

Tracer: Enforcing Mandatory Access Control in Commodity OS with the Support of Light-Weight Intrusion Detection and Tracing

Zhiyong Shan[†]
zhiyongshan@gmail.com

Xin Wang[‡]
xwang@ece.sunysb.edu

Tzi-cker Chiueh^{*†}
chiueh@cs.sunysb.edu

[†]Key Laboratory of Data Engineering and Knowledge Engineering (Renmin University of China), MOE

[‡]Stony Brook University

^{*}Industrial Technology Research Institute

ABSTRACT

Enforcing a practical Mandatory Access Control (MAC) in a commercial operating system to tackle malware problem is a grand challenge but also a promising approach. The firmest barriers to apply MAC to defeat malware programs are the incompatible and unusable problems in existing MAC systems. To address these issues, we start our work by analyzing the technical details of 2,600 malware samples one by one and performing experiments over two types of MAC enforced operating systems. Based on the preliminary studies, we design a novel MAC model incorporating intrusion detection and tracing in a commercial operating system, named Tracer, in order to disable malware on hosts while offering good compatibility to existing software and good usability to common users who are not system experts. The model conceptually consists of three actions: detecting, tracing and restricting suspected intruders. One novelty is that it leverages light-weight intrusion detection and tracing techniques to automate security label configuration that is widely acknowledged as a tough issue when applying a MAC system in practice. The other is that, rather than restricting information flow as a traditional MAC does, it traces intruders and restricts only their critical malware behaviors, where intruders represent processes and executables that are potential agents of a remote attacker. Our prototyping and experiments on Windows show that Tracer can effectively defeat all malware samples tested via blocking malware behaviors while not causing a significant compatibility problem.

Categories and Subject Descriptors

D.4.6 [Operating System]: Security and Protection—*Invasive software*

General Terms

Security

Keywords

Access Control; Operating System; Compatibility; Usability; Malware; Intrusion Detection;

1. INTRODUCTION

Malicious software (i.e., Malware) has resulted in one of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASIACCS '11, March 22–24, 2011, Hong Kong, China.

Copyright 2011 ACM 978-1-4503-0564-8/11/03...\$10.00.

most severe computer security problems today. A network of hosts which are compromised by malware and controlled by attackers can cause a lot of damages to information systems. As a useful malware defense technology, Mandatory Access Control (MAC) works without relying on malware signatures and blocks malware behaviors before they cause security damage. Even if an intruder manages to breach other layers of defense, MAC is able to act as the last shelter to prevent the entire host from being compromised. However, as widely accepted [2][3][5], existing MAC mechanisms built in commercial operating systems (OS) often suffer from two problems which make general users reluctant to assume them. One problem is that a built-in MAC is incompatible with a lot of application software and thus interferes with their running [2][3][5], and the other problem is low usability, which makes it difficult to configure MAC properly [2]. Thus, enforcing a practical MAC on commercial OS to defend against malware is a promising but challenging task.

In order to devise a new form of MAC to defeat malware, we have performed two preliminary studies. First, we analyzed the technical details of 2,600 samples so as to get a deep and overall view on malware programs. We extracted 30 critical malware behaviors and found three common malware characteristics that can guide anti-malware system design. Second, we investigated the root cause of incompatibility and low usability of existing MAC models through experiments on two types of MAC enforced operating systems. Our observations are as follows. The incompatibility problem is introduced because the security labels of existing MACs are unable to distinguish between malicious and benign entities, which causes a huge number of false positives (i.e. treating benign operations as malicious) thus preventing many benign software from performing legal operations; the low-usability problem is introduced, because existing MACs are unable to automatically label the huge number of entities in OS and thus require tough configuration work at end users.

With these investigation results, we propose a novel MAC model, Tracer, which consists of three actions: detection, tracing and restriction. Each process or executable has two states, suspicious or benign. An executable in this paper represents an executable file with a specific extension, such as .EXE, .COM, .DLL, .SYS, .VBS, .JS, .BAT, or a special type of data file that can contain macro codes, say a semi-executable, such as .DOC, .PPT, .XLS, and .DOT. The actions of detection and tracing change the state of a process or executable to suspicious if it is suspected to be malicious, and the entity marked as suspicious is called a suspicious intruder. The action of restriction forbids a suspicious intruder to perform malware behaviors in order to maintain confidentiality, integrity and availability of the system, as well as to stop malware propagation.

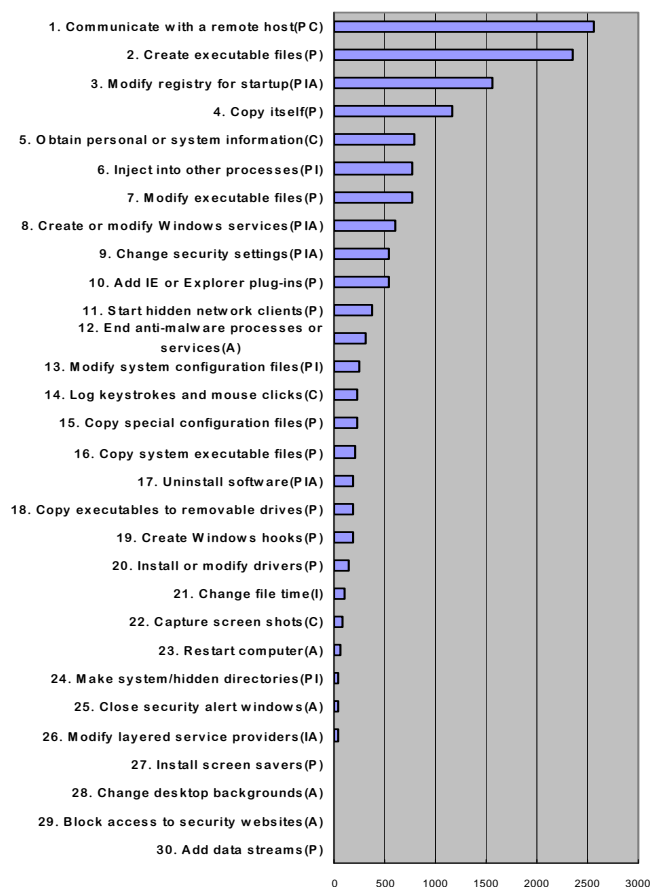


Figure 1. The top 30 Critical malware behaviors

To be precise, once detecting a suspicious process or executable, Tracer labels it to be suspicious and traces its descendent and interacted processes, as well as the executables it generates. Tracer does not restrict any operations of benign processes. Meanwhile, it permits suspicious processes to run as long as possible but only forbids their malware behaviors.

The novelty of Tracer is that, it incorporates light-weight intrusion detection and tracing techniques for configuring security labels, i.e., labeling suspicious OS entities, which is often done manually. Moreover, rather than restricting information flow as a traditional MAC does, it traces suspected intruders and restricts the malware behaviors of suspected intruders, i.e., processes and executables that are potential agents of remote attackers. These novelties lead to two advantages. First, Tracer is able to better identify potentially malicious OS entities and regulate their behaviors, which in turn significantly reduces the false positive (FP) rate which is the root cause of incompatibility in existing MAC-enforced systems. Second, Tracer is able to label OS entities automatically to tackle the low usability problem which is the other major issue of existing MAC systems [2].

We have implemented Tracer on Windows and have been using evolving prototypes of the Tracer system in our lab for a few months. Our experiments on the function of Tracer with a set of real-world malware samples demonstrate that it can effectively block malware behaviors while offering good compatibility to applications and good usability to normal users. Moreover, another experiment shows that Tracer causes much fewer FPs than commercial anti-malware tools and MIC (Mandatory Integrity Control) which is a MAC mechanism on Windows Vista

[4][16]. The contributions of this paper are as follows:

1. We introduce Tracer, a novel MAC mechanism which integrates intrusion detection and tracing techniques to disable malware on a commercial OS in a compatible and usable manner.
2. We have implemented Tracer model on Windows OS to disable malware timely without need of malware signatures. Developing a prototype on Windows is important, because most of the over 236,000 known malware items are designed for the attacks in the Windows environment, only about 700 malware items target for the attack of various Unix/Linux distributions [12].
3. Based on the analysis of 2,600 malware samples, we extract 30 critical malware behaviors and summarize three useful malware characteristics, which will benefit future anti-malware researches.
4. We investigate the root reasons of incompatibility and low usability problems of existing MACs. Although not all the observations are brand new, we believe that understanding these reasons more comprehensively and illustrating them through the design of an actual system are useful for other MAC researchers.

The rest of the paper is organized as follows. Section 2 introduces in details our investigation on various behaviors of malware programs, and our analysis on existing problems in MAC. Section 3 describes Tracer model. Section 4 provides our prototype and tests of Tracer on Windows. Lastly, we present the related research in Section 5 and conclude the work in Section 6.

2. PRELIMINARY STUDIES

2.1 Malware Investigation

Malware contribute to most Internet security problems. Anti-malware companies typically receive thousands of new malware samples every day. An analyst generally attempts to understand the actions that each sample can perform, determines the type and severity of the threat that the sample constitutes, and then forms detection signatures and creates removal procedures. Symantec Threat Explorer [6] is such a publicly available database which stores the analysis results of thousands of malware samples from various sources and is thus valuable to malware researchers. To have a thorough understanding of the philosophies behind malware design, we have spent considerable amount of time analyzing the behaviors of malware programs. Specifically, since 2008, we have read, recorded and analyzed the technical details of 2,600 malware samples of a wide range of formats and varieties, such as viruses, worms, backdoors, rootkits, and Trojan horses. As taking many samples from the same malware family might make the analysis results biased, we have intentionally not chosen multiple samples of a polymorphic malware or similar malware.

Figure 1 depicts the top 30 critical malware behaviors extracted from the samples and ranked in the descending order of their appearance times. For the behavior repeatedly appearing in a single malware, we only count it once. As the analysis is made on a great number of malware samples, we expect the behaviors captured to reflect the popular attacking techniques taken by the community of malware writers. Our performance studies in Section 4.2 have demonstrated that these behaviors are helpful to defend against unknown malware programs. The top 10 behaviors on the list are explained as follows:

1. **Communicate with a remote host.** Most malware samples need to access the network in order to break into hosts, download files, receive hacker commands, upload confidential information, and spread themselves to other hosts and so on. The

most commonly used network protocols are HTTP, POP3, IRC, SMTP, FTP and ICMP. Here, we not only take into account the network communications when a malware sample executes on a host but also the network communications through which a malware sample breaks into a host. **2. Create executable files.** When reaching a system, almost all malware immediately create their executables in various formats, including both binary and script files. A significant part of the created executables comes from the network. **3. Modify registry for startup.** 1,568 malware samples change or create registry entries in order to launch themselves upon restart of the system or applications. The registry entries changed the most are “run”, “Explorer” and “Winlogon” under the path “HKEY_LOCALMACHINE\SOFTWARE\Microsoft\Windows\C-urrentVersion\”. **4. Copy itself.** As a characteristic behavior of malware, more than 1000 malware samples copy themselves to system directories. **5. Obtain personal or system information.** 785 malware samples steal user or system information including password, email address, encryption key, IRC accounts, system version, configuration data, banking data, etc. Particularly, more than half of these malware samples are only interested in system information. **6. Inject into other processes.** 775 malware samples inject malicious codes into other process’ address space and remotely start a thread from the code. The system processes often being injected into include explorer.exe, winlogon.exe, svchost.exe, services.exe. **7. Modify executable files.** 769 malware samples insert codes into executables, particularly system dynamic library files, for propagating, hiding or starting up themselves. **8. Create or modify Windows services.** To start or hide themselves, 608 malware samples create or modify services by calling service API functions or directly changing service related registry entries. **9. Change security settings.** To escape from the checking by security software, 550 malware samples change registry to lower the system security level, disable firewall, bypass firewall, disable system restoring mechanism, alter safe mode setting, as well as disable task manager and registry editor. **10. Add IE or Explorer plug-ins.** 541 malware samples add plug-in into IE and Explorer, for example, Browser Helper Object on IE, Shell Service Object on Explorer.

Moreover, from the details of 2,600 malware samples, we discovered three common characteristics of malware that can guide our subsequent anti-malware design:

(1) **Entrance-Characteristics.** All malware samples break into hosts through two entrances, network and removable drive. Most breaking-ins are via network, mainly through **dangerous protocols** including HTTP, POP3, IRC, SMTP, FTP and ICMP, etc. Using these protocols, malware samples can penetrate the network and host firewalls by disguising themselves as popular software conducting benign network traffic.

(2) **Damage-Characteristics.** Malware behaviors can impose multiple forms of damages, i.e., resulting in problems in confidentiality, integrity and availability. Besides, we consider malware propagation as another type of damage since it can indirectly cause the former three forms of damages and eventually lead the entire host to be taken over. For example, the behavior “Copy itself” does not directly hurt security but is an essential step towards propagating itself and then executing malicious behaviors on a host. Therefore, we evaluate the damages of each behavior and record them in Figure 1, using C, I, A, and P to represent the damages related to confidentiality, integrity, availability and propagation respectively.

(3) **Attack-Characteristics.** Malware samples from the network have two attack patterns. One is that, most malware

samples exploit bugs in network-facing daemon programs or client programs to compromise them, then immediately spawn a shell or back-door process. Next, an attacker typically tries to download and install attacking tools and rootkits, as well as performs some other adversary behaviors. The other attack pattern is that, malware samples increasingly use social engineering methods to lure users into downloading and launching them. After started, a malware sample usually copies itself and makes itself a resident in a host.

2.2 Problems in MAC

Incompatibility is a well-known problem when enforcing a MAC model in a commercial operating system [2][3][5]. To investigate its root reason, in a secure network environment, we set up two machines to run MAC enforced operating systems including SELinux [14] with MLS policy enabled and RSBAC [15] with MAC module enabled. After a few days, we observed that these MAC systems produced a huge number of log records about denied accesses, which indicated that some applications failed and some acted abnormally. As the operation environment is secure without intrusion and malware, these denied accesses are thus “false positive”. In other words, MAC systems consider benign accesses malicious and refuse them. Many FPs together could make the whole system finally unusable. Although part of the FPs can be removed by experts through fine-granular policy configuration, many of them are not removable, and thus the impacted applications need to be modified before running on the MAC enabled systems.

These unremovable FPs are resulted because most MAC models aim to forbid illegal information flow rather than forbid intrusive behaviors directly. An example of such FPs is the self-revocation problem [3] in Low-Water-Mark model, which forbids a process to write a file created by itself if it has read a file with a lower integrity level before the writing. From the perspective of stopping illegal information flow, forbidding the write operation is reasonable. However, from the perspective of stopping intrusion, the write operation should not be denied if the process is actually not serving for an attacker. Another example of such FPs on a BLP-enforced Unix/Linux stems from the access control of the directory “/tmp” shared by the entire system [17]. To prevent illegal information flow, a process with a lower sensitive level can not read from /tmp or a process with a higher sensitive level can not write to /tmp. However, from the view of intrusion prevention, these processes do not necessarily represent intruders so that their “read” or “write” accesses to the /tmp should not be simply denied. Although it is possible to resolve this problem by adding “hiding sub directories” under /tmp, it is still difficult to eliminate the FPs resulting from many other shared entities on an OS, such as shared files, devices, pipes and memories.

Meanwhile, the security labels of MAC models also do not suit for fighting against malware, as they are designed to represent information integrity level or confidentiality level but not to distinguish between malicious and benign entities. In fact, a lower integrity level alone can not indicate that a process is malicious, as “malicious” also has other meanings, e.g., lower confidentiality and the risk of damaging system availability. Similarly, a lower confidentiality level alone cannot indicate that a process is malicious. Moreover, MAC labels are defined before an intrusion happens and can not be changed dynamically to reflect intrusion propagation in an OS. Although some of the MAC models are able to adjust label states, e.g. LOMAC [3] and DTE [9], they are still not flexible enough to track the intrusion propagation at the whole system level. Consequently, MAC labels can not differentiate

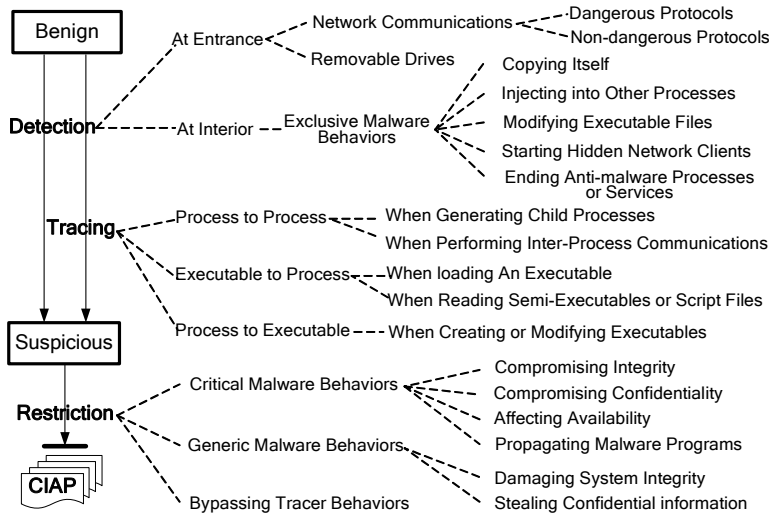


Figure 2. Tracer model

between malicious and benign entities. Relying on these labels, a MAC system often fails to make correct decisions on intrusion blocking which eventually results in many FPs.

Low-usability is another problem in a MAC-enabled system, as it often requires complicated configurations and unconventional ways of usage. In a modern OS, there are a wide range of entities including processes, files, directories, devices, pipes, signals, shared memories and sockets, etc. If just considering the files, there are more than 100,000 files on a typical Windows XP or Linux desktop. Moreover, MAC systems have complex policy interfaces which are difficult to configure. For instance, SELinux has 29 different classes of objects, hundreds of possible operations, and thousands of policy rules for a typical system. Hence, it is cumbersome for a common user to correctly configure labels for all entities without leaving security vulnerabilities. In addition, after enforcing a MAC, users must break their usage convention and learn how to use the MAC. Consequently, the ideal way for MAC to provide good usability is to automatically initialize and change entity labels without changing users' usage convention or requiring extra knowledge.

3. TRACER MODEL

In this section, we present our Tracer model that aims to disable malware in a commodity OS by disallowing malware behaviors. The adversaries of Tracer are malware programs that break into a host through the network and removable drives. As Windows is the most popularly used OS and attractive to hackers, the description of Tracer is based on our design for Windows. We believe the model can also be applied to other operating systems (e.g. Linux) with some changes. Investigating the suitability of Tracer for non-Windows operating systems is beyond the scope of this paper.

3.1 Model overview

The design of an access control model needs to answer three questions. The first is how to define the security label. Based on the analysis in Section 2.2, we introduce a new form of security label called suspicious label for our Tracer model. It has two values: suspicious and benign. A suspicious label indicates that the associated process is potentially serving for an intrusion purpose and thus possible to initiate some malicious behaviors. Meanwhile, Tracer only assigns a suspicious label to a process or an executable,

because a process is possibly the agent of an intruder and an executable determines the execution flow of a process which represents an intruder. All other entities in OS, e.g. non-executables, inter-process communication objects, registry entries, etc. do not need suspicious labels. When a process requests to access these entities, Tracer mainly utilizes their DAC information to make access control decisions, thus a huge amount of configuration work can be reduced while keeping traditional usage conventions unchanged. The second design question is how to configure security labels. As discussed in Section 2.2, in order to achieve good usability, a MAC model must have the capability of automatically deploying security labels. Accordingly, we introduce two types of actions named "detection" and "tracing" to automate the security label deployment progress. The two actions employ intrusion detection and tracing techniques respectively to recognize and mark suspicious processes and executables. The third design question is about access control rules. Being of mandatory access control, the rules of Tracer forcibly refuse the requests from suspicious processes if and only if the requests could lead to a malware behavior.

Figure 2 gives an overview of Tracer which consists of three types of actions, detection, tracing and restriction. Each process or executable has two states, suspicious and benign. The actions of detection and tracing change the state of a process or executable to suspicious if it is identified as a potential intruder. The restriction action forbids a suspected intruder to perform malware behaviors in order to protect CIAP. That is to protect confidentiality, integrity and availability, as well as to stop malware propagation. The three actions work as follows. Once detecting a suspected process or executable, Tracer labels it as suspicious and traces its descendent and interacted processes, as well as its generated executables. Tracer does not restrict benign processes at all, and permits suspicious processes to run as long as possible but stops their malware behaviors that would cause security damages. In addition, Tracer also provides a special system call to allow a user to change the state of a suspicious process or executable back to benign if the user trusts it. In the rest of this section, we describe Tracer model in details, including detecting, tracing and restricting intruders.

3.2 Detecting Intruders

The detecting action is responsible for identifying all potential intruders. We do not intend to design a complex intrusion detection algorithm to achieve a low FP rate at the cost of heavy overhead. Instead, we design a light-weight intrusion detection algorithm that can identify all potential intruders but may have a relatively higher FP rate at the initial step. However, even if the detecting action wrongly denotes a benign process as suspicious, the subsequent actions of Tracer, i.e., tracing and restricting actions, will still allow it to run rather than stop it immediately, but only prevent it from executing featured malware behaviors. In other words, Tracer is built to have a good tolerance to the FPs caused by the detecting action.

As depicted in Figure 2, the detection works at two levels: entrance and interior. The detection at entrance attempts to check all possible venues through which a malware program may break into the system. Network communications is the main type of entrances and most malware programs exploit several common dangerous protocols to compromise hosts as these protocols are usually permitted by firewalls according to the Entrance-Characteristics presented in Section 2.1. Hence, we denote a process as suspicious if it receives network traffic

through dangerous protocols. A frequently used application (e.g., web browser) thus might be denoted as suspicious but its normal running will not be affected, because Tracer doesn't restrict the suspected processes instantly and permits them to perform as long as possible except stopping their critical malware behaviors. The rest of the network protocols, say non-dangerous protocols, are difficult to be exploited by malware programs, because they are not permitted by firewalls since benign software rarely use them. Nevertheless, in order to completely monitor all the network traffic, we denote a process as suspicious if it receives network traffic through a non-dangerous protocol and then exhibits any of the malware behaviors. Instead of only checking non-dangerous network protocols, further checking malware behaviors can reduce the extra high FP rate. The Attack-Characteristics summarized in Section 2.1 supports this point. That is, a process exploited by a malware program from the network necessarily executes at least one critical malware behavior, e.g., launching a shell process or downloading an executable, to propagate the malware program within the system. Although a carefully crafted malware program that subverts a process through a non-dangerous protocol can perform some behaviors before performing a malware behavior, it is difficult for the process to make significant damages on the system. The reason is that the malware behaviors monitored by Tracer include all of the behaviors that can cause significant damages, let alone that malware programs are difficult to attack a host through non-dangerous protocols which are usually blocked by firewalls. The other type of entrances through which malware programs get into the system is removable drives according to the Entrance-Characteristics, hence we denote a process as suspicious when it opens or loads an executable from a removable drive.

With these detecting approaches enforced, however, two types of system maintenance tasks, i.e., updating software through the network and installing software from a removable drive, can not be performed because the processes that perform these tasks are treated as suspicious. As presented in the literature work [2] [3], a MAC policy should have ways to specify exceptions since no simple policy model can capture all accesses that need to be allowed and at the same time forbid all illegal accesses. Hence, we provide two means to facilitate these system maintenance tasks. One is trusted communications through which processes can update software remotely without being marked as suspicious. A communication is considered to be trustful if the three factors associated with it, i.e., "image file of the local process", "communication protocol" and "remote host" are all trusted. Meanwhile, a trusted communication is time limited, i.e., effective only within a predefined time period. Although a trusted protocol, e.g. SSL, is not absolutely secure, a further check of the process' image file, the remote host and the time stamp simultaneously will greatly reduce the attack surface. Moreover, as a computer usually only needs to access a few well-known sites to update its necessary software, the administrator can easily set up the trusted communications required. The other means is a new system call to facilitate a user to manually remove suspicious labels on specific processes or files if the user trusts them. For example, when installing benign software from a CD disk, a user can remove the suspicious labels from the processes which read the executables on the CD disk then Tracer will not affect the installation progress any more. Note that, only a process without a suspicious label has the privilege to use the system call so as to prevent a malware program from bypassing Tracer.

Although bypassing the detection at entrances is difficult, in case that a sophisticated malware program unexpectedly breaks into

the system, we prepare a type of detection at the interior of the system to ambush it. This type of detection monitors the exclusive malware behaviors that a benign program will not exhibit. The current version of Tracer conservatively uses five such behaviors to detect malware programs inside a system as shown in Figure 2. More behaviors can be monitored for malware detection in the interior at the cost of additional FPs. Actually, these behaviors together provide a strong detection capability as they are indispensable to most malware programs, e.g., "Copy itself". In addition, this type of detection will not bring extra performance overhead since the restricting action of Tracer also needs to monitor such behaviors, which will be presented in Section 3.3.

In short, the detection action identifies a process as suspicious if it meets one of the following conditions: *receiving network traffic through dangerous protocols; receiving network traffic through non-dangerous protocols then exhibiting any of the malware behaviors; reading or loading an executable from a removable drive; and exhibiting any of the five exclusive malware behaviors.*

3.3 Tracing Intruders

To track intruders within an operating system, one can use OS-level information flow as done in [18] [26]. However, a major challenge for leveraging OS level information flow to trace suspicious entities is that, file and process tagging usually leads the entire system to be floated with "suspicious" labels and thus incurs too many FPs. To address this issue, we propose the following two methods to limit the number of tagged files and processes in a single OS while preventing malware programs from evading the tracing as much as possible.

For tagging files, unlike the approaches in [18] [26] and the schemes of many malware detection and MAC systems [1][2][5][21] that trace information flow on OS level, Tracer only focuses on the tagging of executables while ignoring non-executables and directories. This is because an executable represents the possible execution flow of the process loading it, thus it should be deemed as an inactive intruder while a process is considered as an active intruder. On the other hand, since there are a huge number of non-executable files and directories within a single OS, not tracing them can prevent the entire file system from being floated with the suspicious labels that mostly are due to FP.

For tagging processes, we observed that the excessive number of tags mainly come from tracing IPC (Inter-Process Communication), i.e. marking a process as suspicious if it receives IPC data from a suspicious process, just as the approaches assumed in [2] [18]. To address this issue, Tracer only tags a process receiving data from dangerous IPCs that can be exploited by a malware program to take control of the process to perform arbitrary malicious behaviors. Note that, dangerous IPCs do not include the other types of vulnerable IPCs that can be used to launch denial-of-service attack, or disclose sensitive information, or escalate the privileges of the processes which send IPC data. Moreover, a dangerous IPC only involves the local IPCs instead of the IPCs over the network, since the detection at entrance can mark a process that receives IPC data from the network as suspicious. In order to identify the dangerous IPCs, we investigated Microsoft Security Bulletins [19], a database storing information about security vulnerabilities on Windows family OS and other Microsoft software. As malware programs usually exploit these vulnerabilities to compromise Windows hosts, Microsoft Security Bulletins become primary sources for analyzing attack vectors of Windows OS as done in [11]. Concretely, we analyzed all vulnerabilities recorded in security bulletins related to named-pipes, local procedure calls, shared

memories, mailslots and Windows messages from 1998 to 2009, as these IPCs send free-formed data that can be crafted to exploit bugs in the receiving process. However, among all of the security bulletins, we only found one dangerous IPC, i.e. MS03-025 [19]. The result reveals that in reality it is quite difficult to propagate malware through local IPCs within a Windows OS since people could only find one dangerous IPC over the period of eleven years. Consequently, Tracer employs a Dangerous-IPC-List to record and trace each type of dangerous IPC since there should be a very limited number of dangerous IPCs in a Windows OS.

Therefore, we mark the following entities as suspicious: *a process spawned by a suspicious process, an executable created or modified by a suspicious process, a process loading an executable with a suspicious label, a process receiving data from a suspicious process through a dangerous IPC, and a process loading a semi-executable or script file with a suspicious label.*

A script file is written in interpreting language, e.g. JavaScript or VBScript, and thus needs execution engine, e.g. wscript.exe or cscript.exe, to load and run it. Accordingly, to defend against a script virus, Tracer should restrict the engine processes that are reading and interpreting a suspicious script file. On the other hand, a semi-executable represents certain types of data files that might contain macro codes, which mainly involves various types of Microsoft Office documents, e.g. Word, Excel, PowerPoint, and template files. Although the macro virus protection in Office software can reduce the chances of macro virus infection, relying on it is very dangerous because crafted macro codes are able to subvert it and cause destructive damages, for example, viruses Melissa and W97M.Dranus. Consequently, Tracer considers a process that reads a suspicious semi-executable as suspicious in order to defeat macro virus.

3.4 Restricting Intruders

In order to disable malware programs on a host, the restricting action monitors and blocks intruders' requests for executing critical malware behaviors listed in Figure 1. Additionally, to reduce FPs, a few malware behaviors that are indispensable to benign programs and do not directly hurt security are not blocked but traced, for instance, "Create executable files" which is often utilized by Internet browsers to download and run ActiveX programs from web pages.

To follow the principle of complete mediation [13] for building a security protection system, Tracer further restricts two special behaviors, called **generic malware behaviors**, to protect security more widely. The first one is "Steal confidential information", which represents all illegal reading of confidential information from files and registry entries. The other is "Damage system integrity", which represents all illegal modifications of the files and registry entries that require preserving integrity. In addition, other behaviors that can be used to bypass Tracer mechanism also need to be monitored and restricted, including "Change file attributes", "Change registry entry attributes", "Execute non-executable files" and "Execute Tracer special system calls". The behavior "Change file attributes" represents changing file extension names to executable or changing file DAC information.

By mediating all these behaviors, Tracer is able to preserve system security and prevent a malware program from propagating itself in the system. To be specific, confidentiality is mainly achieved by blocking the generic behavior "Steal confidential information"; integrity is mainly protected by blocking the generic behavior "Damage system integrity"; availability is defended by blocking the behaviors listed in Figure 1 with the capital letter A

attached; propagation is prevented by blocking the behaviors in Figure 1 with the capital letter P attached.

Meanwhile, blocking these behaviors can also help to defend against unknown malware programs because of two reasons. First, these behaviors are extracted from thousands of malware samples and thus represent popular hacking techniques that are often used in unknown malware programs by malware authors. For example, the behavior "Add IE or Explorer plug-in" is also a popular technique that is frequently used by enormous amount of malware programs both known and unknown to hide and automatically launch themselves, as well as monitor user data. Second, these behaviors are high-level behaviors so that they widely cover various low-level behaviors of various types of malware programs known or unknown. For example, "Communicate with a remote host" involves downloading hacker tools, sending emails to spread malware programs, connecting with a remote host to accept hacker commands, etc. Particularly, the two generic malware behaviors presented previously actually cover all illegal accesses of files/directories and registry entries in the system.

To efficiently restrict these malware behaviors, an issue needs to be addressed. That is how to determine the generic malware behaviors. We identify behaviors "Steal confidential information" and "Damage system integrity" by monitoring illegal reading on read-protected objects and illegal writing on write-protected objects, respectively. However, it is difficult to identify the objects that need protection among a large number of candidates in a Windows OS in order to recognize the generic malware behaviors. A traditional MAC requires users to give every file a security label to identify whether the file needs protection, which in turn becomes a heavy burden on general users. In Tracer, we utilize the existing operating system information "file extension name" and "DAC information" to denote the protected files and registry entries. To be specific, we use file extension names to recognize the executables and configuration files that usually are write-protected. For example, .exe, .com, .dll, .sys, .js, and .vbs represent executables; .ini and .inf represent configuration files. Meanwhile, we use the DAC information of a file to determine whether it is protected. If the user group "Everyone" does not have a read permission on a file, the file is treated as read-protected; if the user group "Everyone" does not have write permission on a file, the file is treated as write-protected. Similarly, we use the permission of user groups "Users", "Everyone" and "Restricted" to recognize protected registry entries. In addition, we do not allow an intruder to change file extension and DAC information. As a result, an attacker is not able to alter a file or registry entry from a protected state to an unprotected state to escape the access control mechanism. With above methods, the configuration work required to identify files and registry entries to be protected is significantly reduced without changing the user's usage convention.

3.5 Limitation

Although Tracer has advantages on compatibility, usability and defense against unknown malware programs, it also has a limitation that requires further research. As based on static malware characteristics and behaviors, detection in Tracer cannot be dynamic over time in the sense of evolving with the emergence of new malware characteristics and behaviors. In future, we will try to improve Tracer to address this limitation.

4. IMPLEMENTATION

To evaluate the effectiveness of Tracer model, we have developed a prototype implementation for Windows XP, and carried out a series of experiments. Although XP is not as new as

Table 1. Decision logic of Tracer. The Benign Process and Suspicious Process columns represent that the processes requesting the behaviors below are benign or suspicious respectively. Ip and Ie indicate setting the related process or executable as suspicious respectively. D indicates denying the behavior request.

Behaviors		Benign Process			Suspicious Process		
		Detect	Trace	Restrict	Detect	Trace	Restrict
1. Communicate with a remote host	Normal Communication	Ip					
	Trusted Communication						
2. Create executable files					Ie		
3. Modify registry for startup						D	
4. Copy itself		Ip				D	
5. Obtain personal or system information						D	
6. Inject into other processes		Ip				D	
7. Modify executable files		Ip				D	
8. Create or modify Windows services						D	
9. Change security settings						D	
10. Add IE or Explorer plug-ins						D	
11. Start hidden network clients		Ip				D	
12. End anti-malware processes or services		Ip				D	
13. Modify system configuration files						D	
14. Log keystrokes and mouse clicks						D	
15. Copy special configuration files						D	
16. Copy system executable files					Ie		
17. Uninstall software						D	
18. Copy executables to removable drives						D	
19. Create Windows hooks						D	
20. Install or modify drivers						D	
21. Change file time						D	
22. Capture screen shots						D	
23. Restart computer						D	
24. Make system/hidden directories						D	
25. Close security alert windows						D	
26. Modify layered service providers						D	
27. Install screen savers						D	
28. Change desktop backgrounds						D	
30. Add data streams						D	
Damage system integrity						D	
Steal confidential information						D	
Mount removable drives		Ie			Ie		
Change file attributes						D	
Change registry entry attributes						D	
Create processes					Ip		
Load suspicious executables					Ip		
Read certain suspicious executables					Ip		
Communicate with local processes					Ip		
Execute non-executable files						D	
Execute Tracer special system calls						D	

Vista, it is enough for verifying the Tracer model since both versions of OS have very similar system calls and Win32 API functions based on which Tracer works. Moreover, if developing the prototype on Vista, the MIC might interfere with Tracer as both schemes attempt to complete MAC tasks.

4.1 Implementation

Tracer implementation consists of two parts: Interception and Decision. Most of the implementations are located in the kernel so that they are difficult to be bypassed. The Interception part monitors Native Windows API functions (i.e. system call) at the kernel level and Win32 API functions (i.e. system library functions) at the application level, then issues behavior requests to the Decision part, and allows or disallows a behavior according to the result returned from the Decision part. The intercepted

behaviors are listed in Table 1.

Most behaviors can be intercepted by monitoring only one essential system call function or a Win32 API function, for example, monitoring NtDeviceIoControlFile() for “Communicate with a remote host”, monitoring NtCreateFile() for “Create executable”, monitoring NtOpenFile() for “Steal confidential information”. Some behaviors consist of more than one system call or Win32 function, for instance, the behavior “Inject into other processes” consists of OpenProcess(), VirtualAllocEx(), WriteProcessMemory(), CreateRemoteThread(), etc. Considering the performance impact, we only intercept the first essential function, i.e. OpenProcess(), and block it if a suspicious process tries to perform an execution, such that the subsequent calls, i.e. WriteProcessMemory() and CreateRemoteThread(), which would cause damages are not executed any more. Moreover, to prevent intended bypassing, Tracer always intercepts a function at the kernel level rather than the application level if possible. Thus, for the behavior “Inject into other processes”, Tracer actually intercepts NtOpenProcess() at the kernel level rather than OpenProcess() at the application level. However, the file-copying like behaviors in Figure 1 can not be identified by intercepting a single system call as they consist of two essential system calls for reading and writing files respectively. Thus, we devised an algorithm to correlate the two system calls by exploiting related file names and buffers.

The Decision part residing in the kernel handles behavior requests from the Interception part. When making a decision, it first reads the Tracer attributes of processes and files, e.g., suspicious flags and DAC information, and then decides whether to permit the behaviors and whether to modify the Tracer attributes according to the Tracer actions presented in Section 3. Table 1 shows the decision logic implemented in the Decision part.

To be permanent, the suspicious flag of an executable is stored in a specially created file stream of the executable file. The suspicious flag of a process, however, is stored in a data structure associated with the process in the memory. The whole implementation is encapsulated in a kernel driver and a DLL. The kernel driver is responsible for intercepting system calls via modifying the system call entry point in the System Service Dispatch Table (SSDT), and implementing the Decision part within the kernel. The DLL is responsible for intercepting Win32 API functions via modifying the library function entry point in the Import Address Table (IAT) of application processes. Note that, our Tracer implementation does not need to impose any modifications on the Windows or application codes, thus it is highly compatible with existing software.

4.2 Evaluation

We evaluate Tracer performance from three important perspectives: its effectiveness in ensuing security, its compatibility with application software, and the overhead added after enabling Tracer on OS.

Security. To verify the capability of Tracer on restricting malware behaviors, we collected 93 real-world malware samples, most of which are obtained from a publicly available website [20]. 32 of the samples are unknown to Tracer, because they can not be found with the same or different names in Symantec Threat Explorer from which the critical malware behaviors are extracted. We also prepared 54 benign samples mostly from two reputable websites, i.e. technet.microsoft.com and www.download.com. To further facilitate the experiments, we prepare a set of monitoring tools to help check experimental results, which include ApiMonitor to record system call and Win32 API,

ProcessExplorer to analyze processes, Regmon to trace registry activities, and Filemon to monitor file operations. Meanwhile, we set up a local network which consists of two servers and two hosts as a testing environment [25]. One server machine, on which the samples to be tested are intentionally placed, runs an IIS web server, a ftp server and an EZ-IRC server. The other server machine, on which only benign samples are placed, runs an IIS web server to act as a trusted site for testing trusted communications. Note that, in reality the trusted sites can be easily recognized by general users because a host only has to connect to several well-known websites to upgrade its important software. The host machines installed with Windows XP run the client programs that are often the attacking vectors for malware samples, including mIRC, MSN Messenger, MS Outlook, eMule, KaZaA, IE and ftp client, etc. To emulate the real-world usage scenarios, we login the hosts and perform various types of tasks, such as browsing the malicious website and ftp server in the local network and downloading samples, sending and receiving malicious instant messages and emails, accessing P2P shared folders or removable drives that contain samples. Thus, the samples are introduced into a host through various channels. With this testing environment, the capability of Tracer to detect, trace and restrict malware behaviors can be thoroughly evaluated.

For every sample, we perform a two-step experiment. First we run a sample on a host without turning on Tracer and record what happens using the monitoring tools above. Then, we enable Tracer protection, run the same sample, and record what happens again. We can determine whether a sample is indeed disabled from two perspectives. First, we deduce whether malware behaviors are successfully executed by comparing the two versions of logs produced by ApiMonitor, Regmon and Filemon without or with protection. Second, we manually check whether the files, registry entries and processes that are created by the sample and recorded in the former logs are exactly present or not in the logs after turning on the Tracer. Moreover, we restart the computer to see if the sample can be enabled automatically.

The testing results are reported in Table 2. For each type of samples, after turning on Tracer, we record the number of false negatives, i.e., FNs, and the number of FPs. We can see that Tracer was able to correctly disable all malware samples including known and unknown ones, as well as block or cancel all their malware behaviors. However, it falsely stopped 3 benign samples by blocking their behaviors. The FPs were a personal firewall program, a file system tool and a process tool, downloaded from the IRC and web server with which we did not set up a trusted communication. By analyzing the logs, we observed that some behaviors of these benign programs closely resemble those of malware, for example, “Create or modify Windows services”, “Modify system configuration files”, “Install or modify drivers”, “Modify registry for startup”, etc. As Tracer relies on the source and behaviors of a program to identify a malware program, the benign programs that come from a remote host through an untrusted communication are tracked and restricted as suspicious ones. However, one still can make the programs work by manually removing the suspicious flags from the program files before running them.

To compare with other anti-malware techniques on Windows, we performed an experiment to test three popular commercial tools: Kaspersky [27], VIPRE [28] and MIC. The former two running on XP are well known anti-malware tools and have modules blocking suspicious behaviors to defend against unknown malware. The anti-malware tools relying only on signatures can not detect unknown malware [1] and thus are inappropriate to

Table 2. Security test results. FP Rate is 5.6% and FN Rate is 0%.

Samples	Programs			Behaviors			
	Total	FNs	FPs	Total	FNs	FPs	
Known malware	Worm	20	0	-	274	0	-
	Trojan	19	0	-	155	0	-
	Backdoor	17	0	-	152	0	-
	Script Virus	2	0	-	65	0	-
	Macro Virus	3	0	-	49	0	-
	Unknown malware	32	0	-	491	0	-
Benign program	Security utilities	11	-	1	103	-	8
	System utilities	10	-	2	83	-	15
	Games	7	-	0	82	-	0
	Multi-media	10	-	0	36	-	0
	Web Pages	16	-	0	99	-	0
	Sum	147	0	3	1589	0	23

compare with Tracer especially on FP rate. MIC is a partial enforcement of BIBA model in Vista kernel [16], which is the only MAC mechanism in Windows OS family. For every anti-malware technique, we tested all of the samples in Table 2. We count a program as a FP if the anti-malware technique abnormally refuses or alarms at least one of its access requests, since this will affect the running of the testing program or annoy the user. We do not count a program as a FP if it fails on Vista but the failure is not caused by MIC. Figure 3 shows the FP rates (FPR) obtained. MIC and the anti-malware tools have FP rates above 34%, whereas, Tracer has FP rate of merely 5.6%. The high FP rate of MIC comes from the no-write-up rule of BIBA model. The modules that block suspicious behaviors contribute to most of FPs of the anti-malware tools. The fundamental reason is that the anti-malware tools identify a suspicious behavior only based on the behavior itself while Tracer further considers the suspicious label of the process requesting the behavior. On the other hand, the FN rates of Kaspersky, VIPRE and Tracer are almost all zero. However, MIC is observed to have a high FN rate of 42%. One possible reason is that MIC does not implement the no-read-down rule of BIBA model [16] in order to avoid a significant impact on the usability and compatibility of Windows which is a commodity OS. As a result, some sophisticated malware programs can manage to bypass it. Nevertheless, with MIC, Vista can still achieve a significant security improvement compared with XP that can not defeat any malware samples by itself.

Compatibility. The requirement for compatibility is that existing Commercial Off-The-Shelf (COTS) software can run on the MAC prototype without causing significant incompatibility problems. On the two hosts with Windows XP installed, we run many commonly used network-dependent applications and local applications e.g. Internet Explorer, MS Outlook Express, MS word, MS excel, MS Power Point, MS Messenger, mIRC and Visual C++. We set the protocols HTTP, POP3, IRC, SMTP, FTP, FastTrack, eDonkey and ICMP as dangerous. We define trusted communications for downloading useful applications and Windows update as follows: {"C:\Windows\system32\wuauclt.exe", "SSL", "update.microsoft.com"}, {"C:\Windows\system32\svchost.exe", "SSL", "update.microsoft.com"}, {"C:\Windows\system32\wbem\wmiprvse.exe", "SSL", "update.microsoft.com"}. On the system tested, we send emails, browse websites through Internet, edit word documents, develop VC++ programs, share files remotely, update Windows and move files through USB disks, etc. The system works well for the past a few months, without need of modifications of existing software or running into failures.




Figure 3. Comparing false positives with commercial anti-malware techniques on Windows

Table 3. Overhead of Tracer (CPU Cycles), The columns Tracer-m, Tracer-bf and Tracer-b show the CPU cycles taken by the malware programs, the benign programs with and without suspicious flags running on Tracer, respectively.

Functions	Native	Tracer-m	Tracer-bf	Tracer-b	Functions	Native	Tracer-m	Tracer-bf	Tracer-b
NtCreateFile	334492	348471(4.2%)	348197(4.1%)	338506(1.2%)	CreateService	6568120	6679778(1.7%)	6679625(1.7%)	6568323(<0.1%)
NtOpenFile	167620	175263(4.6%)	173235(3.3%)	169713(1.2%)	OpenService	5490443	5609379(2.2%)	5609352(2.2%)	5490560(<0.1%)
NtWriteFile	245179	324623(32.4%)	278214(13.5%)	249832(1.9%)	NtSetValueKey	210491	225120(6.9%)	225093(6.9%)	210493(<0.1%)
NtCreateNamedPipeFile	204711	214743(4.9%)	214751(4.9%)	204789(<0.1%)	NtCreateKey	281722	296371(5.2%)	296008(5.1%)	281784(<0.1%)
NtCreatePort	37241	40221(8%)	40180(7.9%)	37275(<0.1%)	NtCreateProcessEx	206458	215435(4.3%)	215426(4.3%)	208849(1.2%)

Performance overhead. The performance overhead of Tracer comes from the overhead of executing additional instructions associated with every intercepted system call and Win32 API function. In the following experiment, we evaluate the additional overhead imposed by Tracer enforcement. The test-bed is a Pentium-4 2.8GHz machine with 1GB memory running Windows XP SP2. We first disable Tracer, run a group of benign and malware programs, and count the average CPU cycles spent in each system call and API function through rdtsc instruction. Second, we enable Tracer, run the malware programs, the benign programs with suspicious flags and without suspicious flags to perform the test again. In all tests, the average CPU cycles of every system call or API function is calculated from 100 invokes. Results are shown in Table 3. With Tracer enabled, the malware programs have 1.7%~32.4% more performance penalty than native, while the benign programs have only 0~13.5%. The highest performance penalty comes from the interception of NtWriteFile() as a result of capturing file-copying behaviors. The overhead incurred on benign programs is lower than 2%. Therefore, the general performance impact from the system call and Win32 API function interception is acceptable.

5. RELATED WORK

The concepts of Tracer model are partially inspired by the MAC models such as DTE [9], LOMAC [3], UMIP [2] and PRECIP [5]. DTE proposed by Lee Badger et al. [9] is a classical MAC model to confine process execution, which groups processes and files into domains and types respectively, and controls accesses between domains and types. Tracer can be regarded as a simplified DTE that has two domains (i.e., benign and suspicious) and four types (i.e., benign, read-protected, write-protected and suspicious). Moreover, Tracer can automatically configure the DTE attributes (i.e., domain and type) of processes and files under the support of intrusion detection and tracing so as to improve usability.

Both LOMAC and UMIP evolve from BIBA model, which aims to add usable and compatible mandatory integrity protections into mainstream operating systems. The LOMAC model deals with the pathological cases in the Low-Water Mark model's behaviors to decrease its partial compatibility cost. The UMIP model is designed to preserve system integrity in the face of network-based attacks in a highly usable manner. The PRECIP model addresses several practical issues that are critical to contain spyware that intends to leak sensitive information. Tracer, however, differs from these

MAC models in that, it traces suspected intruders and restricts their behaviors rather than restricts information flow. With this novel concept, it is able to considerably reduce FPs and automatically deploy security labels, which result in good compatibility and usability. Meanwhile, the philosophy of Tracer is roughly similar to the risk-adaptive access control [10] that targets to make access control more dynamic and flexible so as to achieve a better tradeoff between risk and benefit. Tracer dynamically changes the security labels of the processes that derive from dangerous sources to reduce the risk of executing malware behaviors, while not restricting other behaviors and processes at all to preserve the benefits of compatibility and usability.

Most existing anti-malware technologies are based on detection [22][23][24]. Tracer tries to combine detection and access control so that it not only can detect but also can block malware behaviors before their harming security. Tracer can be used as a complementary technique to malware detection, to act as an automatic intrusion response mechanism of a malware detection system, since it is able to further trace intruders and restrict their behaviors after malware detection. Another anti-malware technology that resembles Tracer is behavior blocking [29], which can confine the behaviors of certain adverse programs that are profiled in advance. However, Tracer does not need to profile program behaviors beforehand, and can confine the adverse programs that execute malware behaviors.

Many commercial anti-malware tools [27][28] also have a behavior-based module to defend against unknown malware programs. The major difference between Tracer and the commercial tools is that Tracer determines a malicious behavior based not only on the behavior itself but also the source of the process requesting the behavior. In other words, Tracer further checks the suspicious label of the process requesting the behavior rather than merely the behavior as anti-malware tools do. As a consequence, Tracer produces much less false positives than that of the commercial-tools as shown in Section 4.2.

6. CONCLUSIONS

In this paper, we propose a novel MAC model that integrates intrusion detection and tracing to defend against malware in a commercial OS. We have extracted 30 critical malware behaviors and three common malware characteristics from the study of 2,600 real-world malware samples and analyzed the root reasons for the incompatibility and low usability problems in MAC, which will benefit other researchers in this area. Based on these

studies, we propose a novel MAC model, called Tracer, to disable malware timely without need of malware signatures or other knowledge in advance. It detects and traces suspected intruders so as to restrict malware behaviors. The novelty of Tracer design is two-fold. One is to use intrusion detection and tracing to automatically configure security labels. The other is to trace and restrict suspected intruders instead of information flows as done by traditional MAC schemes. Tracer doesn't restrict the suspected intruders right away but allows them to run as long as possible except blocking their critical malware behaviors. This design produces a MAC system with good compatibility and usability. We have implemented Tracer in Windows OS and the evaluation results show that it can successfully defend against a set of real-world malware programs, including unknown malware programs, with much lower FP rate than that of commercial anti-malware techniques.

7. ACKNOWLEDGMENT

We would like to thank all the anonymous reviewers for their insightful comments and feedback. This work is supported by Natural Science Foundation of China under grants No. 60703103, No. 60833005 and No. 60873213, US National Science Foundation under grants CNS-0751121, CNS-0751121 and CNS-0628093.

8. REFERENCES

- [1]. Heng Yin, Dawn Song, Manuel Egele, Christopher Kruegel, and Engin Kirda. 2007. Panorama: capturing system-wide information flow for malware detection and analysis. In *Proceedings of the 14th ACM conference on Computer and communications security*. ACM, New York, NY, USA, 116-127.
- [2]. Ninghui Li, Ziqing Mao, and Hong Chen. Usable Mandatory Integrity Protection for Operating Systems. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Washington, DC, USA, 164-178.
- [3]. Timothy Fraser. 2000. LOMAC: Low Water-Mark Integrity Protection for COTS Environments. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy (SP '00)*. IEEE Computer Society, Washington, DC, USA, 230-245.
- [4]. Microsoft, Mandatory Integrity Control, [http://msdn.microsoft.com/en-us/library/bb648648\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/bb648648(VS.85).aspx).
- [5]. XiaoFeng Wang, Zhuwei Li, Jong Youl Choi, Ninghui Li. PRECIP: Towards Practical and Retrofittable Confidential Information Protection. In *Proceedings of 15th Network and Distributed System Security Symposium*, 2008.
- [6]. Symantec, Inc, http://www.symantec.com/business/security_response/threatexplorer/threats.jsp.
- [7]. D.E. Bell and L. J. LaPadula. Secure computer system: Unified exposition and Multics interpretation. Technical Report ESD-TR-75-306, MITRE Corp. MTR-2997, Bedford, MA, 1975. Available as NTIS AD-A023 588.
- [8]. K. J. Biba. Integrity considerations for secure computer systems. Technical Report MTR-3153, MITRE, April 1977.
- [9]. L. Badger, D. F. Sterne, D. L. Sherman, K. M. Walker, and S. A. Haghghat. 1995. Practical Domain and Type Enforcement for UNIX. In *Proceedings of the 1995 IEEE Symposium on Security and Privacy (SP '95)*. IEEE Computer Society, Washington, DC, USA, 66-77.
- [10]. Pau-Chen Cheng, Pankaj Rohatgi, Claudia Keser, Paul A. Karger, Grant M. Wagner, and Angela Schuett Reninger. 2007. Fuzzy Multi-Level Security: An Experiment on Quantified Risk-Adaptive Access Control. In *Proceedings of the 2007 IEEE Symposium on Security and Privacy (SP '07)*. IEEE Computer Society, Washington, DC, USA, 222-230.
- [11]. Michael Howard, Fending Off Future Attacks by Reducing Attack Surface, <http://msdn.microsoft.com/en-us/library/ms972812.aspx>, 2003.
- [12]. Marius van Oers, OSX Malware not taking off yet, <http://www.avertlabs.com/research/blog/index.php/2007/03/20/osx-malware-not-taking-off-yet/>, March, 2007.
- [13]. Jerome Saltzer and Michael Schroeder. The protection of information in computer systems. *Communications of the ACM*, 17(7), 1974.
- [14]. Peter Loscocco and Stephen Smalley. 2001. Integrating Flexible Support for Security Policies into the Linux Operating System. In *Proceedings of the FREENIX Track: 2001 USENIX Annual Technical Conference*, Clem Cole (Ed.). USENIX Association, Berkeley, CA, USA, 29-42.
- [15]. Marek Jawurek, RWTH Aachen. RSBAC - a framework for enhanced Linux system security, <http://www.rsbac.org/documentation>, 2006.
- [16]. Microsoft, Windows Vista Integrity Mechanism, [http://msdn.microsoft.com/en-us/library/bb625964\(v=MSDN.10\).aspx](http://msdn.microsoft.com/en-us/library/bb625964(v=MSDN.10).aspx).
- [17]. Virgil D. Gligor, C. S. Chandrasekaran, Robert S. Chapman, Leslie J. Dotterer, Matthew S. Hecht, Wen-Der Jiang, Abhai Johri, Gary L. Luckenbaugh, and N. Vasudevan. 1987. Design and Implementation of Secure Xenix. *IEEE Trans. Softw. Eng.* 13, 2, 208-221.
- [18]. Samuel T. King and Peter M. Chen. 2003. Backtracking intrusions. In *Proceedings of the nineteenth ACM symposium on Operating systems principles*. ACM, New York, NY, USA, 223-236.
- [19]. Microsoft Security Bulletins, <http://www.microsoft.com/technet/security/current.aspx>.
- [20]. Offensive Computing, <http://www.offensivecomputing.net/>.
- [21]. Francis Hsu, Hao Chen, Thomas Ristenpart, Jason Li, and Zhendong Su. 2006. Back to the Future: A Framework for Automatic Malware Removal and System Repair. In *Proceedings of the 22nd Annual Computer Security Applications Conference*. IEEE Computer Society, Washington, DC, USA, 257-268.
- [22]. Engin Kirda, Christopher Kruegel, Greg Banks, Giovanni Vigna, and Richard A. Kemmerer. 2006. Behavior-based spyware detection. In *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15 (USENIX-SS'06)*, Vol. 15. USENIX Association, Berkeley, CA, USA.
- [23]. Lorenzo Martignoni, Elizabeth Stinson, Matt Fredrikson, Somesh Jha, and John C. Mitchell. A Layered Architecture for Detecting Malicious Behaviors. In *Proceedings of the 11th international symposium on Recent Advances in Intrusion Detection*, Springer-Verlag, Berlin, Heidelberg, 78-97.
- [24]. Andrea Lanzi, Monirul Sharif, and Wenke Lee. K-Tracer: A System for Extracting Kernel Malware Behavior. In *Proceedings of Network and Distributed System Security Symposium*, February 2009.
- [25]. Zhihong Zhang, Jianfeng Zhan, Yong Li, Lei Wang, Dan Meng, Bo Sang. "Precise request tracing and performance debugging for multi-tier services of black boxes". *IEEE/IFIP International Conference on Dependable Systems & Networks*, pp.337-346, June 2009.
- [26]. Ashvin Goel, Kenneth Po, Kamran Farhadi, Zheng Li, and Eyal de Lara. The taser intrusion recovery system. In *Proceedings of the twentieth ACM symposium on Operating systems principles (SOSP '05)*. ACM, New York, NY, USA, 163-176.
- [27]. Kaspersky Lab. <http://www.kaspersky.com/>.
- [28]. Vipre, Inc, <http://www.vipre.com/vipre/>.
- [29]. Carey Nachenberg. Behavior Blocking: The Next Step in Anti-Virus Protection. <http://www.securityfocus.com/infocus/1557>, March 2002.