recipient specified in the inbound email. However, if the sender email address does not match a verified email address in the particular keyman profile, the computer system can: withhold the email from the recipient; quarantine the email, such as by diverting the email to a quarantine database; and notify 25 an email administrator of the quarantined email. The email administrator may further investigate validity of the email and its sender and then determine whether to deliver the quarantined email to the recipient.

30 To notify the email administrator of the quarantined email, the computer system can: generate a notification email containing a hyperlink to access the quarantine database within a web portal; and deliver the notification email to the email administrator. Upon receiving the notification email, the email administrator can: select the hyperlink to 35 automatically open a web browser and to navigate to the web portal containing the quarantine database; view the quarantined email to determine the validity of the sender; and select whether to deliver the quarantined email to the recipient or 40 discard the quarantined email based on results of her investigation. Additionally and/or alternatively, the computer system can: populate a notification email with contents and 45 sender data of the inbound email and a hyperlink to release the email into the notification email; deliver the notification email to the email administrator; and automatically deliver the email to the recipient upon selection of this hyperlink by the email administrator (e.g., "one-click" release). Yet alternatively, the computer system can generate an alert—linked

to this quarantined email—and insert this alert into a security alert feed at an ISOC affiliated with the organization.

4. Onboarding & Whitelist

The computer system interfaces with employees, associates, or other representatives of an organization (e.g., a clinic, a law firm, a company) to access and aggregate contact information for these employees (e.g., via a web portal) and leverages contact information provided for these employees to investigate the validity of email senders sending emails to employees within the organization.

In one implementation, an email administrator within an organization identifies a select group of “keyman” or “very important people” within the organization who exhibit greater risk of email spoofing. For example, the email administrator may generate a list of organization keymen including a CEO, a COO, a president, vice presidents, and senior managers. The computer system can then verify that inbound emails that include display names that match (or are otherwise similar to) names of these keymen are inbound from email addresses associated with these keyman before releasing these emails to their designated recipients, as described below. Therefore, the computer system can interface with an email administrator during a setup period to aggregate contact information of keymen and/or entities in the organization and to compile the associated contact information into a whitelist of verified contact information for these keymen and/or entities within the organization.

In another implementation, the computer system interfaces with keymen in an organization directly to access contact information of these keymen. For example, the CEO of a first company may upload her contact information, including both work and personal contact information (e.g., name, email addresses, phone numbers). The computer system can then receive this contact information and generate a keyman profile for this CEO. Therefore, at a future time, the computer system can verify inbound emails containing display names similar or identical to the CEO’s names based on the CEO’s keyman profile. In particular, upon receiving an email with a display name and from a particular email address, the computer system can: access the whitelist containing keyman profiles for a group of employees included in the whitelist; and search for a keyman profile containing a name or display similar or identical to the display name contained in the inbound email. Then, upon identifying the CEO’s keyman profile that contains a name that matches the display name of the inbound email, the computer system can compare the sender email address of the inbound email to a set of verified email addresses contained in the CEO’s keyman profile. If the sender email address matches a verified email address in the CEO’s keyman profile, the computer system can verify the sender as the CEO and deliver the email to its designated recipient. Otherwise, the computer system can flag and quarantine this email.

In one variation, the whitelist is generated manually for a select group of employees or associates within the organization and is uploaded to the computer system via a web portal. For example, the computer system can initially receive a manually generated whitelist of contact information for a select group of employees, and at a later time incorporate contact information (e.g., names, email addresses, display names) of additional employees in the organization as emails are sent to and/or sent by these additional employees and approved by an email administrator, as described below.

In one variation, the computer system can add all email addresses within a particular domain to the whitelist for an organization. For example, the computer system can whitelist all email addresses of employees within an organization domain (e.g., Company ABCD with an email domain “@ABCD.com”). Therefore, the computer system can automatically authorize transmission of emails sent from email addresses within the organization domain without further checks for authenticity according to the method, thereby reducing overhead and computational power scanning these internal emails for spoofing attempts.

5. Inbound Email Check

15 Once the whitelist is generated, the computer system can receive (or “intercept”) inbound emails from senders and verify the validity of these senders based on this whitelist before releasing these inbound emails to their designated recipients. In particular, the computer system can receive an 20 inbound email from an inbound email address (hereinafter a “sender email address”) and extract contact information from the inbound email, including: a display name of the sender (e.g., a name of the sender that is visible to a recipient of the email); a username of the sender email address (e.g., a local component of the sender email address); and a domain of the sender email address. The computer system can then: compare the display name to verified names of employees contained in the whitelist; and compare the sender email address and/or its components (e.g., the user- 25 name) to verified email addresses contained in keyman or entity profiles in the whitelist if the display name of the sender email address matches the name or display name of a keyman or entity profile in the whitelist.

30 For example, upon receiving an email from a sender at a sender email address, the email including a display name and designating a target recipient, the computer system can: access a whitelist including names of employees at an organization; in response to the whitelist including a verified name of an employee that matches the display name included in the email, compare a set of verified email 35 addresses corresponding to the employee and included in the whitelist to the sender email address; in response to the whitelist containing the sender email address, deliver the email to the target recipient; in response to the whitelist omitting the sender email address, quarantine the email in a 40 quarantine database; and notify an email administrator of the email for further investigation. Therefore, the computer system can deliver emails sent with display names excluded from the whitelist without further investigation and flag emails with display names corresponding to employees within an organization for authentication in order to ensure the legitimacy of emails allegedly sent from key employees 45 within the organization.

50 In one variation, the computer system can initially access the whitelist to compare a sender email address with verified email addresses of employees in the whitelist. If the sender email address exactly matches a verified email address of an employee in the whitelist, the computer system can automatically authenticate the email and authorize transmission 55 of the email to a designated target recipient. For example, the computer system can: intercept an email addressed to a target recipient within an organization, the email received from a sender at a sender email address and including an inbound display name; access a set of contact information corresponding to an employee within the organization, the set of contact information including a set of verified email 60 addresses associated with the employee and a verified dis-

play name associated with the employee; compare the inbound email address to first set of verified email addresses associated with the employee. Then, in response to identifying the sender email address in the set of verified email addresses, the computer system can authorize transmission of the email to the target recipient. However, in response to the set of verified email addresses excluding the sender email address, the computer system can further compare the inbound display name with the verified display name to verify validity of the sender. Therefore, the computer system can automatically deliver emails sent from email addresses included in the whitelist without further investigation and further investigate emails sent from email addresses excluded from the whitelist.

In one variation, upon receiving a new email, the computer system can initially access a blacklist, containing email addresses of known fraudulent senders, to automatically withhold and discard emails from these email addresses. For example, the computer system can receive an email from an email address; access a blacklist containing blocked email addresses; in response to the blacklist containing the email address as a blocked email address, automatically withhold the email from a designated recipient and discard the email. Therefore, the computer system performs no further investigation if an email address is already contained in the blacklist and does not notify the email administrator, thus increasing efficiency of email investigation and reducing a workload of the email administrator. If the email address is not found on the blacklist, the computer system can implement the method S100 described above to continue the investigation process.

6. Sender Verification

The computer system can confirm email legitimacy and verify the identity of a sender by checking a display name visible to a user viewing the email and/or an email address from which the email is sent. For example, upon receiving an email from a sender at a first email address, the computer system can: identify the display name displayed in the email; access the whitelist containing contact information of a select group of employees within an organization; in response to the display name matching a verified name of a particular employee in the whitelist, compare the first email address to a set of verified email addresses contained in the whitelist for the particular employee; in response to the first email address matching a verified email address within the set of email addresses contained in the whitelist, deliver the email to a recipient as designated in the email; in response to the first email address not matching a verified email address within the set of email addresses, flag the email for quarantine; and notify an email administrator of the email for additional investigation.

Therefore, the computer system can quarantine emails including inbound display names matching and/or similar to display names included in the whitelist if the email address from which the email was sent is not also included in the whitelist. The computer system can then notify an email administrator of the quarantined email to further investigate the validity of the sender. For example, upon receiving an email with a display name matching an employee contained in the whitelist but from an email address not contained in the whitelist, the computer system can: quarantine the email in an email quarantine database and notify an email administrator of the email for further investigation. The computer

system can then receive inputs from the email administrator to determine whether to deliver the email to a designated recipient.

The computer system can notify the email administrator of a quarantined email from an unverified email address with a display name of an employee within an organization, and the email administrator can further investigate the validity of this sender. The email administrator may manually investigate the validity of the email address and the sender to confirm that the email is legitimate and not a phishing attempt by corroborating with the employee corresponding to the display name contained in the email. Alternatively, the email administrator may read the original email and determine its validity based on content contained in the email.

The computer system can characterize difference (or “display name differences”) between inbound display names and verified display names included in the whitelist to identify inbound display names that closely resemble (or “imitate”) display names of employees and/or entities included in the whitelist. For example, the computer system can characterize a display name difference on a scale of zero percent (e.g., inbound display name equivalent to verified display name) to 100 percent (e.g., no overlap between inbound display name and verified display name). In another example, the computer system can characterize the display name difference as either “low”, “intermediate”, and/or “high.” The computer system can then leverage this display name difference to assess whether a spoofing attempt is possible and/or likely and thus whether to further investigate sender validity.

6.1 Similar Display Names

In one variation, the computer system can check whether the inbound display name included in the email approximates a verified display name included in the whitelist. In this variation, the computer system can leverage pattern matching techniques to extract differences between characters (e.g., combination, sequence, quantity) in the inbound display name and the verified display name. Based on these differences, the computer system can characterize a display name difference between the inbound display name and the verified display name and leverage this display name difference to extract insights into sender validity.

In one variation, the computer system can identify inbound display names closely resembling verified display names, such as an inbound display name corresponding to a different spelling of a verified display name (e.g., “Sarah” versus “Sara”). For example, the computer system can: intercept an email addressed to a target recipient within an organization, the email from a sender at an inbound email address and including an inbound display name of “John Troy”; access a whitelist including contact information corresponding to a first employee within the organization, the whitelist including a verified display name of “Jon Troy” and a verified email address corresponding to the first employee. The computer system can leverage pattern matching techniques to characterize a display name difference between “John Troy” and “Jon Troy” as a “low” display name difference indicating the inbound display name is similar to the verified display name of the first employee. Based on this “low” display name difference, the computer system can further investigate validity of the sender of the email by checking whether the inbound email address exactly matches the verified email address of the first employee. Additionally and/or alternatively, for a second email including an inbound display name of “Jessica Roy”, the computer system can again leverage pattern matching techniques to characterize a display name difference

between “Jessica Roy” and “Jon Troy” as a “high” display name difference indicating the inbound display name is distinct from the verified display name in the whitelist and therefore less likely to be a spoofing attempt of the first employee. Therefore, in this example, the computer system can authorize transmission of the second email without further investigation—with respect to the first employee. However, the computer system can similarly compare the inbound display name to other verified display names compared in the whitelist before authorizing transmission of the second email to a target recipient.

In one variation, the computer system can identify inbound display names corresponding to partial extractions of verified display names in the whitelist, such as an inbound display name corresponding to a set of initials extracted from a verified display name (e.g., “Jon Troy” versus “J Troy”, “Jon T”, and/or “JT”). For example, the computer system can: intercept an email from a sender at an inbound email address and including an inbound display name of “JT”; access a whitelist including contact information corresponding to a first employee and including a verified display name of “Jon Troy” and a verified email address corresponding to the first employee. The computer system can leverage pattern matching techniques to characterize a display name difference between “J Troy” and “JT” of twenty percent. Therefore, to characterize the display name difference, the computer system can identify whether the inbound display name corresponds to an abbreviated and/or shortened version of the verified display name included in the whitelist.

In one variation, the computer system can weight different characters within display names differently to characterize the display name difference. For example, in the preceding example, the computer system can: assign a weight of twenty percent to a first character of the forename; assign a weight of twenty percent to a first character of the surname; and distribute a weight of sixty percent evenly among each of the other characters in the verified display name. The computer system can then characterize the display name difference according to the assigned weights, such that the first characters of both the forename and the surname more heavily impact the display name difference than the other characters in the verified display name.

In one variation, the computer system can identify inbound display names corresponding to rearrangements of verified display names in the whitelist, such as an inbound display name corresponding to an inverse (or “reverse”) of a verified display name in the whitelist (e.g., “Jon Troy” versus “Troy Jon”). For example, the computer system can extract a forename and a surname from the inbound display name and compare this to a verified forename and a verified surname of the verified display name in the whitelist. If these do not match, the computer system can compare the forename of the inbound display name to the verified surname of the verified display name and the surname of the inbound display name to the verified forename of the verified display name. If these match, the computer system can identify overlap between the inbound display and the verified display name. Therefore, to characterize the display name difference, the computer system can identify whether the inbound display name corresponds to a rearrangement of the verified display name.

6.2 Nicknames

In one variation, the computer system can identify inbound display names corresponding to common name variations (or “nicknames”) associated with verified display names (e.g., “John” versus “Johnny” or “Jack”). The com-

puter system can leverage pattern matching techniques to identify nicknames corresponding to minor changes in display names, such as an inbound display name of “Johnny” compared to a verified display name of “John.” For nicknames corresponding to less obvious changes, the computer system can explicitly compare the inbound display name to a list of nicknames corresponding to the verified display name. For example, during onboarding, the computer system can prompt a user to input a set of nicknames for each employee in the whitelist. Later, the computer system can compare an inbound display name to both a verified display name and a set of nicknames assigned to the verified display name.

In another example, the computer system can access a name graph to identify a set of nicknames corresponding to a verified display name of an employee. More specifically, in response to intercepting an email received from a first sender at an inbound email address and including an inbound display name, the computer system can: access the whitelist including a verified display name and a set of verified email addresses corresponding to an employee within an organization; and characterize a display name difference between the inbound display name and the verified display name. To characterize the display name difference, the computer system can access a name graph including a corpus of names distributed about the graph based on correlations between names, such that a set of names proximal and/or surrounding a particular name are highly correlated to the particular name, thus possibly representing nicknames (or related names) of the particular name. In this example, the computer system can: access the name graph; search the name graph for the verified display name (e.g., a forename); and extract a set of names within a set radius of the verified display name within the name graph. The computer system can then compare the inbound display name to each name in the set of names extracted from the name graph. Therefore, the computer system can characterize the first display name difference based on whether the inbound display name matches a name in the set of names extracted from the name graph.

6.3 Symbols Imitating Authentic Characters

In one variation, the computer system can identify whether an inbound display name includes unauthentic characters excluded from verified display names (e.g., “John” versus “John”) and verified entity display names (“IT Department” versus “iT Department”). For example, the computer system can compare a set of characters corresponding to the inbound display name to a set of authentic characters corresponding to a verified display name. In response to identifying a first number in replacement of a first letter (e.g., a “0” in replacement of an “o”) in the inbound display name, the computer system can characterize the display name difference based on this replacement. Alternatively, in one variation, in response to identifying symbols (e.g., numbers) within the inbound display name in replacement of authentic characters (e.g., letters) within the verified display name, such as—“1T Department”, “Accounting T3am”, etc.—the computer system can automatically withhold transmission of the email and flag the email for authentication.

6.4 Base64 Encoding

In one variation, the computer system can identify different encoding methods within the inbound display name. The computer system can characterize the display name difference based on a type of encoding identified. For example, the computer system can identify an inbound display name encoded in Base64 and automatically flag this

email for authentication by the email administrator. Alternatively, the computer system can translate the inbound display name accordingly and compare this inbound display name to a verified display name in the whitelist.

In one variation, the computer system can implement Base64 decoding and Base32 decoding to compare the inbound display name to the verified display name. For example, the computer system can first translate an inbound display name according to Base64 decoding. If, based on the translation, the display name difference exceeds the threshold display name difference and an inbound email address does not match a verified email address corresponding to the verified display name, the computer system can withhold transmission of the email and flag the email for authentication. If, however, the display name difference falls below the threshold display name difference, the computer system can translate the inbound display name according to Base32 decoding and compare this translated inbound display name to the verified display name accordingly.

7. Email Quarantine

In one variation, the computer system can deliver an email notification to the email administrator including a hyperlink that, when selected by the email administrator, automatically opens a web browser with access to a web portal and the quarantined email for investigation. The email administrator may investigate the quarantined email and determine whether the sender is legitimate. Upon receiving verification of the sender by the email administrator via the web portal, the computer system can deliver the email to a designated recipient. Alternatively, if the email administrator determines the sender is not authentic, the computer system can withhold the email from the designated recipient.

In one variation, the computer system can deliver an email notification to the email administrator including: the original quarantined email (e.g., the text of the quarantined email in the content of the notification email or via a hyperlink that opens a web browser that downloads the quarantined email); a hyperlink that, when selected, triggers the computer system to deliver the quarantined email to a designated recipient; and a hyperlink that, when selected, triggers the computer system to discard the quarantined email. Additionally, the notification email can include a hyperlink that, upon selection, triggers the computer system to discard the quarantined email and add a sender of the quarantined email to the blacklist.

Alternatively, the computer system can notify the email administrator of a quarantined email and enable the email administrator to deliver or withhold the email. For example, the computer system can: notify the email administrator of a new quarantined email; prompt the email administrator to investigate the validity of the email and sender; prompt the email administrator to either forward the email to a designated recipient or discard the email.

In one variation, a particular email address may send out multiple emails to multiple recipients within an organization. In this variation, the computer system can combine these emails into one notification to the email administrator. For example, in response to receiving multiple emails from a particular sender with a display name found in the whitelist but from an email address not contained in the whitelist, the computer system can: flag each email from this sender for quarantine; merge these emails into a single email notification and deliver the email notification to the email administrator; and receive verification or denial of these emails or

a subset of these emails from the email administrator and distribute these emails or withhold these emails accordingly.

7.1 Quarantine Portal

In one variation, as shown in FIG. 4, the computer system can withhold flagged emails for further investigation of sender validity within an online portal (or “quarantine portal”) accessible by the email administrator. The email administrator may access an instance of the quarantine portal (e.g., via a native application operating on her mobile phone, at a webpage operating on her laptop computer) to view, sort, and/or verify authenticity of emails flagged by the computer system.

Upon flagging an email for authentication, the computer system can automatically add the email to a quarantined email list viewable to the email administrator within the quarantine portal. The email administrator may access the quarantine portal to view the updated quarantined email list and select the email to view an inbound email address and an inbound display name associated with the email. The email administrator may then investigate authenticity and, upon determination of an authentic sender, transmit authentication of the email to the computer system (e.g., via selection of a corresponding “authenticate” hyperlink). Alternatively, upon determination of an unauthentic sender (e.g., a spoofing attempt), the email administrator may transmit confirmation of a spoofing attempt to the computer system (e.g., via selection of a corresponding “spoof attempt” hyperlink). In response to receiving authentication of the email from the email administrator, the computer system can authorize transmission of the email to a target recipient designated in the email. Alternatively, in response to receiving confirmation of an unauthentic sender, the computer system can withhold transmission of the email to the target recipient and/or discard the email.

In one variation, the computer system can deliver an email notification to the email administrator including a hyperlink that, when selected by the email administrator, automatically opens a web browser with access to a web portal containing an email quarantine database. Within the web portal, the email administrator may view all emails flagged for authentication by the computer system. The computer system can then receive inputs from the email administrator via the web portal indicating whether to authorize transmission of quarantined emails to target recipients or discard these emails.

For example, the computer system can serve an email notification including a hyperlink to an email administrator for a new quarantined email. Upon selection of the hyperlink by the email administrator, the computer system can enable viewing of the quarantine portal at which the email administrator may view the quarantined email and select whether to deliver or discard the quarantined email. The computer system can receive this selection and act accordingly. Therefore, the computer system can automatically inform the email administrator of a new quarantined email available within the quarantine portal, thus enabling the email administrator to quickly and more easily access the quarantine portal when a new email is added to the quarantine list, such that emails from authentic (e.g., legitimate) senders may be quickly authenticated and delivered to their designated target recipients.

In one variation, the email administrator may sort emails within the quarantined email list by sender, email address, entity name, order received (e.g., timestamp), and/or reasons for flagging in order to prioritize investigation of particular emails and/or more efficiently investigate these emails. For example, the email administrator may sort the quarantined email list by inbound email address and select whether to

deliver or discard a set of emails, within the quarantined email list, sent from a particular inbound email address. More specifically, the email administrator may determine that each email, in the set of emails, sent from this email address includes a different display name. Thus, the email administrator may determine that this particular inbound email address corresponds to an unauthentic sender. Additionally, the email administrator may input a selection indicating addition of this particular email address to the blacklist.

In one variation, the computer system can automatically sort emails within the list of quarantined emails when adding a new email to this list. In this variation, the computer system may sort emails within the list of quarantined emails according to priority (e.g., sender priority, target recipient priority, order received). For example, in response to flagging an email for authentication, the computer system can access the whitelist to extract a rank (e.g., Rank I, Rank II, and Rank III) of an employee associated with the inbound email address and/or display name of the email, the rank corresponding to a position of the employee within the organization. In response to the whitelist specifying a high rank (e.g., corresponding to the CEO of the organization), the computer system can automatically insert the email at a top of the quarantined email list to enable prompt investigation of this email.

In one variation, the email administrator can input a selection within the quarantine portal indicating that the email administrator may be unsure of whether an inbound email address corresponds to an authentic sender or an unauthentic sender. Upon receiving this selection, the computer system can generate a notification requesting confirmation of validity of the email and transmit this notification to an employee (e.g., at a verified email address) or an individual within an entity associated the email from which the email was allegedly sent. Upon receiving confirmation of validity of the email from the employee or individual associated with the entity, the computer system can automatically authorize transmission of the email to a target recipient. Alternatively, upon receiving confirmation of an unauthentic sender from the employee, the computer system can automatically discard the email.

8. Authenticated Email

Upon receiving verification of an email—initially flagged for quarantine (not found in the whitelist)—from the email administrator via the web portal, the computer system can deliver the email to the target recipient. Alternatively, the computer system can notify the email administrator of the email flagged for quarantine, and the email administrator may manually forward the email to a target recipient upon verification of the sender or withhold the email if the sender is not verified.

In one variation, the computer system can update the whitelist to include a verified email address for an employee or entity associated with this email address. For example, the computer system can: receive verification of an email from an email address initially flagged for quarantine from the email administrator via the web portal; deliver the email to a designated recipient; and add the email address to the whitelist. Therefore, at a future time, when an employee or individual within the entity associated with this email address sends an email, the computer system will find the email address in the whitelist and not initiate further investigation of validity, thus decreasing latency between sending emails from this email address and receiving these emails,

decreasing a workload of the email administrator, and eliminating the need to manually upload new contact information.

In one variation, the computer system can include a verified notification to the recipient in the email to communicate to the recipient that the email is from a verified sender. For example, the computer system can: receive verification of the email from the email administrator via the web portal, add a tag (e.g., a notification) in the email indicating the email has been verified and the sender is legitimate, and deliver the email to a designated recipient. Therefore, the computer system can increase confidence of the recipient that the sender and the contents contained in the email are legitimate. Thus, the computer system can leverage the ability to verify the identity of email senders to increase trust and confidence of both senders and recipients of emails, and therefore enable employees to engage with or act on contents contained in emails more efficiently.

8.1 Authentic Email Notification

In one variation, upon receiving verification of the email sender from the email administrator, the computer system can generate a notification for transmission to the target recipient of the email indicating authentication of the email sender. For example, the computer system can: intercept an email addressed to a target recipient, the email from a first sender at an inbound email address and including an inbound display name. Upon inspection of the inbound display name and the inbound email address, the computer system can withhold transmission of the email to the target recipient and flag the email for authentication by the email administrator. Later, upon receiving authentication of the email from the email administrator, the computer system can: generate a notification indicating authentication by the email administrator of the email sender; insert the notification into a bottom portion of the email; and authorize transmission of the email to the target recipient. Therefore, the computer system can enable email recipients to feel confident that email senders are legitimate and thus respond accordingly and confidently to emails from verified senders.

9. Sender Error: Invalid Email/Spoofing Attempt

The computer system may receive confirmation from the email administrator via the web portal that an email from a particular email address is not verified, invalid, or a spoofing attempt. Upon receiving this confirmation, the computer system can withhold the email from its designated recipient and instead discard the email.

In one variation, the computer system can generate a notification detailing this spoof attempt for delivery to the target recipient of the discarded email. Additionally and/or alternatively, the computer system can generate a notification detailing this spoof attempt for delivery to an employee associated with the verified display name copied or imitated in the spoofing attempt by the email sender.

9.1 Blacklist

In one variation, the computer system can additionally add an email address to a database containing email addresses of phony senders in near real-time upon receiving confirmation from the email administrator that the email address is invalid, not verified, or a spoofing attempt. Later, the computer system can access this list when checking other incoming emails. For example, the computer system can: at a first time, receive a first email from a first sender addressed to a first recipient, the first email containing a display name and sent via a first email address with a first display name and at a first domain; access a whitelist containing a set of verified email addresses; in response to

the whitelist containing the display name, comparing the first email address to a verified email address in the set of verified email addresses associated with the display name; in response to the verified email address not matching the first email address, quarantine and flag the first email for further investigation; and notify the email administrator. In response to receiving a denial of the first email from the email administrator via a web portal, the computer system can record the first email address to a blacklist representing blocked email addresses associated with phishing schemes; at a second time, receive a second email sent via the first email address; access the blacklist to search for the first email address; and, in response to the blacklist containing the first email address, automatically quarantine the second email and withhold the second email from a designated recipient. Therefore, by populating a blacklist and searching for email addresses in the blacklist before checking the whitelist, the computer system can: minimize notifications sent to the email administrator; decrease workload of the email administrator by eliminating notifications from previously denied senders (e.g., known phishing emails); minimize latency between the sending of an authentic email by an unverified email address and delivery to the recipient.

The computer system can generate a unique blacklist for each organization (e.g., company) based on denied emails sent to employees of the organization. Alternatively, the computer system can access a global blacklist containing denied email addresses from multiple organizations to further reduce latency between sending of legitimate emails and receiving these emails. For example, the computer system can receive confirmation from the email administrator via the web portal that an email from a particular email address is malicious (e.g., a spoofing attempt). Upon receiving this confirmation, the computer system can: withhold the malicious email from a designated recipient; discard the malicious email; and add the sender email address from this malicious email to a blacklist containing email addresses associated with malicious spoofing emails. For example, the computer system can populate this blacklist with email address associated with malicious spoofing emails inbound to multiple (e.g., many, thousands of) organizations and automatically discard inbound emails from any sender email address on this blacklist to recipients within all of these organizations. Therefore, at a future time, in response to receiving a new inbound email, the computer system can check the blacklist for an email address that matches the sender email address of this new inbound email and then automatically discard this new inbound email responsive to detecting such as a match.

In one variation, the computer system can populate the blacklist by adding email addresses that have sent multiple emails exhibiting different display names. For example, the computer system can: at a first time, receive a first email with a first display name from a first email address; access the blacklist for the first email address; in response to the blacklist omitting the first email address, access the whitelist to compare the first display name in the first email with verified employee names contained in the whitelist; in response to the whitelist containing a first verified employee name corresponding to the first display name in the first email, compare the first email address with a set of verified email addresses corresponding to the first verified employee name; in response to the first email address not matching a verified email address contained in the whitelist, quarantine the first email; notify the email administrator; receive a denial of the first email from the email administrator via a web portal; at a second time, receive a second email with a

second display name from the first email address; in response to the second display name not matching the first display name, notify the email administrator; and add the first email address to the blacklist.

10. Employee Ranking

In one implementation, the computer system can screen emails differently depending on the type or level of the employee (e.g., CEO, manager, entry-level employee) with which the display name in an email is associated. For example, an organization may collect contact information for all employees and upload this contact information to the computer system via a web portal and specify employee rankings within the organization. The computer system can: at a first time, receive the whitelist for the organization containing contact information of employees at the company; assign each employee a threat level (e.g., low or high threat level, Level I, II, or III) based on employee rankings at the organization as specified in the whitelist; at a second time, receive an email containing a first display name and addressed to a first recipient; access the whitelist of the organization to check for the first display name; in response to the first display name corresponding to the CEO of the company, the CEO previously assigned a high threat level, quarantine the email and withhold the email from the first recipient; flag the email for further investigation and notify the email administrator. If an email administrator determines that the email is legitimate, the computer system can then receive verification of the email from the email administrator and deliver the email to the recipient.

Alternatively, upon receiving an email with a display name corresponding to a lower level employee within the company and from a first email address, the computer system can: access the whitelist of the organization to check for the display name; in response to the display name corresponding to an entry level employee of the company, the entry level employee previously assigned a low threat level, check the whitelist to compare the first email address to a verified email address of the entry level employee; in response to the first email address not matching the verified email address of the entry level employee, flag the email as an unverified email from an unverified sender; and deliver the email to a designated recipient including an unverified sender notification.

Therefore, in this variation, the computer system can verify email senders for all employees within an organization while allocating additional resources to higher ranking employees (e.g., CEO, COO, president) and minimizing resources allocated to lower ranking employees. For example, an organization may define three employee levels (e.g., Level I, II, and III) and assign employees to these levels based on employee position with the organization. The computer system can then implement different email sender review methods based on the level of the employee associated with the display name of an email. Thus, the computer system can: upon receiving an email from a first sender with a display name corresponding to a Level I employee, automatically quarantine the email and notify the email administrator; upon receiving an email from a first sender with a display name corresponding to a Level II employee, check the whitelist for the email address associated with the email, and quarantine the email and notify the email administrator if the email address is not found on the whitelist; upon receiving an email from a first sender with a display name corresponding to a Level III employee: check the whitelist for the email address associated with the email,

flag the email as sent from an unverified sender if the email address is not found on the whitelist, and deliver the email to a designated recipient including a notification of the unverified sender. In this variation, the computer system enables organizations to set minimum thresholds for which investigation of email senders is required.

11. Common Names

In one variation, the computer system can verify and quarantine emails from senders with exact or similar names as employees included in the whitelist and, upon further investigation and approval by the email administrator, add the contact information for these senders to the whitelist. Upon serving a notification of a quarantined email to the email administrator, the computer system can receive a secondary verification of the email address indicating a verified contact other than the verified employee contained in the whitelist.

For example, a vendor for the organization may have a name similar or identical to the name of a keyman (e.g., the CEO) within the organization. In this example, the computer system can: receive an email from a sender with a display name “James Smith” and with a first email address including a first username “J.Smith” and a first domain; access a whitelist of an organization; find a keyman profile with a name “James Smith” in the whitelist and with a second email address including a second username “JSmith” and a second domain; and quarantine the email and notify the email administrator. In response to receiving a secondary verification of the email from the email administrator as from a vendor of the organization and not the CEO, the computer system can: add the vendor to the whitelist; and deliver the email to its designated recipient. Therefore, the computer system can verify and deliver emails from legitimate senders with similar or exact names as employees in the whitelist, and incorporate their contact information into the whitelist to prevent the delay or quarantine of emails from these senders.

12. Whitelist Update

In one variation, the computer system can update the whitelist over time in response to verification—by the email administrator or other security analyst—of flagged or quarantined emails sent from email addresses not previously contained in keyman profiles in the whitelist.

For example, the computer system can: quarantine an inbound email containing a display name contained in a particular keyman profile in the whitelist but sent from an email address not contained in this particular keyman profile; and notify the email administrator of this quarantined email, such as in an email quarantine database or security alert feed. Upon verification of the sender and release of the first email to a designated recipient by the email administrator, the computer system can prompt the email administrator to indicate whether the sender email address from this released email corresponds to a new email address for the particular keyman; if so, the computer system can add the sender email address to the particular keyman profile associated with this particular keyman in the whitelist.

Later, upon receiving a second inbound email containing this same display name and identifying this new email address, the computer system can verify this second inbound email based on the updated keyman profile of this particular

keyman and deliver this second inbound email to its designated recipient without prompting further investigation by the email administrator.

Therefore, during a setup period, the computer system can aggregate and indicate a set of contact information for a first select group of employees (e.g., leaders and managers within an organization), generate initial keyman profiles for these employees, and generate a whitelist containing these keyman profiles. The computer system can then intelligently add new contact information to these keyman profiles as the new contact information for these keymen are verified—ad hoc—by the email administrator or other security analyst over time.

The computer system can also add sender email addresses—verified by the email administrator or security analyst but not associated with a keyman in the organization—to the whitelist. For example, the computer system can: receive an email from an email address with a domain external to the organization; access a whitelist for this organization; and implement methods and techniques described above to check the sender email address and display name for this inbound email. If the computer system determines that the display name in this inbound email matches a display name in a keyman profile in the whitelist but the sender email address does not match a verified email address in a corresponding keyman profile in the whitelist, the computer system can flag the email for further investigation and notify the email administrator of the email. If the email administrator then verifies that the email is valid but not sent by a keyman in the organization, the computer system can generate a new verified sender profile containing this display name and sender email address and add this new verified sender profile to the whitelist.

13. Variation: Exact Email Address

In one variation, as shown in FIG. 2, the computer system can initially compare the inbound email address associated with an email to a set of verified email addresses included in the whitelist, prior to searching for the inbound display name. Upon finding an exact match of the inbound email address in the set of verified email addresses, the computer system can automatically authorize transmission of the email to a target recipient. Alternatively, the computer system can still compare the inbound display name to the verified display name in the whitelist as an additional check for authentication.

In one variation, the computer system can prompt a user to update a verified display name corresponding to the user in response to intercepting an email from the user at a verified email address but including an alternate display name. For example, the computer system can: intercept an email received from a first sender at a first inbound email address and including a first inbound display name; access a whitelist including a first set of contact information corresponding to a first employee within an organization, the set of contact information including a verified display name associated with the first employee and a set of verified email addresses associated with the first employee; identify the first inbound email address within the set of verified email addresses associated with the first employee; and, in response to identifying a difference between the first inbound display name and the first verified display name, generate a prompt for transmission to the first employee, the prompt requesting confirmation of the inbound display name

as a new verified display name of the first employee and/or presenting instructions to update the first verified display name.

14. Variation: Email Verification for an Organization

As shown in FIGS. 5A-5C, 6A-6C, 7A, and 7B, one variation of the method S100 includes, during a first time period: intercepting a first email addressed to a first target recipient within an organization, the first email received from a first sender at a first inbound email address and including a first inbound display name in Block S110; and accessing a whitelist including a first set of contact information corresponding to a first entity associated with the organization, the first set of contact information including a first verified display name associated with the first entity and a first set of verified email addresses associated with the first entity in Block S120. The method S100 further includes, during the first time period, in response to the first set of verified email addresses omitting the first inbound email address: characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first entity in Block S130; and, in response to the first display name difference falling below a threshold difference, withholding transmission of the first email to the first target recipient and flagging the first email for authentication in Block S170. In this variation, the method S100 further includes, during a second time period: intercepting a second email addressed to a second target recipient within the organization, the second email received from a second sender at a second inbound email address and including a second inbound display name in Block S110; and, in response to the first set of verified email addresses including the second inbound email address, authorizing transmission of the second email to the second target recipient in Block S160.

One variation of the method S100 includes, during a first time period: intercepting a first email addressed to a first target recipient within an organization, the first email received from a first sender at a first inbound email address and including a first inbound display name in Block S110; accessing a whitelist including a first set of contact information corresponding to a first entity associated with the organization, the first set of contact information including a first verified display name associated with the first entity and a first set of verified email addresses associated with the first entity in Block S120; characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first entity in Block S130; and, in response to the first display name difference falling below a threshold difference and in response to the first set of verified email addresses omitting the first inbound email address, withholding transmission of the first email to the first target recipient and flagging the first email for authentication in Block S170. In this variation, the method S100 further includes, during a second time period: intercepting a second email addressed to a second target recipient within the organization, the second email received from a second sender at a second inbound email address and including a second inbound display name in Block S110; characterizing a second display name difference between the second inbound display name and the first verified display name associated with the first entity in Block S130; and, in response to the second display name difference falling below the threshold difference and in response to the first set of verified email addresses including

the second inbound email address, authorizing transmission of the second email to the second target recipient in Block S160.

One variation of the method S100 includes, during a first time period: intercepting a first email addressed to a first target recipient within an organization, the first email including a first inbound display name and received from a first sender at a first inbound email address including a first domain in Block S110; and accessing a whitelist including a first set of verified domains—associated with inbound emails received by recipients within the organization—and a first set of contact information corresponding to a first entity associated with the organization, the first set of contact information including a first verified display name associated with the first entity and a first set of verified email addresses associated with the first entity in Block S120. The method S100 further includes, during the first time period, in response to the set of verified domains omitting the first domain and in response to the first set of verified email addresses omitting the first inbound email address: characterizing a first display name difference between the first inbound display name and the first verified display name in Block S130; and, in response to the first display name difference falling below a threshold difference, withholding transmission of the first email to the first target recipient and flagging the first email for authentication in Block S170. In this variation, the method S100 further includes, during a second time period: intercepting a second email addressed to a second target recipient within the organization, the second email received from a second sender at a second inbound email address including a second domain and including a second inbound display name; and, in response to the set of verified domains including the second domain, authorizing transmission of the second email to the second target recipient in Block S160.

One variation of the method S100 includes: during a first time, intercepting a first email addressed to a first target recipient within an organization, the first email received from a first sender at a first inbound email address and including a first inbound display name in Block S110; accessing a whitelist including a first set of contact information corresponding to a first entity within the organization in Block S120, the first set of contact information including a first verified display name associated with the first entity, and a first set of verified email addresses associated with the first entity; and characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first entity in Block S130. In this variation, the method S100 further includes: in response to the first inbound display name difference falling below a threshold difference, comparing the first inbound email address to the first set of verified email addresses associated with the first entity in Block S140; in response to the first set of verified email addresses omitting the first inbound email address, withholding transmission of the first email to the first target recipient; and flagging the first email for authentication in Block S170. This variation of method S100 further includes: during a second time, intercepting a second email addressed to the first target recipient within the organization, the second email received from the first sender at a second inbound email address and including a second inbound display name in Block S115; accessing the whitelist including the first set of contact information corresponding to the first entity within the organization in Block S120; characterizing a second display name difference between the second inbound display name and the second verified display name associated with the first entity in Block S130; in

response to the second display name difference falling below the threshold difference, comparing the second inbound email address to the first set of verified email addresses associated with the first entity in Block S140; and, in response to identifying the second inbound email address in the first set of verified email addresses, authorizing transmission of the second email to the first target recipient in Block S160.

14.1 Variation: Email Verification for an Organization

Generally, in this variation, the computer system can verify an identity of an email sender—associated with and/or imitating a particular entity (e.g., a department, a team, a particular group of employees) within an organization (e.g., a company, a business)—before passing an email to its designated recipient (e.g., associated with the organization) in order to detect and suppress email spoofing attempts.

For example, a spoof or spammer may: leverage or gain access to an organization directory to identify various entities (e.g., departments, employee groups, divisions) within an organization and/or select generic department or group titles — such as “HR,” “finance,” “engineering,” “billing,” “admin”, and/or “IT” — commonly associated with organizations; and create a false email account with a display name identical and/or related to a particular entity (e.g., “Finance” or “Billing” or “Human Resources”). In this example, the spoof or spammer may deliver an email to an employee, group of employees, or associate of the organization—from this false email account—with an urgent request, such as an email requesting immediate payment for an outstanding invoice sent from a false email account with a display name of “Billing Department” or an email requesting login credentials sent from a false email account with a display name “Human Resources.”

In this variation, the email administrator within the organization identifies an entity (e.g., a department, a team) within the organization that exhibits greater risk of email spoofing. For example, the email administrator may define an entity, such as—an accounting department, an HR department, a healthcare unit (e.g., Radiology, Family Practice), and/or a technology service group (e.g., IT department), etc. The computer system can then verify that inbound emails that include display names that match (or are otherwise similar to) verified display names and/or names (e.g., titles) of these entities are inbound from verified email addresses associated with these entities—and corresponding verified display names—before releasing these emails to their designated recipients.

14.2 Example

In one example, the computer system can receive contact information—such as including email addresses, display names, phone numbers, etc.—for each entity (e.g., a department, a division, a branch, a group of employees, a particular employee) associated with an organization (e.g., a company, a business). The computer system can then store this information in a contact profile generated for the particular entity.

For example, the computer system can generate: a first contact profile including verified contact information associated with a finance department of an organization; a second contact profile including verified contact information associated with an HR department of the organization; a third contact profile including verified contact information associated with an engineering department of the organization; a fourth contact profile including verified contact information associated with an executive department (e.g., or “C-suite”) of the organization; etc. The computer system

can then compile the first, second, and third contact profiles into a whitelist for the organization.

In particular, in this example, the first contact profile—associated with the finance department—can include: a set of verified email addresses, such as “FinanceDep@companyXYZ.com”, “FinanceDepartment@companyXYZ.com”, and/or “Finance@companyXYZ.com”; a verified display name—such as “Finance”, “Finance Department”, “XYZ Finance Division”, etc.—corresponding to each verified email address in the set of verified email addresses corresponding to the finance department.

Later, upon receiving a first email from a particular sender (e.g., a potential spoof)—sent from an inbound email address and specifying an inbound display name—addressed to a target recipient (e.g., an employee) associated with the organization, the computer system can access a whitelist including a set of contact information corresponding to the finance department. In this example, the set of contact information can include the verified display name, “Finance Department,” and the set of verified email addresses associated with the finance department, “FinanceDep@companyXYZ.com”, and “FinanceDepartment@companyXYZ.com.”

The computer system can then characterize a display name difference between a first inbound display name (e.g., “Financial Department”) and the verified display name associated with the finance department. In this example, in response to the first display name difference falling below a threshold difference, the computer system can compare the first inbound email address, such as “financialdepartment@companyXYZ.com” to the set of verified email addresses associated with the first entity, “FinanceDep@companyXYZ.com”, and “FinanceDepartment@companyXYZ.com”. Then, in response to the set of verified email addresses omitting the first inbound email address, the computer system can withhold transmission of the first email to the target recipient and/or flag the first email for authentication by an administrator overseeing the contact profiles.

The computer system can later: receive a second email from a second sender at a second inbound email address of “FinanceDepartment@companyXYZ.com” and including a second inbound display name of “Finance Department”; and characterize a second display name difference between the second inbound display name and the first verified display name of “Finance Department” associated with the finance department. In response to the second display name difference falling below the threshold difference, the computer system can: compare the second inbound email address to the set of verified email addresses associated with the finance department; and, in response to identifying the second inbound email address in the set of verified email addresses, authorize transmission of the second email to the target recipient.

14.3 Onboarding

Generally, the computer system interfaces with an email administrator of the organization to obtain contact information for each entity within the organization and/or a representative (e.g., a manager) of each entity. For example, the email administrator and/or entity representative may upload contact information—such as including email addresses, display names, phone numbers, etc.—for each entity associated with an organization (e.g., a company, a business). The computer system can then store this information in a contact profile (e.g., a whitelist) generated for the particular entity. In particular, for an HR Department, the computer

system store a contact profile including: a set of verified display names, such as including “HR Department” and/or “HR Administrator”; and a set of verified email addresses, such as including “HR@companyXYZ.com,” “humanresources@companyXYZ.com,” and/or “HRadmin@companyXYZ.com.” Therefore, at a future time, the computer system can verify an inbound email including an inbound display name (e.g., “HR,” human resources,” or “HRadmin”) similar or identical to the set of verified display names defined for the HR department based on the contact profile.

In one implementation, upon receiving an email including a display name and from a sender at a particular sender address, the computer system can: access the whitelist containing entity profiles; and search for an entity profile including a set of display names that match or are identical to the display name based on the inbound email. Then, upon identifying that the HR department’s profile contains a display name (e.g., “HR Department”) that matches the display name of the inbound email, the computer system can compare the sender email address of the inbound email to a set of verified email addresses contained in the HR department’s profile. If the sender email address matches a verified email address in the HR department’s profile, the computer system can verify the sender as an individual in the HR department and deliver the email to its designated recipient. Otherwise, the computer system can flag and quarantine the email for further investigation.

14.3.1 Variation: External Entities

In one variation, the computer system can interface with the email administrator of the organization to obtain contact information for a set of entities associated with and external the organization, such as including customers, vendors, clients, etc., of the organization. Generally, in this variation, the computer system interfaces with the email administrator (e.g., via an administrator portal) to: obtain contact information for each external entity associated with the organization; generate a contact profile (e.g., as described above) for each external entity; and store the contact profile in the whitelist generated for the organization.

Furthermore, in one variation, the computer system can automatically update the whitelist to include contact information associated with external entities (e.g., external the organization) based on email engagement of these entities with internal entities (e.g., internal the organization), such as including employees and/or departments within the organization. For example, the computer system can: intercept an email sent by a sender—to a target recipient within the organization—at an inbound email address of “billing@company.abc” and including a display name of “ABC Billing Department”; scan an email database for historical emails sent from the inbound email address; derive a quantity of inbound emails sent from the inbound email address during a particular time period; derive a quantity of outbound emails sent to the inbound email address during the particular time period; and, based on the quantity of inbound emails and the quantity of outbound emails, characterize engagement—such as represented as a qualitative (e.g., high, moderate, low) or quantitative (e.g., 10%, 50%, 95%) engagement score—of an external entity associated with the inbound email address with the organization. Then, in response to characterizing engagement as relatively high, the computer system can: retrieve the whitelist for the organization; generate a new contact profile—associated with the external entity—including the inbound email address and the inbound display name; and store the new contact profile in the whitelist. Alternatively, in response to

characterizing engagement as relatively low, the computer system can: omit the inbound email address from the whitelist and continue to verify future emails sent from this inbound email address.

5 Additionally or alternatively, the computer system can: generate a notification including the new contact profile and a prompt to review the new contact profile; and, in response to receiving verification of the new contact profile from the email administrator, update the whitelist to include the new 10 contact profile. However, in response to receiving rejection of the new contact profile from the email administrator, the computer system can discard the new contact profile and, therefore, continue to verify emails sent from this inbound 15 email address.

14.3.2 Verified Domains

The computer system can thus whitelist all email addresses of entities (e.g., an employee, a department) within an organization domain (e.g., Company ABCD with an email domain “@ABCD.com”). Therefore, the computer 20 system can automatically authorize transmission of emails sent from email addresses within the organization domain without further checks for authenticity according to the method, thereby reducing overhead and computational power scanning these internal emails for spoofing attempts.

25 Furthermore, in one variation, in which the computer system verifies inbound emails from senders external the organization, the computer system can similarly add all email addresses within a particular domain—such as associated with an external customer or vendor of the organization—to the whitelist. For example, the computer system can access a whitelist including: a set of verified domains (e.g., email domains) associated with the organization; a first 30 set of contact information corresponding to a first entity—such as an internal department (e.g., human resources, finance, billing, administrative, engineering)—within the organization—and including one or more verified display names (e.g., associated with the first entity) and one or more verified email addresses including a first verified domain (e.g., “@company.xyz”, “@hr.company.xyz”, “@eng.com- 35

pany.xyz”) in the set of verified domains; and a second set 40 of contact information corresponding to a second entity—such as an external customer or vendor—associated with the organization and including one or more verified display names (e.g., associated with the second entity) and one or 45 more verified email addresses including a second verified domain (e.g., “@company.abc”, “@sales.company.abc”, “@billing.company.abc”) in the set of verified domains. In this example, the computer system—such as in response to receiving confirmation from the email administrator—can therefore automatically authorize transmission of emails sent from both internal entities and external entities (e.g., external the administration) sent from email addresses including verified domains in the whitelist.

14.4 Inbound Email Check

55 Upon generation of the whitelist, the computer system can receive (or “intercept”) inbound emails from senders—associated with and/or spoofing a particular entity within the organization—and verify the validity of these senders based on contact information stored in the whitelist (e.g., email addresses, display names).

In particular, the computer system can: receive an email from a sender email address and addressed to a recipient within the organization; extract a display name and a sender email address from the email; compare the display name (e.g., “HR”, “Finance Department”) to verified display 60 names—included in a whitelist generated for the organization—of various entities within the organization; and, in

response identifying a verified display name in the whitelist that is similar to or identical the display name for the inbound email, compare the sender email address to a set of verified email addresses (e.g., one or more verified email addresses)—corresponding to the verified display name—in the whitelist. The computer system can then selectively transmit and/or withhold transmission of the email to the recipient based on omission and/or inclusion of the sender email address in the set of verified email addresses in the whitelist.

In one implementation, upon receiving an email including a display name and designating a target recipient from a sender at a sender email address, the computer system can: access a whitelist including contact information (e.g., email addresses, display names) of entities associated with the organization; in response to the whitelist including a verified display name of an entity that matches the display name included in the email, compare a set of verified email addresses specified for the entity in the whitelist to the sender email address; in response to the whitelist including the sender email address, deliver the email to the target recipient; and, in response to the whitelist omitting the sender email address, quarantine the email in a quarantine database and notify an email administrator of the email for further investigation. Therefore, if the computer system detects that the inbound display name matches a verified display name (e.g., in the whitelist), the computer system can verify whether the inbound email address matches a verified email address corresponding to the verified display name.

Then, in the preceding implementation, if the computer system identifies that the inbound display name matches (and/or is similar to) a verified display name and the inbound email address matches the verified email address specified for this verified display name, the computer system can withhold additional authentication by the administrator and send the email to the recipient. However, if the computer system identifies that the inbound display name matches (and/or is similar to) a verified display name—and the inbound email address fails to match the verified email address—the computer system can quarantine the email for authentication by an administrator prior to delivering the email to the recipient.

Alternatively, in the preceding implementation, if the computer system detects that the inbound display name fails to match a verified display name on the whitelist associated with the entity, the computer system can predict that the email is not a spoofing attempt and therefore automatically release the email for transmittal to the recipient.

14. Sender Verification: Email Domain

In one variation, the computer system can verify whether a domain of the inbound email address corresponds to a verified domain defined by the whitelist.

In this variation, the computer system can verify sender identity by comparing an inbound domain name (e.g., “@companyXYZ”)—extracted from the inbound email address from which the email is sent—to a set of verified domains (e.g., email domains) associated with the organization and/or a particular entity within the organization. For example, upon receiving an email from a sender at an inbound email address, the computer system can: extract a first domain from the first email address; access the whitelist containing a set of verified domains associated with verified email senders within and/or associated with the organization; and, in response to the first domain name matching a

verified domain, in the set of verified domains, automatically authorize transmission of the email to a target recipient and/or release the email to the target recipient.

Alternatively, in response to the first domain differing from each verified domain, in the set of verified domains, the computer system can: compare an inbound display name included in the email to a set of verified display names defined for an entity—associated with the organization—within the whitelist; and, in response to the inbound display name corresponding to (e.g., matching, approximating, mimicking) a first verified display name, in the set of verified display names, withhold transmission of the email to the target recipient and flag the email for further investigation and/or authentication by an administrator.

The computer system can thus whitelist all email addresses of entities (e.g., an employee, a department) within an organization domain (e.g., Company ABCD with an email domain “@ABCD.com”). Therefore, the computer system can automatically authorize transmission of emails sent from email addresses within the organization domain without further checks for authenticity according to the method, thereby reducing overhead and computational power scanning these internal emails for spoofing attempts.

14.6 Sender Verification: Display Name

Generally, in response to receiving an email, the computer system can check a display name associated with the email. In particular, the computer system can characterize a difference between the display name and a verified display name associated with an entity in the whitelist. In response to the difference falling below a threshold difference—such that the display name matches, approximates, imitates, and/or corresponds to the verified display name—the computer system can withhold transmission of the email to the target recipient and/or implement further security checks to verify authenticity of the sender.

14.6.1 Verified Display Name Match

In one implementation, the computer system can verify whether the inbound display name matches a verified display name, in a set of verified display names, associated with an entity defined in the whitelist. In particular, in this implementation, the computer system can: intercept a first email addressed to a first target recipient within an organization, the first email received from a first sender at a first inbound email address and including a first inbound display name; access the whitelist including a set of contact information—including a verified display name and a set of verified email addresses—corresponding to an entity within the organization; compare the inbound display name to a first verified display name in the set of verified display names. Then, in response to the inbound display name matching the first verified display name, the computer system can selectively withhold sending of the first email to the first target recipient and/or verify whether the inbound email address matches a verified email address in the set of verified email addresses.

Alternatively, in the preceding example, in response to the inbound display name differing from the first verified display name, the computer system can repeat this process—such as for a second, third, and/or fourth verified display name in the set of verified display names—to compare the inbound display name to each verified display name in the set of verified display names or until identifying a match between the inbound display name and a corresponding verified display name in the set of verified display names.

Alternatively, in response to the inbound display name differing from each verified display name, in the set of verified display names, the computer system can: authorize

transmission of the first email to the first target recipient; and/or further verify authenticity of the sender of the inbound email, such as based on similarities between the inbound display name and the verified display name, as described below.

14.6.2 Word Similarity

In one implementation, the computer system can identify inbound display names closely related in meaning to the verified display names associated with an entity, such as an inbound entity display name corresponding to a different spelling of a verified entity display name (e.g., “HR Department” versus “Human Resources Department”).

In one example, the whitelist can include contact information for a first entity—corresponding to a human resources department—within an organization. In particular, the whitelist can include: a verified display name of “Human Resources Department” specified for the human resources department; and a verified email address—such as `HR@companyXYZ.com`—associated with the verified display name.

In this example, the computer system can: intercept an email addressed to a target recipient within an organization, the email from a sender at an inbound email address and including an inbound display name of “HR Department”; access the whitelist to compare the inbound display name to the verified display name specified for the human resources department; and leverage pattern matching techniques and acronym correlation techniques to characterize a display name difference between “HR Department” and “Human Resources Department” as a “low” display name difference—such as represented by a relatively “high” display name correlation—indicating the inbound display name is similar in meaning to the verified display name of the first entity (e.g., the human resources department). In this example, in response to characterizing the display name difference as a “low” display name difference, the computer system can automatically compare the sender email address to a set of verified email addresses—stored for the first entity within the whitelist—and selectively transmit and/or withhold transmittal of the email to the recipient based on inclusion and/or omission of the inbound email address in the set of verified email addresses stored for the first entity.

Additionally and/or alternatively, in the preceding example, the computer system can similarly: intercept an email addressed to a target recipient within the organization, the email from a sender at an inbound email address and including an inbound display name of “HR D3pARTM3NT”; access the whitelist to compare the inbound display name to the verified display name specified for the human resources department; and leverage pattern matching and acronym correlation techniques to characterize a “low” or “moderate” display name difference between the inbound entity display name and the verified entity display name, thereby indicating that the inbound display name is similar or derived from the verified entity display name and therefore at increased risk of corresponding to a spoofing attempt. Therefore, in response to characterizing the display name difference as a “low” or “moderate” display name difference, the computer system can compare the sender email address to the set of verified email addresses—stored for the first entity (e.g., the human resources department) within the whitelist—and selectively transmit and/or withhold transmittal of the email to the recipient based on inclusion and/or omission of the inbound email address in the set of verified email addresses stored for the first entity.

Additionally and/or alternatively, in the preceding example, the computer system can similarly: intercept an email addressed to a target recipient within the organization, the email from a sender at an inbound email address and including an inbound display name of “Sender”; access the whitelist to compare the inbound display name to the verified display name specified for the human resources department; and leverage pattern matching techniques and acronym correlation techniques to characterize a display name difference between “Sender” and “Human Resources Department” as a “high” display name difference—such as represented by a relatively “low” display name correlation—thereby indicating that the inbound display name is distinct and/or relatively unsimilar from the verified entity display name and therefore at relatively low risk of corresponding to a spoofing attempt (e.g., impersonation of the human resources department by a spammer). Therefore, in this example, in response to characterizing the display name difference as a “high” display name difference, the computer system can automatically authorize transmittal of the email to the recipient, such as without further investigation. Additionally and/or alternatively, the computer system can similarly compare the inbound display name to each verified display name—corresponding to each entity (e.g., human resources department, finance department, management branch) within the organization—prior to authorizing transmittal of the email to the designated recipient.

Additionally or alternatively, in another example, the computer system can: identify a set of terms or phrases related to a term or phrase included in the verified display name of the entity; and scan the inbound display name for this set of terms or phrases. In particular, in this example, the computer system can: access a word graph including a population of words distributed about the graph based on correlations between words in the population of words, such that a set of words proximal and/or surrounding a particular word are highly correlated to the particular word, thus possibly representing synonyms (or related words) of the particular word; search the word graph for a first word (or first combination of words) in the verified display name, such as by representing the verified display name in a container (e.g., a vector, a matrix) defining a set of coordinates and projecting the container onto the graph; extract a set of words falling within a threshold distance (e.g., a radius of the threshold distance) of the first word; and, for each word, in the set of words, characterize a difference between the word and a word and/or combination of words forming the inbound display name.

For example, in response to the verified display name corresponding to “Finance Department,” the computer system can: search the word graph for the term “finance”; and extract a set of words—including “billing,” “invoice,” and “payment”—falling within a threshold distance of “finance” within the graph. Then, in response to the inbound display name corresponding to “Billing Department,” the computer system can characterize the difference between the verified display name and the inbound display name as relatively low. Alternatively, in response to the inbound display name corresponding to “Scheduling,” the computer system can characterize the difference between the verified display name and the inbound display name as relatively high, such as based on a distance between “scheduling” and “finance” on the word graph.

Similarly, in another example, the computer system can: access a word graph defining a population of words distributed within the word graph based on correlations between words in the population of words; represent the

inbound display name within the word graph based on a set of features (e.g., words, phrases, characters) extracted from the inbound display name; extract a subset of words, in the population of words, falling within a threshold distance of the inbound display name within the word graph; and, for each word, in the subset of words, characterize a difference between the word and the verified display name. The computer system can therefore: search the word graph for words and/or phrases associated with the inbound display name; and compare these words and/or phrases to the verified display name accordingly.

14.6.3 Variation: Inbound Display Name+Verified Email Address

In one variation, the computer system can verify whether the inbound display name included in the email matches and/or approximates a verified email address associated with an entity defined in the whitelist.

In particular, a spoof or spammer may access an organization directory and identify a verified email address associated with an entity (e.g., a department, a group of employees, a particular employee) within the organization. Therefore, the spoof may send an email—from an inbound email address distinct from the verified email address associated with the entity—including a verified display name matching and/or approximating the verified email address. Therefore, a recipient of the email viewing the mail (e.g., on her mobile device) may view the inbound display name—matching and/or approximating the known, verified email address associated with the entity—and may therefore interpret this email as legitimate. Therefore, to prevent transmission of emails including inbound display names distinct from verified display names associated with an entity—but matching and/or similar to verified email addresses (e.g., complete email address, domain, username) associated with the entity—the computer system can: compare an inbound display name to a set of verified email addresses associated with the entity; and selectively withhold or authorize transmission and/or implement further security checks accordingly.

In one implementation, the computer system can: intercept a first email—received from a first sender at a first inbound email address and including a first inbound display name—addressed to a first target recipient within an organization; access the whitelist; extract a first verified email address, in a set of verified email addresses, associated with a first entity defined in the whitelist; and, characterize a difference between the first inbound display name and the first verified email address.

For example, the computer system can: leverage pattern matching techniques to extract differences between characters (e.g., combination, sequence, quantity) in the inbound display name and the first verified email address; and characterize a display name difference between the inbound display name and the verified email address based on these differences. In particular, in one example, the computer system can: receive an email including an inbound display name of “finance@company.xyz” or ““finance@company.xyz” or ““finance@company”; access the whitelist to identify a verified email address of “finance@company.xyz” for a first entity associated with the organization; and characterize a relatively low difference between the inbound display name and the verified email address, such as based on minimal character differences between the inbound display name and the verified email address. The computer system can therefore identify instances of emails including inbound display names that match or approximate the verified email address, such as

corresponding to a different spelling, including a subset of replacement characters (e.g., “l” in place of “i,” “3” in place of “e”), and/or omitting characters or including additional characters.

5 Additionally or alternatively, in another implementation, the computer system can identify an inbound display name corresponding to a username of a verified email address. In particular, in this implementation, the computer system can: intercept a first email—received from a first sender at a first inbound email address and including a first inbound display name—addressed to a first target recipient within an organization; access the whitelist; extract a first verified email address (e.g., “finance@company.xyz.com”), in a set of verified email addresses, associated with a first entity defined in the whitelist; extract a first username (e.g., “finance”) from the first verified email address; and characterize a difference between the first inbound display name and the first username. For example, to characterize the difference, the computer system can: leverage pattern matching techniques to extract differences between characters (e.g., combination, sequence, quantity) in the inbound display name and the first username; and/or leverage a word chart to identify words exhibiting a relatively high correlation (e.g., within a threshold distance on the word graph) to the username, as described above.

10 Additionally or alternatively, in another implementation, the computer system can identify an inbound display name corresponding to a domain of a verified email address. In particular, in this implementation, the computer system can: intercept a first email—received from a first sender at a first inbound email address and including a first inbound display name—addressed to a first target recipient within an organization; access the whitelist; extract a first verified email address (e.g., “finance@company.xyz.com”), in a set of verified email addresses, associated with a first entity defined in the whitelist; extract a first domain (e.g., “company.xyz”) from the first verified email address; and characterize a difference between the first inbound display name and the first domain. For example, to characterize the difference, the computer system can leverage pattern matching techniques to extract differences between characters (e.g., combination, sequence, quantity) in the inbound display name and the first domain as described above.

14.7 Sender Verification: Inbound Email Address

15 In one implementation, in response to locating the sender display name and/or a similar display name in the whitelist associated with the entity, the computer system can compare the inbound email address to a set of verified email addresses associated with the entity and linked to the display name. The computer system can then: authorize transmission of an email to the recipient in response to identifying the inbound email address in the set of verified email addresses; withhold transmission of the email—and/or and flag the email for further investigation by an administrator—in response to identifying absence of the inbound email address in the set of verified email addresses.

14.8 Entity Ranking

20 In one implementation, the computer system can selectively screen emails based on the entity (e.g., HR, Payroll, Docketing, Marketing) with which the display name in an email is associated. For example, the organization may collect contact information for all employees within the organization (e.g., law firm, company, clinic), upload the contact information to the computer system via a web portal, allocate sets of employees into respective entities, and specify entity rankings within the organization. The computer system can: at a first time, receive the whitelist for the

organization containing contact information of entities in the company; assign each entity an investigation priority level (e.g., low or high investigation priority level, Level I, II, or III) based on entity rankings in the organization as specified in the whitelist (e.g., “low” for “Human Resources”, “high” for “Research and Development”); at a second time, receive an email containing a first entity display name and addressed to a first recipient; access the whitelist of the organization to check for the first entity display name; in response to the first entity display name corresponding to an entity assigned a high investigation priority level, quarantine the email and withhold the email from the first recipient; and flag the email for further investigation and notify the email administrator. If an email administrator determines that the email is legitimate, the computer system can then receive verification of the email from the email administrator and deliver the email to the recipient.

In one variation, upon receiving an email with an entity display name corresponding to a lower investigation priority level within the company and from a first email address, the computer system can: access the whitelist of the organization to check for the entity display name; in response to the entity display name corresponding to a lower investigation priority level entity of the company, search the whitelist to identify that the first email address corresponds to a verified email address of the lower investigation priority level entity; in response to the first email address not matching the verified email address of the lower investigation priority level entity, flag the email as an unverified email from an unverified sender; and deliver the email to a designated recipient including an unverified sender notification. Thus, in this variation, the computer system can verify email senders for entities within the organization while allocating additional resources to higher priority entities (e.g., “R&D”, “Management”, “Legal”) and reducing resources allocated to lower priority entities (e.g., “Creative Services”, “HR”).

14.9 Email Risk

In one variation, the computer system can selectively authorize and/or withhold transmission of an email based on predicted risk associated with this email.

In particular, in this variation, the computer system can extract a set of features—such as including a display name difference, a particular entity associated with the inbound display name and/or corresponding to the inbound email address, and/or a set of email signals (e.g., text and/or language signals, image signals) linked to risk—from an email; and characterize risk associated with the email based on the set of features. For example, the computer system can extract a set of email signals including: a set of text in a body or subject of the email associated with high-risk topics, such as finance, identity, security, credentials, etc.; presence of an attachment linked to the email; presence of a hyperlink within the email; etc.

In one implementation, the computer system can: intercept a first email—addressed to a first target recipient within an organization—received from a first sender at a first inbound email address and including a first inbound display name; and access a whitelist including a first set of contact information corresponding to a first entity associated with the organization and including a first verified display name associated with the first entity and a first set of verified email addresses associated with the first entity. Then, in response to the first set of verified email addresses omitting the first inbound email address, the computer system can: characterize a first display name difference between the first inbound display name and the first verified display name; and, in response to the first display name difference exceeding a

threshold difference, characterize a first risk score—representing risk associated with the first email—based on the first entity, the first display name difference, and a first set of email signals extracted from the first email. Finally, based on the first risk score, the computer system can selectively withhold or authorize transmission of the first email to the first target recipient. For example, in response to the first risk score falling below a threshold risk score, the computer system can authorize transmission of the first email to the first target recipient. Alternatively, in response to the first risk score exceeding the threshold risk score, the computer system can withhold transmission of the first email to the first target recipient and flag the first email for authentication (e.g., by the email administrator).

Therefore, in the preceding implementation, the computer system can further verify whether risk associated with the email exceeds a defined maximum risk, such as prior to authorizing transmission of the email to the target recipient.

Additionally or alternatively, in another implementation, the computer system can: selectively withhold and/or authorize transmission of emails including inbound display names exhibiting display name differences—between the inbound display name and a verified display name on the whitelist—exceeding a first threshold difference and falling below a second threshold difference exceeding the first threshold difference; and automatically authorize transmission of emails including inbound display names exhibiting display name differences exceeding the second threshold difference.

In particular, in this implementation, the computer system can: intercept a first email—addressed to a first target recipient within an organization—received from a first sender at a first inbound email address and including a first inbound display name; and access a whitelist including a first set of contact information corresponding to a first entity associated with the organization and including a first verified display name associated with the first entity and a first set of verified email addresses associated with the first entity. Then, in response to the first set of verified email addresses omitting the first inbound email address, the computer system can: characterize a first display name difference between the first inbound display name and the first verified display name; and, in response to the first display name difference exceeding the first threshold difference and falling below the second threshold difference, characterize a first risk score—representing risk associated with the first email—based on the first entity, the first display name difference, and a first set of email signals extracted from the first email. Then, based on the first risk score, the computer system can selectively withhold or authorize transmission of the first email to the first target recipient. Alternatively, in response to the first display name difference exceeding the second threshold difference, the computer system can automatically authorize transmission of the first email to the first target recipient.

The systems and methods described herein can be embodied and/or implemented at least in part as a machine configured to receive a computer-readable medium storing computer-readable instructions. The instructions can be executed by computer-executable components integrated with the application, applet, host, server, network, website, communication service, communication interface, hardware/firmware/software elements of a user computer or mobile device, wristband, smartphone, or any suitable combination thereof. Other systems and methods of the embodiment can be embodied and/or implemented at least in part as a machine configured to receive a computer-readable medium storing computer-readable instructions. The

instructions can be executed by computer-executable components integrated by computer-executable components integrated with apparatuses and networks of the type described above. The computer-readable medium can be stored on any suitable computer readable media such as RAMs, ROMs, flash memory, EEPROMs, optical devices (CD or DVD), hard drives, floppy drives, or any suitable device. The computer-executable component can be a processor but any suitable dedicated hardware device can (alternatively or additionally) execute the instructions.

As a person skilled in the art will recognize from the previous detailed description and from the figures and claims, modifications and changes can be made to the embodiments of the invention without departing from the scope of this invention as defined in the following claims.

I claim:

1. A method comprising:

during a first time period:

intercepting a first email addressed to a first target recipient within an organization, the first email received from a first sender at a first inbound email address and comprising a first inbound display name; 25 accessing a whitelist comprising a first set of contact information corresponding to a first entity associated with the organization, the first set of contact information comprising:

a first verified display name associated with the first entity; and

a first set of verified email addresses associated with the first entity; and

in response to the first set of verified email addresses omitting the first inbound email address:

characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first entity; and

in response to the first display name difference 40 falling below a threshold difference:

withholding transmission of the first email to the first target recipient; and

flagging the first email for authentication; and

during a second time period:

intercepting a second email addressed to a second target recipient within the organization, the second email received from a second sender at a second inbound email address and comprising a second inbound display name; and

in response to the first set of verified email addresses including the second inbound email address, authorizing transmission of the second email to the second target recipient.

2. The method of claim 1, further comprising, during a 55 third time period:

intercepting a third email addressed to a third target recipient within the organization, the third email received from a third sender at a third inbound email address and comprising a third inbound display name; 60 and

in response to the first set of verified email addresses omitting the third inbound email address:

characterizing a second display name difference between the third inbound display name and the first 65 verified display name associated with the first entity; and

in response to the second display name difference exceeding the threshold difference, authorizing transmission of the third email to the third target recipient.

5. 3. The method of claim 2, wherein authorizing transmission of the third email to the third target recipient in response to the second display name difference exceeding the threshold difference comprises, in response to the second display name difference exceeding the threshold difference:

10 characterizing a third display name difference between the third inbound display name and a first verified email address, in the set of verified email addresses, associated with the first entity;

in response to the third display name difference falling below the threshold difference:

withholding transmission of the third email to the third target recipient; and

flagging the third email for authentication; and

in response to the third display name difference exceeding the threshold difference, authorizing transmission of the third email to the third target recipient.

4. The method of claim 2:

wherein accessing the whitelist comprising the first set of contact information comprises accessing the whitelist comprising the first set of contact information comprising:

the first verified display name and a second verified display name associated with the first entity; and

the first set of verified email addresses; and

wherein authorizing transmission of the third email to the third target recipient in response to the second display name exceeding the threshold difference comprises in response to the second display name exceeding the threshold difference:

characterizing a third display name difference between the third inbound display name and the second verified display name; and

in response to the third display name difference exceeding the threshold difference, authorizing transmission of the third email to the third target recipient.

5. The method of claim 2:

wherein withholding transmission of the first email to the first target recipient and flagging the first email for authentication in response to the first display name difference falling below the threshold difference comprises, in response to the first display name difference falling below the threshold difference:

characterizing a first risk score for risk associated with the first email based on the first entity, the first display name difference, and a first set of email signals extracted from the first email; and

in response to the first risk score exceeding a threshold risk score, withholding transmission of the first email to the first target recipient and flagging the first email for authentication;

wherein authorizing transmission of the third email to the third target recipient in response to the second display name exceeding the threshold difference comprises, in response to the second display name exceeding the threshold difference:

characterizing a second risk score for risk associated with the second email based on the first entity, the second display name difference, and a second set of email signals extracted from the third email; and

in response to the second risk score falling below the threshold risk score, authorizing transmission of the third email to the third target recipient.

6. The method of claim 1:
 wherein intercepting the first email received from the first sender at the first inbound email address comprises intercepting the first email received from the first sender at the first inbound email address comprising a first domain;
 wherein accessing the whitelist comprising the first set of contact information comprises accessing the whitelist comprising:
 a set of verified domains associated with the organization;
 the first set of contact information corresponding to the first entity and comprising the first verified display name and the first set of verified email addresses comprising a first verified domain in the set of verified domains; and
 a second set of contact information corresponding to a second entity associated with the organization and comprising:
 a second verified display name associated with the second entity; and
 a second set of verified email addresses associated with the second entity and comprising a second verified domain in the set of verified domains;
 wherein characterizing the first display name difference in response to the first set of verified email addresses omitting the first inbound email address comprises characterizing the first display name difference in response to the set of verified domains omitting the first domain;
 wherein intercepting the second email received from the second sender at the second inbound email address comprises intercepting the second email received from the second sender at the second inbound email address comprising a second domain; and
 wherein authorizing transmission of the second email to the second target recipient in response to the first set of verified email addresses including the second inbound email address comprises authorizing transmission of the second email to the second target recipient in response to the set of verified domains including the second domain.

7. The method of claim 6, wherein accessing the whitelist comprising the first set of contact information corresponding to the first entity and the second set of contact information corresponding to the second entity comprises accessing the whitelist comprising:
 the first set of contact information corresponding to the first entity comprising an internal department within the organization; and
 the second set of contact information corresponding to the second entity comprising an external customer associated with the organization.

8. The method of claim 1:
 wherein intercepting the first email received from the first sender at the first inbound email address and comprising the first inbound display name comprises intercepting the first email received from the first sender at the first inbound email address and comprising the first inbound display name comprising a first set of characters;
 wherein accessing the whitelist comprising the first set of contact information comprising the first verified display name and the first set of verified email addresses comprises accessing the whitelist comprising the first set of contact information comprising the first verified display name and the first set of verified email

addresses, the first verified display name comprising a set of target characters; and
 wherein characterizing the first display name difference between the first inbound display name and the first verified display name comprises characterizing the first difference between the first set of characters and the target set of characters.

9. The method of claim 1, wherein characterizing the first difference comprises, in response to the first inbound display name differing from the verified display name:
 accessing a word graph defining a population of words distributed within the word graph based on correlations between words in the population of words;
 representing the first inbound display name within the word graph based on a set of features extracted from the first inbound display name;
 extracting a subset of words, in the population of words, falling within a threshold distance of the first inbound display name within the word graph; and
 for a first word, in the subset of words, characterizing the first difference between the word and the verified display name.

10. The method of claim 1:
 wherein accessing the whitelist comprising the first set of contact information comprises accessing the whitelist comprising the first set of contact information and a second set of contact information corresponding to a second entity associated with the organization, the second set of contact information comprising:
 a second verified display name associated with the second entity; and
 a second set of verified email addresses associated with the second entity; and
 wherein withholding transmission of the first email to the first target recipient and flagging the first email for authentication further comprises:
 in response to the second set of verified email addresses omitting the first inbound email address, characterizing a second display name difference between the first inbound display name and the second verified display name associated with the second entity; and
 in response to the second display name difference falling below the threshold difference:
 withholding transmission of the first email to the first target recipient; and
 flagging the first email for authentication.

11. The method of claim 1:
 wherein flagging the first email for authentication comprises:
 generating a notification comprising a hyperlink to the first email and a prompt to verify authenticity of the first sender; and
 transmitting the notification to an email administrator associated with the organization; and
 further comprising, in response to receiving verification of the first sender from the email administrator:
 authorizing transmission of the first email to the first target recipient; and
 appending the first set of contact information associated with the first entity to include the first inbound display name and the first inbound email address.

12. A method comprising:
 during a first time period:
 intercepting a first email addressed to a first target recipient within an organization, the first email received from a first sender at a first inbound email address and comprising a first inbound display name;

37

accessing a whitelist comprising a first set of contact information corresponding to a first entity associated with the organization, the first set of contact information comprising:

a first verified display name associated with the first entity; and

a first set of verified email addresses associated with the first entity;

characterizing a first display name difference between the first inbound display name and the first verified display name associated with the first entity; and

in response to the first display name difference falling below a threshold difference and in response to the first set of verified email addresses omitting the first inbound email address:

withholding transmission of the first email to the first target recipient; and

flagging the first email for authentication; and

20 during a second time period:

intercepting a second email addressed to a second target recipient within the organization, the second email received from a second sender at a second inbound email address and comprising a second inbound display name;

characterizing a second display name difference between the second inbound display name and the first verified display name associated with the first entity; and

in response to the second display name difference falling below the threshold difference and in response to the first set of verified email addresses including the second inbound email address, authorizing transmission of the second email to the second target recipient.

25

13. The method of claim 12, further comprising, during a third time period:

intercepting a third email addressed to a third target recipient within the organization, the third email received from a third sender at a third inbound email address and comprising a third inbound display name;

characterizing a third display name difference between the third inbound display name and the first verified display name associated with the first entity; and

in response to the third display name difference exceeding the threshold difference, authorizing transmission of the third email to the third target recipient.

30

14. The method of claim 13, wherein authorizing transmission of the third email to the third target recipient in response to the third display name difference exceeding the threshold difference comprises, in response to the third display name difference exceeding the threshold difference:

characterizing a fourth display name difference between the third inbound display name and a first verified email address, in the set of verified email addresses, associated with the first entity;

in response to the fourth display name difference falling below the threshold difference:

withholding transmission of the third email to the third target recipient; and

flagging the third email for authentication; and

in response to the fourth display name difference exceeding the threshold difference, authorizing transmission of the third email to the third target recipient.

35

40

45

50

55

60

65

38

15. The method of claim 12:

further comprising, during the first time period:

characterizing a third display name difference between the first inbound display name and a first email address in the set of verified email addresses; and

in response to the third display name difference falling below a second threshold difference and in response to the first set of verified email addresses omitting the first inbound email address:

withholding transmission of the first email to the first target recipient; and

flagging the first email for authentication;

10 wherein characterizing the first display name difference between the first inbound display name and the first verified display name comprises in response to the first difference exceeding the second threshold difference, characterizing the first display name difference between the first inbound display name and the first verified display name; and

15 further comprising, during the second time period:

characterizing a fourth display name difference between the second inbound display name and the first email address in the set of verified email addresses; and

in response to the fourth display name difference falling below the second threshold difference and in response to the first set of verified email addresses omitting the second inbound email address:

withholding transmission of the second email to the second target recipient; and

flagging the second email for authentication.

20

16. The method of claim 12, further comprising, in response to the second display name difference falling below the threshold difference and in response to the first set of verified email addresses including the second inbound email address:

generating a notification comprising a prompt to verify the second display name for emails sent from the second email address;

transmitting the prompt to the second sender at the second email address; and

in response to receiving verification of the second display name from the second sender, appending the whitelist to include the second display name associated with the first entity.

25

17. The method of claim 12:

wherein accessing the whitelist comprising the first set of contact information comprising the first verified display name and the first set of verified email addresses comprises accessing the whitelist comprising the first set of contact information comprising the first verified display name and the first set of verified email addresses comprising a first verified domain associated with the first entity; and

30

wherein characterizing the first display name difference between the first inbound display name and the first verified display name comprises:

extracting a first domain of the first inbound email address;

in response to the first domain differing from the first verified domain, characterizing the first display name difference between the first inbound display name and the first verified display name; and

in response to the first domain matching the first verified domain, authorizing transmission of the first email to the first target recipient.

35

40

45

39

18. A method comprising:
 during a first time period:
 intercepting a first email addressed to a first target
 recipient within an organization, the first email com-
 prising a first inbound display name and received
 from a first sender at a first inbound email address
 comprising a first domain;
 5
 accessing a whitelist comprising:
 a set of verified domains associated with inbound
 emails received by recipients within the organization;
 10 and
 a first set of contact information corresponding to a
 first entity associated with the organization, the
 first set of contact information comprising:
 a first verified display name associated with the
 first entity; and
 15 a first set of verified email addresses associated
 with the first entity; and
 in response to the set of verified domains omitting the
 first domain and in response to the first set of verified
 email addresses omitting the first inbound email
 address:
 20 characterizing a first display name difference
 between the first inbound display name and the
 first verified display name; and
 in response to the first display name difference
 falling below a threshold difference:
 withholding transmission of the first email to the
 first target recipient; and
 25 flagging the first email for authentication; and
 during a second time period:
 intercepting a second email addressed to a second
 target recipient within the organization, the second
 email:
 30

40

received from a second sender at a second inbound
 email address comprising a second domain; and
 comprising a second inbound display name; and
 in response to the set of verified domains including the
 second domain, authorizing transmission of the sec-
 ond email to the second target recipient.
 19. The method of claim 18, further comprising, during a
 third time period:
 intercepting a third email addressed to a third target
 recipient within the organization, the third email com-
 prising a third inbound display name and received from
 a third sender at a third inbound email address;
 15 in response to the set of verified domains omitting the
 third domain and in response to the first set of verified
 email addresses including the third inbound email
 address, authorizing transmission of the third email to
 the third target recipient.
 20. The method of claim 18, further comprising, during a
 third time period:
 intercepting a third email addressed to a third target
 recipient within the organization, the third email com-
 prising a third inbound display name and received from
 a third sender at a third inbound email address; and
 in response to the set of verified domains omitting the
 third domain and in response to the first set of verified
 email addresses omitting the third inbound email
 address:
 characterizing a second display name difference
 between the third inbound display name and the first
 verified display name; and
 in response to the second display name difference
 exceeding the threshold difference, authorizing
 transmission of the third email to the third target
 recipient.

* * * * *