DECEPTION-BASED RESPONSES TO SECURITY ATTACKS

RELATED APPLICATIONS

[0001] This patent application claims priority to U.S. Utility patent application entitled "Deception-Based Responses to Security Attacks" with Serial No. 13/784,720, filed March 4, 2013, which is fully incorporated herein by reference.

BACKGROUND

[0002] With Internet use forming an ever greater part of day to day life, security exploits that steal or destroy system resources, data, and private information are an increasing problem. Governments and businesses devote significant resources to preventing intrusions and thefts related to these security exploits. Security threats come in many forms, such as computer viruses, worms, trojan horses, spyware, keystroke loggers, adware, and rootkits. These threats typically employ security exploits, which are the weaponization of an attack against a specific vulnerability in software. These threats are delivered in or through a number of mechanisms, such as spearfish emails, clickable links, documents, executables, or archives. Some of the threats posed by security exploits are of such significance that they are described as cyber terrorism or industrial espionage.

[0003] To meet the threat posed by these security exploits, many security solutions, such as antivirus software, have been developed. Typically, these solutions scan a computing device, determine if the device is affected by a

security threat, and block or remove the security threat. While blocking and removing counter the immediate threat, they are easy for an adversary purveying the security vulnerability to overcome. The affected user is then left always one step behind, always reacting to actions taken by an adversary rather than taking steps to prevent future intrusions and thefts by the adversary.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The detailed description is set forth with reference to the accompanying figures. In the figures, the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The use of the same reference numbers in different figures indicates similar or identical items or features.

[0005] FIG. 1 illustrates an example network connecting a security service to client devices, the security service providing deception-based techniques for responding to attacks affecting the client devices.

[0006] FIG. 2 illustrates an example process for transitioning an attack to device monitored device posing as the computing device impacted by the attack and enabling the attack to obtain deceptive information from the monitored device.

[0007] FIG. 3 illustrates an example process for configuring a document to report identifying information of an entity opening the document to enable determination of whether an unauthorized entity is in possession of the document.

[0008] FIG. 4 illustrates an example process for determining that a domain specified in a domain name request is associated with malicious activity and responding to the request with a network address of an monitored device to cause the requesting process to communicate with the monitored device in place of an adversary server.

[0009] FIG. 5 illustrates an example process for monitoring dormant domains names associated with malicious activity and, in response to a change, responding with an alert or a configuration update.

[0010] FIG. 6 illustrates a component level view of a computing device capable of acting as a security service device, a client device, or a monitored device.

DETAILED DESCRIPTION

Overview

[0011] This disclosure describes, in part, deception-based techniques for responding to attacks. The techniques include transitioning an attack to a monitored computing process or device (referred to herein as a "monitored device") that is posing as the computing device impacted by the attack and enabling the adversary associated with the attack to obtain deceptive information from the monitored device. The techniques also include identifying the adversary associated with the attack by configuring a document to report identifying information when an unauthorized entity opens the document. Further, the techniques include determining that a domain specified in a domain name request is associated with malicious activity and responding

to the request with a network address of a monitored computing process or device (referred to herein as a "monitored server" or a "monitored device") to cause the requesting process to communicate with the monitored server in place of an adversary server. Additionally, the techniques include monitoring dormant domains names associated with malicious activity and, in response to a change in name resolution or registration, responding with an alert or a configuration update.

[0012] In various embodiments, a security service may implement or enable any one or more of the deception-based techniques. In a first set of techniques, the security service may utilize monitored devices to run and monitor attacks and to use those attacks to provide deceptive information to an adversary. The adversary may then spend time and effort working on the deceptive information without realizing that the attack has been blocked. Future attacks are thus inhibited as the attacked gains the advantage over the attacker.

[0013] When first detecting an attack, a security agent or the security service may block processing of the attack by the impacted device. The security service then transitions the attack to the monitored device and processes the attack on that monitored device. The monitored device poses as the impacted computing device and is configured with a virtual image of the impacted computing device to enable a convincing pose. The security service then monitors activities and communications of the attack and uses information gained from the monitoring to improve deception capabilities and security agent configurations. The security service also enables an entity associated

with the impacted computing device to provide deceptive information to be included on the monitored device. For example, if the adversary is a negotiation partner for a business deal attempting to illicitly acquire an entity's negotiating strategy, the entity could provide a false strategy for inclusion on the monitored device.

[0014] In a second set of techniques, the security service uses beaconing to identify the adversary, thus enabling better targeting of deceptive information. The security service embeds executable instructions in a document that causes the document to report identifying information to the security service when opened. The security service may also or instead embed a link in the document that causes a request to a monitored network address (e.g., when the document is opened and the link is clicked). The security service is then able to receive the report or request to the monitored network address, determine from it whether the opener is an unauthorized adversary, and determine characteristics of the adversary, such as geographic location. These characteristics can be used to craft better deceptive information or more convincing monitored devices or monitored servers. The security service may also alert a user or security agent program of the unauthorized possession or update a security agent program configuration.

[0015] In a third set of techniques, the security service detects the presence of a attack through the attack's specification of a suspicious domain name in a domain name resolution request. In response, the security service deceives the adversary associated with the attack by responding to the request with a

network address of an monitored server, the monitored server posing as an adversary server. The monitored server then gathers information from the attack that may be utilized to enhance the effectiveness of other deception techniques, such as improving the deceptive information made available through a monitored device. Also, the security service may alert at least one of a security agent or users of an entity impacted by the attack. In some embodiments, the security service may then transition the attack to a monitored device and perform the above described first set of deception techniques.

[0016] In a fourth set of techniques, the security service monitors dormant domain names that have previously been used in connection with malicious activity. By monitoring the domain names for changes in name resolution or registration, the security service may be prepared, having monitored devices or monitored servers ready to use when a attack utilizes one of these monitored domain names. With such techniques, the security service may also perform additional actions, such as detecting other dormant domains and monitoring them, or responding to the change in name resolution or registration by providing an alert or a configuration update.

Example Network

[0017] FIG. 1 illustrates an example network connecting a security service to client devices, the security service providing deception-based techniques for responding to attacks impacting the client devices. As illustrated in FIG. 1, a network 102 connects client devices 104 of a client network 106 to a security

service 108 of a security service network 110. The security service 108 may configure the client devices 104 with security agents 112, the security agents 112 being capable of detecting attacks 114 of an adversary 116 that are directed at one or more of the client devices 104. In response, the security service 108 may transition the attack to a monitored device 118 included in the client network 106 or an monitored device 120 included in the security service network 110. The monitored device 118/120 may pose as the client device 104 that the attack is directed at and may be configured with a virtual image of that client device 104. The monitored device 118/120 may also include deceptive information 122 for the attack 114 to obtain and provide to the adversary 116. The attack 114 may provide the deceptive information 122 to one or both of a command-and-control (C2) server 124 of the adversary 116 or an exfiltration server 126 of the adversary 116. In some embodiments, the security service 108 may also receive a redirection domain name resolution request made by the attack 114 and may respond to the request by identifying an monitored server 128 included in the client network 106 or an monitored server 130 included in the security service network 110, the monitored server 128/130 posing as the C2 server 124 or exfiltration server 126 to the attack 114.

[0018] In various embodiments, devices of the security service 108 may include modules and data 132-146 that enable the security service 108 to perform the operations described herein. These modules and data 132-146 include a monitoring module 132, a deception module 134, a beaconing module 136, a domain name service (DNS) redirection module 138, an

dormant domain repository 140, a response module 142, an analysis module 144, and a configuration module 146.

[0019] In some embodiments, the network 102, client network 106, and security service network 110 may each include any one or more networks, such as wired networks, wireless networks, and combinations of wired and wireless networks. Further, the network 102, client network 106, and security service network 110 may include any one or combination of multiple different types of public or private networks (e.g., cable networks, the Internet, wireless networks, etc.). For example, the network 102 may be a public network and the client network 106 and security service network 110 may each be a private network. In some instances, computing devices communicate over the network 102, client network 106, and security service network 110 using a secure protocol (e.g., https) and/or any other protocol or set of protocols, such as the transmission control protocol/Internet protocol (TCP/IP). Further, each of the client network 106 and security service network 110 may be assigned a block of network addresses, such as IP addresses.

[0020] In various embodiments, the client devices 104, the computing devices of the security service 108, the monitored device 118/120, and the monitored server 128/130 may each be or include a server or server farm, multiple, distributed server farms, a mainframe, a work station, a personal computer (PC), a laptop computer, a tablet computer, a personal digital assistant (PDA), a cellular phone, a media center, an embedded system, or any other sort of device or devices. In one implementation, the computing devices

of the security service 108 represent a plurality of computing devices working in communication, such as a cloud computing network of nodes. When implemented on multiple computing devices, the security service 108 may distribute the modules and data 132-146 of the security service 108 among the multiple computing devices. In some implementations, one or more of the client devices 104, the computing devices of the security service 108, the monitored device 118/120, and the monitored server 128/130 represent one or more virtual machines implemented on one or more computing devices. An example computing device capable of serving as a client device 104, a computing device of the security service 108, the monitored device 118/120, or the monitored server 128/130 is illustrated in FIG. 6 and described below with reference to that figure.

[0021] In various embodiments, the client devices 104 and client network 106 may be associated with an entity, such as a business, government, or other organization, or may be associated with a single user or groups of users. That entity or those user(s) may subscribe for security services with the security service 108, which may monitor activities on client devices 104 of the entity/user(s). In some embodiments, that monitoring may involve installing security agents 112 on the client devices 104 to monitor execution activities of the client devices 104 and to report on those activities to the security service 108. In other embodiments, the security service 108 may instead rely on scanning the client devices 104 from a security service scanner or rely on a

third party security product or service to detect attacks and communicate them to the security service 108.

[0022] Each client device 104 may be any sort of computing device, as described above. Each client device 104 may be configured with an operating system, applications, folders, files, and desktop screens. Further, each client device 104 may be associated with a user name of a user of that client device 104, a machine name of the client device 104, an operating system version, desktop screens, folder names, preloaded files, and computer firmware versions.

[0023] In some embodiments, the security agent 112 may be a kernel-level security agent that observes and acts upon execution activities of its corresponding client device 104. The security agent 112 may be configurable by the security service 108, receiving, and applying while live, reconfigurations of filters, components, models, etc. of the security agent 112. Based on the observed execution activities, the security agents 112 may generate security information which the security agent 112 may act upon and/or provide to the security service 108. An example security agent 112 is described in greater detail in U.S. patent application serial number 13/492,672, entitled "Kernel-Level Security Agent" and filed on June 8, 2012.

[0024] As mentioned, attacks 114 may be directed at client devices 104. Such attacks 114 comes in many forms, such as computer viruses, worms, trojan horses, spyware, keystroke loggers, adware, and rootkits. These attacks 114 are delivered in or through a number of mechanisms, such as spearfish

emails, clickable links, websites, drive by exploits, QR codes, Near Field Communications (NFC) triggered links, documents, executables, removable drives, or archives. The attacks 114 may be detected by the security agents 112 or other security mechanism and may be blocked or stopped from further processing. As mentioned further herein, the attacks 114 may then be transitioned to a monitored device 118/120 by the security service 108.

[0025] In various embodiments, the security service 108 may be a provider of information security services to client entities, such as maintenance and configuration of the kernel-level security agents 112, threat modeling, and/or remediation. In some embodiments, the security service 108 may also provide social aspects to the security services, forming groups of those client entities and automatically sharing security information among the client entities 104 constituting a group.

[0026] In addition, as described above in some detail, the security service 108 implements any of a number of deception techniques to respond to attacks 114. The modules and data 132-146, monitored devices 118/120, and monitored servers 128/130 may be operated and used by the security service 108 to carry out these techniques.

[0027] In various embodiments, the monitoring module 132 may be configured to perform monitoring operations associated with a number of the techniques. For example, in associated with the first set of deception techniques described above, the monitoring module 132 may be configured to detect the presence or operation of attacks 114 on client devices 104 or to

receive automated or manual notifications from security agents 112 of the presence or operation of attacks 114. In one embodiment, receiving the automated or manual notification may include retrieving the attack from an information sharing system or portal to which the attack has been submitted. The information sharing system or portal may have received the attack from another entity that is different from the entity associated with client devices 104. In some embodiments, the attack 114 may be included in an email (e.g., as a link or attachment), and a user of the client device 104 may forward the email to the security service 108. In response to detection or received notice, the monitoring module 132 may invoke the deception module 134.

[0028] The deception module 134 may take a number of actions to counter the attack 114. First, the deception module 134 may block processing (or halt further processing) of the attack 114 by the affected client device 104. In some embodiments, however, such blocking or halting may be accomplished by the security agent 112 and need not involve the deception module 134.

[0029] The deception module 134 may then configure a device on the client network 106 as a monitored device 118 or a device on the security service network 110 as a monitored device 120. To configure the device, the security service 108 may capture and apply a virtual image of the client device 104. In some embodiments, the monitored device 118 may be implemented in a virtual machine or partition of the client device 104 (e.g., as an isolated sandbox). Such a virtual image may include a number of attributes of the client device 104, such as its user name, machine name, operating system version, desktop

screens, folder names, preloaded files or computer firmware versions. In other embodiments, the security service 108 may receive this information from a user of the client device 104 rather than capturing the virtual image. If configuring a device on the security service network 110, the deception module may utilize the Border Gateway Protocol (BGP) to have an IP address associated with the client network 106 assigned to the monitored device 120, thereby enabling the monitored device 120 to convincingly pose as a device of the client network 106. After configuring the device to create a monitored device 118/120, the deception module 134 may transition the attack 114 from the client device 104 to the monitored device 118/120 and may process the attack 114 on the monitored device 118/120.

[0030] By processing the attack 114 on the monitored device 118/120, the deception module 134 enables monitoring of attack activities and deception of the adversary 116. In some embodiments, the monitoring module 132 may monitor the activities of the attack 114 as it executes. For example, the monitoring module 132 may monitor commands and events issued by the adversary or monitor network activity, file activity, process activity, execution activity, registry activity, operating system activity, firmware updates, kernel extensions, or loaded drivers of the monitored device 118/120. The commands and events issued may include keystrokes, mouse activity, or command line interface activity. Further, the monitoring module 132 may intercept and decode communications from the attack 114 to the adversary 116. The monitoring module 132 may decode the communications by determining the

protocol that the communications are using. If the protocol is not familiar, the monitoring module 132 may invoke the analysis module 144 to process the intercepted communications and determine the protocol used.

[0031] In some embodiment, the monitoring module 132 may further invoke the response module 142 to provide an alert to the security agents 112 or provide human-consumable intelligence to the client entity/user(s). Alternatively or additionally, the monitoring module 132 or response module 142 may invoke the configuration module 146 to reconfigure the security agents 112 to address the attack 114. The reconfiguration may be based on the activity of the attack 114 captured by the monitoring module 132. In some embodiments, the configuration module 146 may correlate attack activity information received from a number of monitored devices 118/120 and use these correlations in reconfiguring the security agents 112. In some embodiments, the different monitored devices may be associated with different entities.

[0032] While the monitoring module 132 monitors activities of the attack 114, the deception module 134 may take actions to deceive the adversary 116. First, the deception module 134 may configure the monitored device 118/120 to mimic user activity by for example, opening a browser window, generating keystrokes, etc. By mimicking user activity, the monitored device 118/120 ensures that its pose will be convincing to an adversary 116.

[0033] In addition to configuring the client device 104 to mimic user activity, the deception module 134 enables the client entity/user(s) to load

deceptive information 122 onto the monitored device 118/120. As mentioned above, the adversary 116 could be a competitor or enemy engaged in espionage, and the deceptive information 122 could pose as real information, causing the adversary 116 to waste time and effort analyzing the deceptive information 122. The deception module 134 also enables the attack 114 to obtain the deceptive information 122 and to provide the deceptive information 122 to the adversary 116.

[0034] In some embodiments, the deception module 134 may maintain the monitored device 118/120 and associated deception over an extended period of time (e.g., weeks, months, or years), feeding new deceptive information to the adversary 116 through its attack 114 throughout the extended period.

[0035] In various embodiments, the beaconing module 136 may operate independently of the monitoring module 132 and deception module 134. When a document is created on a client device 104, the beaconing module 136 may configure the document to report identifying information of the device opening the document. Such identifying information may include at least one of a network address, a geographic location, a universally unique identifier (UUID), domain information, or derived/upstream network data. Derived/upstream network data may include upstream router network addresses, network address translation (NAT) data, etc. The beaconing module 136 may achieve this configuration by embedding executable instructions in the document that open a connection to the security service 108 when the document is opened. Alternatively, the beaconing module 136 may achieve this configuration by

embedding a link in the document that causes a request to be submitted to a monitored network address. In some embodiments, the executable instructions may specify expected identifying information, such as an expected network address or range or network addresses, and an expected geographic location or range of geographic locations. In such embodiments, the document may only establish a connection to the security service 108 or cause a request to be submitted to the monitored network address when the determined network address or geographic location of the document opener differs from the expected address/location. In other embodiments, rather than the beaconing module 136 configuring the documents of the client device 104, the security agents 112 may perform the configuring.

[0036] In further embodiments, the beaconing module 136 may receive the reports or requests sent to the monitored network address from the documents and, in response, determine whether the adversary 116 or some other unauthorized entity/user is in possession of the documents. The beaconing module 136 may perform this determination based on expected identifying information such as an expected network address or expected range of network addresses, or based on expected geographic location(s). Such expected information may be received from the client entity/user(s) at the time of registration with the security service 108 or may be received from the security agents 112.

[0037] In some embodiments, when the document is determined to be possessed by the adversary 116 or another unauthorized entity, the beaconing

module 136 may invoke the response module 142 to provide an alert to the security agents 112 or provide human-consumable intelligence to the client entity/user(s). Alternatively or additionally, the beaconing module 136 or response module 142 may invoke the configuration module 146 to reconfigure the security agents 112 to address the exfiltration of the document.

[0038] In various embodiments, the DNS redirect module 138 may operate independently of the beaconing module 136. The DNS redirect module 138 may configure a DNS server of the client network 106 to treat the DNS redirect module 138 as a higher-level DNS and to redirect domain name resolution requests to the DNS redirect module 138. The DNS redirect module 138 may then receive domain name resolution requests and invoke the analysis module 144 to determine whether the domain names included in the domain name resolution requests are associated with malicious activity. In some embodiments, the analysis module 144 makes the determination by determining whether a domain name is included in a list of known malicious or suspicious domains. The analysis module 144 may also or instead determine that the domain name is unfamiliar or is associated with a specific geographic location (e.g., a known geographic location of a known adversary 116 or class or group of adversaries 116).

[0039] In some embodiments, when the analysis module 114 informs the DNS redirect module 138 that a domain name is associated with malicious activity, the DNS redirect module 138 invokes the deception module 134 to configure a device in the client network 106 as an monitored server 128 or a

device in the security service network 110 as an monitored server 130. The monitored server 128/130 may pose as a C2 server 124 or an exfiltration server 126 of an adversary 116. To make the pose convincing, the deception module 134 may configure the monitored server 128/130 with a plurality of adversary protocols that may be used in decoding communications from a attack 114 that made the domain name resolution request. The DNS redirect module 138 may then respond to the domain name resolution request with a network address of the monitored server 128/130 to cause the attack 114 that made the domain name resolution request to communicate with the monitored server 128/130.

[0040] In further embodiments, the DNS redirect module 138 may invoke the response module 142 to provide an alert to the security agents 112 or provide human-consumable intelligence to the client entity/user(s). Alternatively or additionally, the DNS redirect module 138 or response module 142 may invoke the configuration module 146 to reconfigure the security agents 112 to address the exfiltration of the document.

[0041] Additionally, in some embodiments, the DNS redirect module 138 may invoke the deception module to configure a monitored device 118/120 and to transition the attack 114 to the monitored device 118/120. Such transitioning may take place after the attack 114 has communicated with the monitored server 128/130 for some time period.

[0042] In various embodiments, the dormant domain name repository 140 may be a repository of dormant domain names that were previously associated with malicious activity. The security service 108 may create the dormant

domain name repository 140 from other available repositories of dormant domain names previously associated with malicious activity, by constructing the dormant domain name repository 140 as suspicious dormant domain names are detected, or some combination of both. The security agents 112 may be configured to identify as a suspicious, dormant domain name any locally-resolving domain name (e.g., resolves to network address 127.0.0.1) or any irresolvable domain name. The security agents 112 may inform the security service 108 of these, and the security service 108 may add them to the dormant domain name repository 140.

[0043] In some embodiments, the monitoring module 132 may monitor for changes to name resolutions or registrations of the domain names included in the dormant domain name repository 140. The monitoring module 132 may perform this monitoring for changes periodically or continuously. In response to detecting a change, the monitoring module 132 may further invoke the response module 142 to provide an alert to the security agents 112 or provide human-consumable intelligence to the client entity/user(s). Alternatively or additionally, the monitoring module 132 or response module 142 may invoke the configuration module 146 to reconfigure the security agents 112 to appropriately handle the now-active domain names.

[0044] In some embodiments, the deception module 134 may also prepare monitored devices 118/120 or monitored servers 128/130 for each dormant domain name in the dormant domain names repository 140, in the manner described above.

[0045] In some instances, any or all of the client devices 104, the security service 108, the monitored device 118/120, or the monitored server 128/130 may have features or functionality in addition to those that FIG. 1 illustrates. For example, some or all of the functionality described as residing within any or all of the client devices 104, the security service 108, the monitored device 118/120, or the monitored server 128/130 may reside remotely from that/those device(s).

Example Processes

[0046] FIGs. 2-5 illustrate example processes 200, 300, 400, and 500. These processes are illustrated as logical flow graphs, each operation of which represents a sequence of operations that can be implemented in hardware, software, or a combination thereof. In the context of software, the operations represent computer-executable instructions stored on one or more computer-readable storage media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular abstract data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be combined in any order and/or in parallel to implement the processes.

[0047] FIG. 2 illustrates an example process for transitioning an attack to an monitored device posing as the computing device impacted by the attack and

enabling the adversary associated with the attack to obtain deceptive information from the monitored device. The process 200 includes, at 202, receiving, by a security service, automated or manual notification of an attack directed at a computing device. In some embodiments, receiving the automated or manual notification may comprise retrieving the attack from an information sharing system or portal to which the attack has been submitted. The attack may be one of a spearfish email, a clickable link, a website, a drive by exploit, a QR code, a Near Field Communications (NFC) triggered link, a document, an executable, a removable drive, or an archive. Further, the notification may be received from a user forwarding the attack or from a security agent monitoring the computing device.

[0048] At 204, the security service blocks processing of the attack by the computing device that the attack is targeted at.

[0049] At 206, the security service transitions the attack to a monitored device, the monitored device configured to pose as the affected computing device to an adversary associated with the attack. The monitored device may be a physical machine or virtual machine configured with a virtual image of the computing device. The virtual image may include at least one or more of a user name, a machine name, an operating system version, desktop screens, folder names, preloaded files or computer firmware versions associated with the affected computing device. Also, monitored device may be located on a network of an entity associated with the impacted computing device or implemented remotely by the security service. If implemented remotely, the

monitored device may be assigned a network address associated with the entity. Further, the monitored device may mimic user activity.

[0050] At 208, the security service processes the attack on the monitored device. While processing the attack, the security service may, at 210, monitor commands and events issued by the adversary or monitor network activity, file activity, process activity, execution activity, registry activity, operating system activity, firmware updates, kernel extensions, or loaded drivers of the monitored device. The commands and events issued may include keystrokes, mouse activity, or command line interface activity. At 212, the security service correlates information obtained from monitoring attacks on multiple monitored devices. At 214, while monitoring, the security service intercepts and decodes communications from the attack to the adversary system. At 216, the security service updates configuration of security agent implemented on the affected device based at least in part on the monitoring or provides human-consumable intelligence.

[0051] At 218, the security service enables an entity associated with the affected computing device to load deceptive information onto the monitored device posing as that impacted device. At 220, the security service may then enable the adversary associated with the attack to obtain the deceptive information. In some embodiments, the monitoring of the attack and the enabling of it to provide deceptive information may occur over an extended period of time (e.g., weeks, months, or years). While performing the enabling

at 220, the security service may receive intelligence from the adversary about at least one of the adversary's tools, tactics, techniques, or procedures.

[0052] FIG. 3 illustrates an example process for configuring a document to report identifying information of an entity opening the document to enable determination of whether an unauthorized entity is in possession of the document. The process 300 includes, at 302, configuring a document to report identifying information of a device opening the document to the security service. Such identifying information may include at least one of a network address, a geographic location, a universally unique identifier (UUID), domain information, or derived/upstream network data. In some embodiments, the configuring includes embedding executable instructions in the document. The executable instructions embedded in the document may be configured to open a connection to the security service and to perform the reporting based at least in part on whether the identifying information differs from identifying information embedded in the document. Alternatively, the configuring may include embedding a link in the document that causes a request to be submitted to a monitored network address.

[0053] At 304, the security service may receive a report associated with the document, the report including the identifying information.

[0054] At 306, the security service may determine, based on the at least one of the identifying information, that unauthorized entity is in possession of the document.

[0055] At 308, the security service may alert an entity associated with the document of the possession by the unauthorized entity.

[0056] FIG. 4 illustrates an example process for determining that a domain specified in a domain name request is associated with malicious activity and responding to the request with a network address of an monitored server to cause the requesting process to communicate with the monitored server in place of an adversary server. The process 400 includes, at 402, redirecting a domain name resolution request from a domain name server associated with a device or entity implementing a process that submitted the domain name resolution request.

[0057] At 404, a security service receives the redirected domain name resolution request.

[0058] At 406, the security service determines that a domain name included in the domain name resolution request is indicative of malicious activity. In some embodiments, the determining comprises determining whether the domain name is included in a list of known malicious or suspicious domains. The determining may also or instead comprise determining that the domain name is unfamiliar, is associated with a specific geographic location, or is associated with a specific entity.

[0059] At 408, the security service responds to the domain name resolution request with a network address of a monitored server to cause the requesting process to communicate with the monitored server in place of an adversary server. In some embodiments, the monitored server poses as an adversary

command-and-control system or an adversary exfiltration system to the requesting process, the requesting process being a attack. Also, the monitored server may decode the communications from the requesting process. Further, the monitored server may determine that the requesting process is utilizing a new protocol to encode the communications and analyze the communications to learn the new protocol.

[0060] At 410, the security service alerts at least one of a security agent or users of an entity affected by an attack associated with the requesting process.

[0061] At 412, the security service transitions an attack associated with the requesting process to an monitored device, the monitored device posing as the computing device impacted by the attack.

[0062] FIG. 5 illustrates an example process for monitoring dormant domains names associated with malicious activity and, in response to a change, responding with an alert or a configuration update. The process 500 includes, at 502, detecting suspicious domain names to add to a repository of dormant domain names by identifying whether the domain name resolves to local network address, whether the domain name resolves to non-routable network address, or whether the domain name fails to resolve and then ascertaining whether the domain name is associated with malicious activity. Dormant domain names associated with malicious activity are then added to the repository.

[0063] At 504, a security service associated with the repository monitors for changes to name resolutions or registrations of domain names included in the

repository. The monitoring may include continuously or periodically resolving the dormant domain names to determine if their name resolutions have changed.

[0064] At 506, the security service associates a monitored server with one of the dormant domain names and, in response to the dormant domain name becoming active, responds to domain name resolution requests specifying the now-active domain name with a network address of the monitored server.

[0065] At 508, the security service may perform at least one of updating a security agent configuration or alerting a security service user based on the detected changes.

Example System

[0066] FIG. 6 illustrates a component level view of a computing device capable of acting as a security service device, a client device, an monitored device, or an monitored server. As illustrated, computing device 600 comprises a system memory 602 storing modules and data 604. Also, computing device 600 includes processor(s) 606, a removable storage 608 and non-removable storage 610, input device(s) 612, output device(s) 614 and communication connections 616 for communicating with other computing devices 618.

[0067] In various embodiments, system memory 602 is volatile (such as RAM), non-volatile (such as ROM, flash memory, etc.) or some combination of the two. The modules or data 604 stored in the system memory 602 may

comprise methods, threads, processes, applications or any other sort of executable instructions, such as the instructions utilized to perform the operations of the client devices 104, security service 108, monitored devices 118/120, or monitored servers 128/130. The modules and data 604 may also include files and databases.

[0068] In some embodiments, the processor(s) 606 is a central processing unit (CPU), a graphics processing unit (GPU), or both CPU and GPU, or other processing unit or component known in the art.

[0069] Computing device 600 also includes additional data storage devices (removable and/or non-removable) such as, for example, magnetic disks, optical disks, or tape. Such additional storage is illustrated in FIG. 6 by removable storage 608 and non-removable storage 610. Tangible computer-readable media may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data. System memory 602, removable storage 608 and non-removable storage 610 are all examples of computer-readable storage media. Computer-readable storage media include, but are not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be

accessed by the computing device 600. Any such tangible computer-readable media may be part of the computing device 600.

[0070] Computing device 600 also has input device(s) 612, such as a keyboard, a mouse, a touch-sensitive display, voice input device, etc., and output device(s) 614 such as a display, speakers, a printer, etc. These devices are well known in the art and need not be discussed at length here.

[0071] Computing device 600 also contains communication connections 616 that allow the computing device 600 to communicate with other computing devices 618, such as others of the security service devices, the client devices, monitored devices, monitored servers, or adversary systems.

CONCLUSION

[0072] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described. Rather, the specific features and acts are disclosed as exemplary forms of implementing the claims.

CLAIMS

WHAT IS CLAIMED IS:

1. A computer-implemented method comprising:

receiving automated or manual notification of an attack;

transitioning the attack to a monitored computing process or device that is configured to pose to an adversary as a computing device impacted by the attack; and

enabling the adversary to obtain deceptive information from the monitored computing process or device.

2. The method of claim 1, wherein the attack is one of a spearfish email, a clickable link, a website, a drive by exploit, a QR code, a Near Field Communications (NFC) triggered link, a document, an executable, a removable drive, or an archive.

3. The method of claim 1, wherein receiving the automated or manual notification comprises retrieving the attack from an information sharing system or portal to which the attack has been submitted, and transitioning the attack comprises installing the attack on a monitored computing process or device.

4. The method of claim 1, wherein the information sharing system or portal receives or retrieves the attack a different entity than the entity performing the receiving, transitioning, and enabling.

5. The method of claim 1, wherein the notification is received from a person forwarding the attack or from a security agent or service monitoring the impacted computing device.

6. The method of claim 1, further comprising blocking processing of the attack on the impacted computing device.

7. The method of claim 1, further comprising processing the attack on the monitored computing process or device.

8. The method of claim 1, wherein the monitored computing process or device is a physical machine or virtual machine.

9. The method of claim 8, wherein the physical machine or virtual machine is configured with an image of the impacted computing device.

10. The method of claim 8, wherein the machine is configured with an image that includes at least one or more of a user name, a machine name, an operating system version, desktop screens, folder names, preloaded files or computer firmware versions associated with the impacted computing device.

11. The method of claim 1, wherein the monitored computing process or device is an isolated sandbox on the impacted computing device.

12. The method of claim 1, wherein the monitored computing process or device is located on a network of an entity associated with the impacted computing device or is implemented remotely by a security service.

13. The method of claim 12, wherein the monitored computing process or device is implemented remotely by the security service while being assigned a network address or domain associated with the entity.

14. The method of claim 1, wherein the monitored computing process or device mimics user activity.

15. The method of claim 1, further comprising monitoring commands and events issued by the adversary or monitoring network activity, file activity, process activity, execution activity, registry activity, operating system activity,

firmware updates, loaded drivers, or kernel extensions of the monitored computing process or device.

16. The method of claim 15, wherein the commands and events issued include keystrokes, mouse activity, or command line interface activity.

17. The method of claim 15, further comprising updating a configuration for a security agent of the impacted computing device or providing human-consumable intelligence based at least in part on intelligence derived from the monitoring.

18. The method of claim 1, further comprising enabling an entity associated with the impacted computing device to provide deceptive information for inclusion on the monitored computing process or device.

19. The method of claim 1, further comprising monitoring the attack on the monitored computing process or device and providing deceptive information to the adversary repeatedly over a time period.

20. The method of claim 1, further comprising correlating information obtained from monitoring attacks on multiple monitored computing processes or devices.

21. The method of claim 20, wherein the monitored computing processes or devices are operated by different entities.

22. The method of claim 1, further comprising intercepting and decoding communications associated with the attack .

23. The method of claim 1, further comprising receiving intelligence from the adversary about at least one of the adversary's tools, tactics, techniques, or procedures.

24. One or more tangible computer-readable media storing computer-executable instructions configured to program one or more computing devices to perform operations comprising:

configuring a document to report identifying information of a device opening the document to the one or more computing devices;

receiving a report associated with the document, the report including the identifying information;

determining, based on the identifying information that an unauthorized entity is in possession of the document; and

alerting an entity associated with the document of the possession by the unauthorized entity.

25. The one or more tangible computer-readable media of claim 24, wherein the configuring includes embedding executable instructions or a link in the document.

26. The one or more tangible computer-readable media of claim 25, wherein the link embedded in the document is configured to cause a request to a monitored network address.

27. The one or more tangible computer-readable media of claim 25, wherein the executable instructions embedded in the document are configured to open a connection to the one or more computing devices.

28. The one or more tangible computer-readable media of claim 25, wherein the executable instructions embedded in the document perform the reporting based at least in part on whether the identifying information differs from identifying information embedded in the document.

29. The one or more tangible computer-readable media of claim 24, wherein the identifying information comprises at least one of a network address, a geographic location, a universally unique identifier (UUID), domain information, or derived/upstream network data.

30. A computer-implemented method comprising:

receiving a domain name resolution request from a requesting process;

determining that a domain name included in the domain name resolution request is indicative of malicious activity; and

responding to the domain name resolution request with a network address of a monitored computing process or device to cause the requesting process to communicate with the monitored computing process or device in place of an adversary server.

31. The method of claim 30, wherein the receiving comprises redirecting the domain name resolution request from a domain name server associated with a device or entity implementing the requesting process.

32. The method of claim 30, wherein the determining comprises determining whether the domain name is included in a list of known malicious or suspicious domains.

33. The method of claim 30, wherein the determining comprises determining that the domain name is unfamiliar, is associated with a specific geographic location, or is associated with a specific entity.

34. The method of claim 30, wherein the monitored computing process or device poses as an adversary command-and-control system or an adversary exfiltration system to the requesting process, the requesting process being part of an attack, or intercepts communications of the requesting process to the adversary command-and-control system.

35. The method of claim 30, wherein the monitored computing process or device decodes the communications from the requesting process.

36. The method of claim 30, wherein the monitored computing process or device determines that the requesting process is utilizing a specific protocol to encode the communications and selects an appropriate communications protocol for decoding or attempts to learn the specific protocol.

37. The method of claim 30, further comprising alerting at least one of a security agent or users of an entity affected by an attack associated with the requesting process.

38. The method of claim 30, further comprising transitioning an attack associated with the requesting process to a monitored computing process or device, the monitored computing process or device posing as a computing device impacted by the attack.

39. A system comprising:

one or more processors;

a repository of dormant domain names associated with malicious activity;

a monitoring module operated by the one or more processors and configured to monitor for changes to name resolutions or registrations of domain names included in the repository; and

a response module operated by the one or more processors and configured to perform at least one of updating a security device or agent configuration or alerting a security service user based on the detected changes.

40. The system of claim 39, wherein the monitoring for changes includes continuously or periodically resolving the dormant domain names to determine if their name resolutions have changed.

41. The system of claim 39, further comprising a detection module to identify domain names that resolve to local network addresses, domain names that resolve to non-routable network addresses, or domain names that fail to resolve, to ascertain whether the domain names are associated with malicious activity, and to add the domain names to the repository based at least in part on the ascertaining.

42. The system of claim 39, further comprising a deception module to associate a monitored computing process or device with one of the dormant domain names and, in response to the dormant domain name becoming active, to respond to domain name resolution requests specifying the now-active domain name with a network address of the monitored computing process or device.