

5 **SOFTWARE VULNERABILITY EXPLOITATION SHIELD**

 CROSS-REFERENCE TO RELATED APPLICATIONS

 This application claims priority to and the benefit of U.S. Provisional
Application No. 60/785,723 filed March 24, 2006, and entitled SOFTWARE
10 VULNERABILITY EXPLOITATION SHIELD, which is incorporated herein by
reference in its entirety.

 BACKGROUND

 The Internet is a colossal computing network open to hundreds of millions of
15 computers and users. The Internet's community of users is constantly engaged in
research, innovation and dialog about computer security vulnerabilities.

 A vulnerability is a weakness in a piece of software that can be exploited by
("malicious code"). Some users ("malefactors") seek to use malicious code to exploit
vulnerabilities. But the crafting of malicious code is not easy. Usually, only a very
20 talented programmer is capable of writing malicious code that is effective at
exploiting a given vulnerability, and usually that programmer must toil tediously for
some time to create competent malicious code.

 Often, when a talented programmer creates an example of malicious code, he
or she posts it publicly on the Internet so that others including malefactors can
25 inspect, test, critique and use it. Within the community of Internet users, the examples
of malicious code effective at exploiting any specific vulnerability are usually few in
number. A single example of malicious code is commonly borrowed and used by
many different malefactors who incorporate it (e.g., "cut and paste" it) into a number
of diverse packages, such as an image file or a web site designed to perform a "drive-
30 by download". They may also bundle it with a number of different payloads, such as a
rootkit or a keystroke logger.

 In theory, many vulnerabilities exist inside most complex software products
such an operating system (e.g., Windows® published by Microsoft Corporation). A
given vulnerability, however, is unknown until a researcher discovers it. Nevertheless,
35 upon discovering a particular vulnerability, the natural inclination of a researcher is to
share the discovery publicly, either by talking about it or by attempting to exploit it in

5 a way that can be monitored publicly. Rarely is the discovery of a particular vulnerability kept a secret for long. News spreads and becomes available on Internet bulletin boards and chat rooms or from public monitoring of new exploits.

Typically, once the publisher of the affected software product learns of the vulnerability, they undertake to design, write, test and release a software “patch” to
10 close off this known vulnerability. The patching process, however, is fraught with delay and hence a “risk window” arises.

For example, as a patch is released, the publisher makes it available to computer owners via Internet downloading or other distribution so that they may install it on their computers. For owners of computers running vulnerable software,
15 the preparation, distribution, and installation of a patch can be a painfully slow and arduous process. For instance, time passes as the software publisher works on the patch. Then the installation of the patch may require manual action on the part of the owner, which can be a burden on an enterprise owning hundreds or thousands of computers. Further, owners such as enterprises often must test the patch before
20 installing it on many of their computers and sometimes a patch will conflict with desired functions of other applications. Additional time may also pass if a publisher discovers after it releases a patch that it is faulty and must be recalled, revised, and reissued.

After a vulnerability becomes known in the community, well-crafted
25 malicious code capable of infecting particular computers is developed and deployed before the software on them is patched. Malefactors seek to trespass through this risk window and some programmers endeavor to help them. Accordingly, once a vulnerability becomes known, talented programmers typically race to create and publish malicious code so as to maximize the time available within the risk window.

30 Until now, many types of security software such as antivirus software have attempted to stop a particular example of malicious code by looking for a sequence of bytes in the binary file (i.e., signature) that uniquely identifies the structure of the programming code in the example. Typically, such examination of files or other code occurs at the application layer in the Open Systems Interconnection (OSI) model.
35 More specifically, the scanning of viruses is performed on the files after being stored in the file system or while held in memory on the computing system. Although these

5 systems work well for most viruses, through the use of exploited vulnerabilities, malicious code can be entered into the targeted computer at deeper levels than the application layer. As such, the malefactors exploiting such software vulnerabilities can cause the malicious code to be executed before typical security software can detect it and take action.

10

BRIEF SUMMARY

The above-identified deficiencies and drawback of current software security systems are overcome through example embodiments of the present invention. For example, embodiments described herein provide for narrowing the risk window for
15 computers that are potential targets of malefactors by monitoring for malicious code at a transport layer. Note that this Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in
20 determining the scope of the claimed subject matter.

One embodiment provides a mechanism for minimizing the exploitation of vulnerabilities on software installed on a computing system by inspecting network traffic thereto and identifying malicious code before it can be executed and/or installed. At a transport layer (e.g., transmission communication protocol (TCP)
25 sockets layer), network traffic is monitored using a security component installed thereon. When a message destined for the computing system is received, data included in the message is compared with exploit evidence used to identify malicious code. The exploit evidence is provided to the security component by a security service that gathers information about the malicious code. Such exploit evidence may include
30 such things and a signature of the code, Internet Protocol (IP) address associated with the code, unique resource locator (URL) of a website known to provide malicious code, and other factors.

Based on the comparison of data in the message with the exploit evidence, rules are identified that instruct the security component to take action on the message
35 received. For example, if the comparison indicates that the message includes or potentially includes malicious code, the security component may block the message

5 received from entering the computing system, while allowing other traffic to flow thereto. Alternatively, or in addition, the security component may notify the user (e.g., through a user interface) about the potential risk of the message as including malicious code, thus allowing the user to take the appropriate action. In another embodiment, the security component can disable any harmful features of the
10 malicious code. If, on the other hand, the comparison at the transport layer of the message to the exploit evidence does not indicate that the message includes or is otherwise associated with malicious code, then the message may be allowed to pass to the computing system.

Additional features and advantages of the invention will be set forth in the
15 description which follows, and in part will be obvious from the description, or may be learned by the practice of the invention. The features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and
20 appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order to describe the manner in which the above-recited and other advantageous features of the invention can be obtained, a more particular description
25 of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the
30 accompanying drawings in which:

Figure 1 illustrates a computing network with a security service and security component used to minimize exploitation of software vulnerabilities in accordance with example embodiments; and

Figure 2 illustrates flow diagrams for implementing a method of reducing a
35 risk window for computers that are potential targets of malefactors in accordance with example embodiments.

5 DETAILED DESCRIPTION

The present invention extends to methods, systems, and computer program products for narrowing the risk window for computers that are potential targets of malefactors by monitoring for malicious code at a transport layer. The embodiments of the present invention may comprise a special purpose or general-purpose computer including various computer hardware or modules, as discussed in greater detail below.

As previously mentioned, embodiments of the present invention overcome the above identified drawbacks and deficiencies of security software systems by narrowing the risk window for computers that are potential targets of malefactors. Under this embodiment, the owners of target computers install Security Components of an ongoing Security Service, which monitors news, discussion, technical developments and other information on the Internet or elsewhere about new vulnerability discoveries and new examples of malicious code. When the Security Service finds an example of malicious code, it gathers information about the malicious code in order to appropriately identify it. For example, the Security Service may inspect its data structure seeking to ascertain a signature for it. In addition, the Security Service can determine the source of the malicious code by identifying its IP address, URL, or other electronic address. In fact, there may be numerous information that the Security Service can use to identify the malicious code.

Once information about the malicious code is gathered, this information is used to create security updates including exploit evidence and security commands that are then distributed to installations of the Security Component. Each Security Component uses the exploit evidence and security commands for inspecting Internet or network data traffic flowing to the computer at the transport layer (e.g., TCP socket level), looking for the messages that match the exploit evidence. By inspecting data at the transport level, the Security Component can spot malicious code before any application or other software on the computer can execute or install it. Further, the Security Component can monitor and identify malicious code designed to exploit software vulnerabilities of most any application installed on the computing system.

If the Security Component finds a data transmission that matches the exploit evidence, it can execute any of a number of protective measures. For example, it may block that particular transmission from entering the computer, while allowing other

5 transmissions to pass and/or it may inform the computer user about the data transmission so that the user can take appropriate action. It can also modify the transmission so that any harmful features are disabled.

As time passes after the Security Component is installed, it could become unwieldy. It could be looking for so many signatures, electronic addresses, or other
10 exploit evidence that it slows performance of the computer on which it is installed, and it could potentially conflict with desired functions and transmissions. Accordingly, to minimize this problem, other example embodiments described herein below provide various options to limit the monitoring.

Note that the above embodiments differ significantly from typical firewall
15 systems and Web content filters. For example, firewalls only perform packet inspection to look at the protocols, sources, and destination to make binary decisions of allowance or disallowance of such traffic. Similarly, Web content filters typically focus on white/black lists of IP address or URLs for allowing or disallowing packets, which embodiments herein also contemplate. Embodiments herein, however, inspect
20 the content of messages searching for malicious code targeted at vulnerabilities in software based on security evidence, commands, and rules generated from a Security Service that continually monitors for such malicious code. In other words, embodiments herein do not simply look at a source of a packet and make a binary decision based thereon, but rather provides a more robust system that continually
25 monitors formation of malicious code and then can create evidence, commands, and rules that can dynamically change and be implemented to appropriately handle such content in a variety of ways as described in greater detail below.

Although more specific reference to advantageous features are described in greater detail below with regards to the Figures, embodiments within the scope of the
30 present invention also include computer-readable media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk
35 storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to carry or store desired program code means in the form

5 of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a computer-readable medium. Thus, any such
10 connection is properly termed a computer-readable medium. Combinations of the above should also be included within the scope of computer-readable media.

Computer-executable instructions comprise, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions.
15 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 above. Rather, the specific features and acts described herein are disclosed as example forms of implementing the claims.

20 As used herein, the term "module" or "component" can refer to software objects or routines that execute on the computing system. The different components, modules, engines, and services described herein may be implemented as objects or processes that execute on the computing system (e.g., as separate threads). While the system and methods described herein are preferably implemented in software,
25 implementations in hardware or a combination of software and hardware are also possible and contemplated. In this description, a "computing entity" may be any computing system as previously defined herein, or any module or combination of modules running on a computing system.

Figure 1 illustrates a computing network 100 configured to overcome the
30 above identified drawbacks and deficiencies of current security software systems by narrowing the risk window for computers that are potential targets of malefactors. As shown, target computing systems have installed thereon special Security Components
105 associated with an ongoing Security Service 110. The Security Service 110 continually and/or periodically monitors such things as news 115, discussion 150,
35 technical developments or reports 120, and other information (such as network traffic 125) on the Internet or elsewhere about new vulnerability discoveries and new

5 examples of malicious code. Of course, the sources used to gather information about potential malicious code can vary widely and can include any well known form as indicated by the vertical ellipsis 124.

Further, this monitoring can take various forms, including manual searching and/or reading of Internet or other discussions 150 or automated collection of
10 intelligence such as through probes, honeypots, search bots, etc., that glean relevant statements and code from known or potential circles of malefactors or recorded examples of malefactor behavior, tools, experimentation, etc. Of course, like the sources of information, the mechanisms used for gathering information 155 about malicious code are used herein for illustrative purposes only and are not meant to
15 limit or otherwise narrow embodiments herein unless explicitly claimed.

As the Security Service 110 finds information about malicious code, it can store such data as malicious code data 155. As will be appreciated, the malicious code data 155 can take on many forms as it is gathered from various sources and mechanism. For example, the malicious code data 155 may include information about
20 a data structure 165 for the malicious code. More specifically, Security Service 110 can inspect the data structure of malicious code seeking to ascertain such things as a signature or other representation for it. Each example of malicious code may have a unique data structure 165 or data sequence within it. This unique data can then be used (as described below) for identifying examples of malicious code.

25 Note that the signature or representation that identifies the malicious code can take on many well known forms. For example, the signature could be either the entire code, or only a portion thereof (e.g., a sequence of bytes or binary string). Further, the signature could be a hash of the code or any portion thereof. In fact, malefactors often try to trick security systems by encrypting portions of the code, creating many
30 variations thereof, or using any number of other devious mechanisms. Accordingly, the term "signature" as used herein should be broadly construed to mean any unique data structure or identifier used to represent malicious code.

In another embodiment, the Security Service 110 also monitors for IP/URL addresses 175 associated with efforts to exploit known vulnerabilities. When it does
35 identify such an IP/URL address, these can also be added as the malicious code data 155, which as will be described below can be communicated to the Security

5 Components 195 for blocking data received from such sources (e.g., by blacklisting). Of course, there may be other 170 malicious code data 155 that can be gathered as will be described in greater detail below.

 The malicious code data 155 is used to create security updates 130, which the Security Service 110 distributes to installations of the Security Component 195 for
10 looking for the malicious code. As shown, the security updates 130 may include several pieces of information and code. For example, the security updates may include exploit evidence 135, which can be compared to messages or portions thereof as described in greater detail below. The exploit evidence 135 can include such things as signatures of data structures 165, IP/URL or other electronic addresses 175, and other
15 information 170 gathered from the Security Service 110. The security updates 130 may also include such things as security commands 145 and rules 140, which as described in greater detail below are used by the Security Component 195 to determine such things as to how to use the exploit evidence 135, and what actions to take once malicious code is identified.

20 The Security Component 195 receives the security updates 130 from the Security Service 110 and modifies its library 104 accordingly. The Security Component 195 then inspects Internet or network data traffic 125 flowing to the computer 105. Embodiments provide for the inspection of such traffic at the transport layer 190, which from an OSI model resides between the network layer 185 and the
25 application layer 122 that includes the file system layer 118. Although there are many other layers in the OSI model, and other abstract models that can be used represent communication channels, the ones shown here are used for simplicity reasons only. Nevertheless, there may be other layers and other models that can be used with embodiments described herein.

30 Examples of the transport layer 190 include a TCP socket level, which is the end-point, at the target computer 105, of a two-way communication link between two programs running on a TCP network. By inspecting data at the transport layer 190, the Security Component 195 can spot malicious code before any application or other software on the computer can execute or install it in memory or at the file system
35 layer 118. Further, because most, if not all, network traffic for applications flows through the transport layer 190, the Security Component 195 can monitor and identify

5 malicious code designed to exploit software vulnerabilities of most any application installed on the target computing system 105.

For example, as messages 180 are received at the transport layer 190, Security Component 195 uses its security comparator module 102, and optionally the security commands 112, to compare portions of the message 180 with the exploit evidence
10 106. If the comparator module 102 finds a data transmission matching the exploit evidence 106 or violating policies defined in the security commands 112, the Security Component 195 executes any of a number of protective measures as can be defined by the rules 108. For instance, if the security comparator module identifies message 180 as including a signature, electronic address, or other representations of malicious code
15 as defined by the exploit evidence 106, the Security Component 195 can take the appropriate action defined by rules 108.

In one embodiment, a rule 108 might instruct Security Component 195 to block that particular transmission or message 180 from entering the target computer 105, while allowing other benign transmissions to pass. An alternative embodiment
20 defines a rule 108 that allows the Security Component 195 to inform the computer 105 user about the data transmission so that the user can take appropriate action. Such notification may be in the form of a user interface that provides the user with information about the message (e.g., source, violation, potential for risk, etc.) and an option to take action (e.g., delete, allow, quarantine, hold, etc.). In still another
25 embodiment, a rule 108 may instruct the Security Component 195 to modify the transmission or message 180 so that any harmful features are disabled or otherwise eliminated. For example, the Security Component 195 may automatically remove the portion of the message that contains the malicious code or disable the harmful features within the code itself.

30 Of course, there may be any number of actions that can be taken in any combination as defined by the rules 108. For example, the rules 108 may first remove or disable the harmful features of a message 180 prior to informing the user about the transmission 108. Further, the rules 108 may be set or modified by the user of the target system 105, set or modified by the Security Service 110, be default
35 programming on the Security Component 195, or be set or modified using any well known manner. Accordingly, any specific use of an action or how that action is set or

5 applied is used herein for illustrative purposes only and is not meant to limit or otherwise narrow the scope of embodiments described unless otherwise explicitly claimed.

As mentioned above, as time passes after the Security Component 195 is installed, it could become unwieldy. It could be looking for so many signatures,
10 electronic addresses, and other exploit evidence 106 that it slows performance of the computer 105 on which it is installed, and it could potentially conflict with desired functions and transmissions. Accordingly, to minimize this problem, example embodiments described herein provide various options.

One example option provides for attending to only exploit evidence 106 for
15 malicious code that is especially worrisome. The Security Service 110 evaluates discovered examples of malicious code to determine how likely it is to yield a serious threat to targeted computers. This evaluation might include review of the nature and scope of the relevant vulnerability, the quality of the workmanship used to write the example Code, and so on. The Service 110 then makes and distributes threat ratings
20 160 within security updates 130 and/or any necessary commands 145 for looking for exploit evidence 106 of only those examples of code that the Security Service 110 judges to be most harmful. For example, a threshold threat rating 160 may be used to determine if exploit evidence 135 will or will not be included in the security update 130. Of course, there may be any number of ways of judging what exploit evidence
25 135, security commands 112, or rules 140 should be included or excluded from the security updates 130 based on such threat ratings 160.

Another embodiment provides for cancelling, deleting, temporarily pausing, or performing similar actions on the exploit evidence 106, security commands 112, and/or rules 108 on the Security Component 195 after the relevant risk window has
30 closed or if the threat rating 160 falls below some set threshold. For example, the risk window might close as a result of wide distribution of a competent and trusted patch. Or, it might close because experience reveals the relevant example of malicious code to be harmless. Accordingly, after the Security Service 110 publishes security updates 130 to installations of the Security Component 195, it continues to monitor the
35 subsistence of the relevant vulnerability and malicious code within the Internet community. When the Service 110 makes a judgment that the relevant risk window is

5 minor or nonexistent, the Service 110 issue a threat rating 160 that can be used to send commands 145 to Security Components 195 installations that they should no longer monitor for the exploit evidence 106. Of course, for some signatures, electronic addresses, or other malicious code data 155 the risk window may never close, in which case the Security Service 110 may elect never to cancel the
10 monitoring of that exploit evidence 106.

Of course, there may be many other mechanisms for enhancing the performance of the Security Component 195. For example, the Security Service 110 may send security commands 145 and rules 140 to the Security Component 195 that instruct it to delete, cancel, hold, or perform some other action on exploit evidence
15 106 and/or any security commands 112 or rules 108 associated therewith based on the occurrence of an event. As such, the Security Component 195 monitors the system for the occurrence of such events, and then takes the appropriate action as defined by the commands 145 and/or rules 140. Such events may be based on the expiration of time or on such things as the Security Component 195 determining that the threat rating
20 160 for the exploit evidence 106 or malicious code associated therewith is low. In fact, there may be any number of events that can be used to take action for enhancing the performance of the Security Component 195. Further, there may be any number of actions such as delete, cancel, pause, hold, etc., that may be performed upon the occurrence of the event. Accordingly, any specific event and/or action taken as
25 describe herein is used for illustrative purposes and is not meant to limit or otherwise narrow embodiments described unless otherwise explicitly claimed.

The following describes some specific examples of various implementations that can be used in accordance with example embodiments. These implementations are only examples of how the solution can be applied. They are not the only possible
30 examples or implementations and are therefore not meant to limit or otherwise narrow the scope of embodiments described herein unless explicitly claimed.

Security Component 195 is installed on a user's personal computer 105 running the Microsoft Windows XP® operating system. Security Service 110 is permitted to securely send updates 130 to the Security Component 195. When a
35 researcher announces on the Internet discovery of a previously unknown vulnerability in Microsoft Windows XP®, Security Service 110 reads discussions about it among

5 programmers/hackers on the Internet. The Service 110 observes that one programmer posts code for exploiting the vulnerability. Accordingly, the Security Service 110 tests the code and determines that it appears likely to be a threat to computers running Windows.

10 From the code, the Security Service 110 ascertains a signature, i.e., a unique data structure that identifies the code. The Security Service 110 sends exploit evidence 135, rules 140, and commands 145 to the Security Component 195 installed on the user's computer to monitor for the signature. The Security Component 195 searches for the signature in the Internet data traffic sent to the user's computer. Searching at the transport layer 190 (e.g. TCP socket level) of the data traffic, the
15 Component 195 looks for the signature regardless of the data package it may be in or the payload it may carry. The Security Component 195 finds the signature, combined with a rootkit payload, in incoming e-mail data traffic 180. The Security Component 195 blocks the unit of data including the signature. The data is blocked before any software on the user's computer is able to execute or install the rootkit.

20 By way of another example, a user installs Security Component 195 on her personal computer running the Mozilla Firefox® Internet Browser version 1.5.0.1 for Microsoft Windows® (the "Browser"). Security Service 110 is permitted to securely send updates 130 to the Security Component 195. Service 110 plants and maintains various automated honeypots, probes, and other computers on the Internet to collect
25 examples of malicious code and learn the source and other information about such code. Some of the probes, installed with the Browser, automatically visit many web sites in search of malicious code.

Shortly after a researcher announces discovery of a certain vulnerability in the Browser, a server at a particular IP address (the "Rogue Address") sends to several of
30 the Security Service's probes a Trojan bundled with malicious code that exploits the vulnerability. In response, the Security Service 110 sends security updates 130, which may include exploit evidence 135 and commands 145 to the Security Component 195 installed on the user's computer 105 to monitor for the Rogue Address as the source of data being transmitted to the user's computer 105. The Security Component 195
35 searches, at the TCP socket level 190, for the Rogue Address in the Internet data traffic 180 sent to target computer 105. The publisher of the Browser issues a patch to

5 close the relevant vulnerability. The Security Service 110 determines that the Rogue Address is no longer a threat to a computer running the Browser and transmits security updates 130 to the Security Component 195 to cease searching for the Rogue Address.

10 The present invention may also be described in terms of methods comprising functional steps and/or non-functional acts. The following is a description of steps and/or acts that may be performed in practicing the present invention. Usually, functional steps describe the invention in terms of results that are accomplished, whereas non-functional acts describe more specific actions for achieving a particular result. Although the functional steps and/or non-functional acts may be described or
15 claimed in a particular order, the present invention is not necessarily limited to any particular ordering or combination of steps and/or acts. Further, the use of steps and/or acts in the recitation of the claims—and in the following description of the flow diagram for Figure 2—is used to indicate the desired specific use of such terms.

20 As previously mentioned, Figure 2 illustrates a flow diagram for various exemplary embodiments of the present invention. The following description of Figure 2 will occasionally refer to corresponding elements from Figure 1. Although reference may be made to a specific element from this Figure, such references are used for illustrative purposes only and are not meant to limit or otherwise narrow the scope of the described embodiments unless explicitly claimed.

25 Figure 2 illustrates a flow diagram of a method 230 for narrowing the risk window for computers that are potential targets of malefactors. On the Security Service 200, method 230 includes an act of gathering 205 information about example malicious code. For instance, Security Service 110 may monitor and gather news 115, discussions 150, technical developments 120, network traffic 125, or other
30 information from the Internet or other areas about new (or old) vulnerability discoveries and/or new (or old) examples of malicious code. As previously mentioned, this gathering of information can take on many forms, including manual searching and reading of Internet discussions or automated collection of intelligence such as through probes, honeypots, search bots, etc. that glean relevant statements and
35 software from known or suspected malefactors or recorded examples of malefactor behavior, tools, or experimentation.

5 Based on the information gathered, method 230 further includes an act of generating 210 security updates. For example, Security Service 110 may use the malicious code data 155 to generate security updates 130. Such security updates 130 comprise exploit evidence 135, rules 140, and/or security commands 145, which are used to identify such things as signatures, electronic addresses, or other information
10 about malicious code and take appropriate action thereon.

 Method 230 also includes an act of sending 215 the security updates to the security subscribers. On the Security Component side 250, method 230 further includes an act of receiving 220 the security updates. Based on the security updates, method 230 includes an act of modifying 225 security component 250. For example,
15 Security Service 110 sends the security updates 130 to the target computing system 105, which updates the security component 195 library 104 with the appropriate exploit evidence 106, rules 108, and security commands 112.

 Method 230 further includes an act of monitoring 235 network traffic at a transport layer. For example, before, after, or during, the receipt of the security
20 updates 130, Security Component 195 can monitor network traffic 125 at transport layer 190 (e.g., TCP socket level), which is the end-point, at the target computer 105, of a two-way communication link between two programs running on the network. By inspecting data at the transport layer 190, the Security Component 195 can spot malicious code before any application or other software on the computer can execute
25 or install it in memory or at the file system layer 118. Further, because most, if not all, network traffic for applications flows through the transport layer 190, the Security Component 195 can monitor and identify malicious code designed to exploit software vulnerabilities of most any application installed on the target computing system 105.

 Method 230 also includes a step for identifying 240 message(s) with malicious
30 code based on exploit evidence from security updates. For example, Security Component 195 may perform an act of receiving message 180 at the transport layer 190 and use security comparison module 102 for an act of comparing data within the message 180 with exploit evidence 106. Based on the comparison, Method 230 then includes a step for taking action 245 based on a set of rules. For example, if Security
35 Component 195 performs an act of identifying that the message 180 corresponds to malicious code, then a rule 108 might instruct it to block that particular transmission

5 or message 180 from entering the target computer 105. Even though the rule 180 blocks such message 180 it may allow other benign transmissions to pass. An alternative embodiment provides a rule 108 that allows the Security Component 195 to inform the computer 105 user about the data transmission so that the user can take appropriate action. In still another embodiment, a rule 108 may instruct the Security
10 Component 195 to modify the transmission or message 180 so that any harmful features are disabled or otherwise eliminated.

In order to enhance performance of the Security Component 195 other embodiments are needed. Accordingly, one embodiment for attending to only exploit evidence 106 for malicious code that is above a threat ratings 160, which is a
15 determination of how harmful a particular malicious code might be. Another embodiment provides for cancelling, deleting, holding, or other action to the exploit evidence 106, security commands 112, and/or rules 108 on the Security Component 195 after the relevant risk window has closed or if the threat rating 160 falls below some set threshold.

20 In still another embodiment, the Security Component 195 may receive a command 145 and/or rule 140 that instruct it to delete, cancel, or take some other action on exploit evidence 106 and/or any security commands or rules 108 associated therewith based on the occurrence of an event. As such, upon the occurrence of the desired event, the Security Component 195 can take an appropriate action on the
25 exploit evidence 106, security commands 112 and rules 108 associated therewith such as deleting them, canceling them, temporally pause them, etc.

The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the
30 invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

5 1. At a computing system in a network computing environment, a method of minimizing the exploitation of vulnerabilities on software installed on the computing system by inspecting network traffic thereto and identifying the malicious code before it can be executed and/or installed, the method comprising:

 monitoring, at a transport layer, incoming network traffic of a computing
10 system using a security component installed thereon;

 receiving as part of the network traffic a message at the transport layer identified as destined for the computing system;

 comparing at least a portion of data included in the message received with exploit evidence used to identify malicious code, the exploit evidence provided to the
15 security component by a security service that gathers information about the malicious code;

 based on the comparison with the exploit evidence, identifying one or more rules that instruct the security component to perform one or more actions on the message received.

20 2. The method of claim 1, wherein the transport layer is one or more of a TCP socket level.

 3. The method of claim 1, wherein the comparison identifies the at least a portion of data as corresponding to malicious code, and the one or more rules instruct the security component to do one or more of the following:

25 block the message received from entering the computing system;

 inform a user of the computing system about the correspondence of the message to the malicious code in order to allow the user to take appropriate action; or

 modify the message in order to disable any harmful features of the malicious code.

30 4. The method of claim 3, wherein the one or more rules allow other benign messages to pass to the computing system, while blocking the message received.

 5. The method of claim 3, wherein the one or more rules inform the user of the computing system about the correspondence of the message using a user
35 interface and allowing the user to either accept or reject the message.

5 6. The method of claim 1, wherein the exploit evidence includes a list of known electronic addresses associated with malicious code.

 7. The method of claim 6, wherein the electronic address list includes IP addresses or URLs for websites and wherein an IP address or a URL for a source of the message is compared to the list of known electronic address associated with
10 malicious code.

 8. The method of claim 1, wherein the exploit evidence includes one or more signatures of the malicious code, which are unique data structures that represent the malicious code.

 9. The method of claim 1, further comprising:
15 receiving from the security service a command that indicates that exploit evidence used to identify malicious code should expire based on one or more events; and

 upon occurrence of the one or more events, taking action on the exploit evidence as defined by one or more rules.

20 10. The method of claim 9, wherein the event is one or more of determining that a risk window has passed or that the malicious code is below some threat rating threshold, and wherein the action on the exploit evidence includes one or more of cancelling the comparing of the exploit evidence, deleting the exploit evidence from the security component, or temporarily pausing the comparing of the
25 exploit evidence.

 11. At a computing system in a network computing environment, a computer program product comprising a computer-readable storage medium having encoded thereon computer-readable instructions, the instructions, when executed in a computing environment, perform a method comprising:

30 monitoring, at a transport layer, incoming network traffic of a computing system using a security component installed thereon;

 receiving as part of the network traffic a message at the transport layer identified as destined for the computing system;

 comparing at least a portion of data included in the message received with
35 exploit evidence used to identify malicious code, the exploit evidence provided to the

5 security component by a security service that gathers information about the malicious code;

based on the comparison with the exploit evidence, identifying one or more rules that instruct the security component to perform one or more actions on the message received.

10 12. The computer program product of claim 11, wherein the transport layer is one or more of a TCP socket level.

13. The computer program product of claim 11, wherein the comparison identifies the at least a portion of data as corresponding to malicious code, and the one or more rules instruct the security component to do one or more of the following:

15 block the message received from entering the computing system;

inform a user of the computing system about the correspondence of the message to the malicious code in order to allow the user to take appropriate action; or

modify the message in order to disable any harmful features of the malicious code.

20 14. The computer program product of claim 13, wherein the one or more rules allow other benign messages to pass to the computing system, while blocking the message received.

15. The computer program product of claim 13, wherein the one or more rules inform the user of the computing system about the correspondence of the message using a user interface and allowing the user to either accept or reject the message.

16. The computer program product of claim 11, wherein the exploit evidence includes a list of known electronic addresses associated with malicious code.

17. The computer program product of claim 16, wherein the electronic address list includes IP addresses or URLs for websites and wherein an IP address or a URL for a source of the message is compared to the list of known electronic address associated with malicious code.

18. The computer program product of claim 11, wherein the exploit evidence includes one or more signatures of the malicious code, which are unique data structures that represent the malicious code.

5 19. The computer program product of claim 11, further comprising:
 receiving from the security service a command that indicates that exploit
evidence used to identify malicious code should expire based on one or more events;
and

 upon occurrence of the one or more events, taking action on the exploit
10 evidence as defined by one or more rules.

 20. The computer program product of claim 19, wherein the event is one
or more of determining that a risk window has passed or that the malicious code is
below some threat rating threshold, and wherein the action on the exploit evidence
includes one or more of cancelling the comparing of the exploit evidence, deleting the
15 exploit evidence from the security component, or temporarily pausing the comparing
of the exploit evidence.