Every Secure System Wants to Be a Reference Monitor

By Mark R. Heckman – ISSA member, Sacramento Valley Chapter

This article discusses the reference monitor abstraction that was first introduced in 1972 and is one of the fundamental concepts that makes cybersecurity a field of engineering rather than just an ad hoc set of "best practices."

Abstract

The reference monitor abstraction was first introduced in 1972 and is one of the fundamental concepts that makes cybersecurity a field of engineering rather than just an ad hoc set of "best practices." Originally created to combat the threat of malicious programs, the principles and components of the reference monitor have come to define what it means for a system to be "secure." In fact, most system vulnerabilities can be traced to violations of one or more of the reference monitor principles. Conscious application of the reference monitor concept during requirements specification and system design, however, can help ensure the security of systems.

Outside of studying for the CISSP exam, I'd guess that relatively few cybersecurity practitioners have heard of the reference monitor concept, and fewer still consciously apply it in their work. That is a shame, because the reference monitor is one of the fundamental ideas that makes cybersecurity a field of engineering, rather than just a collection of ad-hoc practices. Moreover, it is a design concept that was specifically developed to combat the threat of malware.

A reference monitor is an abstract model of controls on the access of people (or programs acting on behalf of people) to information stored in a computer system. As originally described in the 1972 Anderson report, "In concept, the reference monitor mediates each reference made by each program in execution by checking the proposed access against a list of accesses authorized for that user. The reference monitor is implemented as a reference validation mechanism."[1]

In the abstract model, a reference monitor controls access by active entities, called *subjects*, to passive, information-containing entities, called *objects*. The reference monitor consults an authorization database to decide if accesses are permitted and records attempted or allowed accesses in an audit log. These components of the reference monitor are depicted in figure 1.

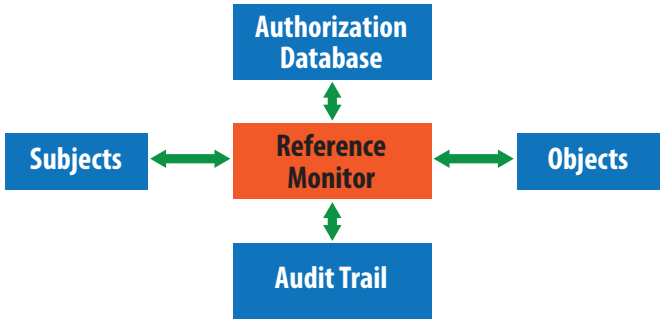

Figure 1—Reference monitor components

Brinkley and Schell describe a library of sensitive paper documents as an example of a reference monitor.[2] Their example posits an organization that has a collection of extremely sensitive documents, such as corporate plans and strategies or classified national security information. In order to protect the documents, the organization places them into a secure room—a bank vault, perhaps. A guard stands guard at the door to the vault and checks the credentials of everyone who attempts to enter the vault. The guard compares the identity of each person against a list that specifies the identities of people who are authorized to access the documents. There also must be some way to ensure individual accountability for each of the people who visits the library, so the system in-

1 J. P. Anderson, "Computer Security Technology Planning Study Volume II," Electronic Systems Division, Air Force Systems Command, Hanscom Field, Bedford, MA, 1972 – http://csrc.nist.gov/publications/history/.

2 D. L. Brinkley and R. R. Schell, "Essay 2. Concepts and Terminology for Computer Security," in *Information Security: An Integrated Collection of Essays*, M. D. Abrams, S. Jajodia and H. J. Podell, Eds., Los Alamitos, California, USA, IEEE Computer Society Press, 1994, pp. 45-48.

cludes a visitor log signed by both visitor and guard. The visitors in the example are instances of subjects in the abstract model of a reference monitor, and the documents are instances of objects. The list of authorized visitors corresponds to the authorization database of the abstract model and the visitor log corresponds to the audit log.

Of course, the guards in the library must also have some means to authenticate the identity of each person such as a valid photo ID. Identification and authentication, and protection and use of the audit log, are "supporting policies" in the reference monitor. They support the primary mission of the reference monitor, which is to enforce the access control policy. Worth noting is that the audit log in a reference monitor is intended to be used for accountability, not intrusion detection, because the reference monitor is expected to correctly enforce the policy.

There are many real-world examples of reference monitors. These include operating systems such as SCOMP and GEMSOS,[3] which were built on security kernels specifically intended to be implementations of reference monitors. These are some of the most secure systems ever developed and were themselves used as components in reference monitor-based, secure networked systems for the US Department of Defense (e.g., BLACKER[4]). But the reference monitor concept is not restricted to these now somewhat historic systems;[5] it is applicable in contemporary networked environments as well.

Consider a corporate network that uses Windows Active Directory. The network controls the access of "security principals"—entities such as user accounts, or threads or processes that run in the context of a user account, or a group of accounts—to various resources such as printers, files, and folders.[6] This is an implementation of a reference monitor that has all of the components of the reference monitor model. The security principals are subjects and the resources are objects. Active Directory audit logs record changes to objects and their attributes. Access control lists and the authentication server database serve as the authorization database used by the reference monitor to enforce a discretionary access control policy.

As an abstract model of security controls, the reference monitor model does not specify a particular access control policy that is to be enforced; it is up to an organization to specify an appropriate policy. The access control policy is the definition of security for the system in terms of allowed access by people to information stored in the system. A bad policy can lead to undesired outcomes, even if reliably enforced by a reference monitor. What if, for example, the example document library in the example above had a policy that permitted any authorized users on their own volition to add other people to the authorized visitor list? The vault and guards would faithfully make sure that only people on the authorized visitor list were given access, but that could not guarantee the safety of the documents when anyone could so easily become authorized.

Similarly, the model does not specify a particular implementation, but any implementation of a reference monitor must satisfy three, fundamental properties:

1. **The reference validation mechanism must be invoked for every reference by a subject to an object.** This property is called "completeness" or "non-bypassability." Consider the case of the document library. If someone could somehow bypass the guard, then there would be no security on the documents. But there is only the one door to our library vault, so anyone who wants to enter must pass by the guard. By contrast, consider a vault in one of the *Mission Impossible* movies, where Tom Cruise was able to enter through a ventilation duct. That vault was not a good implementation of a reference monitor because the reference validation mechanism could be bypassed. And what if the library permitted authorized users to take documents home with them? The reference monitor could not perform reference validation for any documents outside the vault.
2. **The reference validation mechanism must be protected from unauthorized modification.** This property is called "isolation" or "tamperproofness." In the document library, for example, if someone were to bribe the guard,

3 T. Jaeger, *Operating System Security*, Morgan and Claypool Publishers, 2008.

4 "BLACKER: security for the DDN examples of A1 security engineering trades," in *Research in Security and Privacy, 1992. Proceedings*, 1992 IEEE Computer Society Symposium on, Oakland, CA, 1992.

5 GEMSOS is still offered as a commercial product by the Aesec Corporation (http://aesec.com/).

6 "What Are Security Principals?," [Online]. Available: https://technet.microsoft.com/en-us/library/cc780957%28WS.10%29.aspx. [Accessed May 2015].

or somehow distract him or her, or even tamper with the authorized visitor list, there would be no security on the documents. Similarly, if someone were able to break through the walls of the vault and grab the documents, even the most diligent guard could not protect the documents.

3. **The reference validation mechanism must be small and simple enough to be thoroughly analyzed and tested for correctness.** This property is called "verifiability." A flawed reference validation mechanism cannot be depended on to enforce the desired policy, but a large, complex system is impossible to verify through testing or other contemporary techniques.

As the Anderson report noted, "computer systems present to a malicious user a unique opportunity for attempting to subvert through programming the mechanism upon which security depends (i.e., the control of the computer vested in the operating system)." Any example of malware you can cite depends for its success on a violation of at least one of the three reference monitor principles, or a failure of the identification and authentication mechanism, or an inadequate policy. Any system that upholds the three reference monitor principles, has strong identification and authentication mechanisms, and enforces an adequate policy, on the other hand would be highly resistant to malware.

Take, for example, Stuxnet, which infected systems in several ways. One version of Stuxnet spread through USB drives using the Windows autorun feature or through the local network via a print-spooler zero-day exploit.[7] The first of these could be considered to be a poor policy (why should autorun programs be permitted to modify the system configuration and install a rootkit?). The second is due to a violation of the third reference monitor principle (the reputed 45 million lines of code in Windows is too large and complex to thoroughly analyze and test for security). In both cases, the second principle (isolation) was violated. It isn't clear if the access control policy was bypassed (violating the first principle) or was simply inadequate.

7 K. Zetter, "An Unprecedented Look at Stuxnet, the World's First Digital Weapon," *Wired*, 3 November 2014 – http://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/.

The example of Stuxnet suggests that using the reference monitor concept as a framework could be very valuable as a framework when designing a system and identifying necessary security controls. A step-wise approach to requirements gathering and system design based on the reference monitor concept might take the following form:

1. **Create a suitable security policy.** This step requires identifying the subjects and objects in the system and specifying the permitted access of subjects to objects in terms of observation and modification. It is critical that the protected information and allowed access be identified or else there is no definition of what it means for the system to be considered "secure." Even a rock-solid reference monitor can only enforce "security" as defined by the policy.
2. **Identify the identification and authentication controls.** Proper identification and authentication is necessary for the reference monitor to be able to enforce the policy and accountability.
3. **Identify the authorization controls.** This is the reference validation function that is the core of the reference monitor.
4. **Identify controls to enforce the completeness and isolation properties of the reference monitor**. This can be a tough proposition in contemporary systems. A rigorous threat-model can help identify potential weaknesses that might allow the reference monitor to be bypassed or tampered with, and help come up with compensating controls. For example, total disk encryption is a compensating control for environments where the isolation of the reference monitor mechanism cannot be guaranteed (i.e., computers could be lost or stolen) and to help guarantee completeness (no one can access the encrypted data outside of the organization, assuming the keys are kept safely).
5. **Design the audit log mechanism**. The records should contain sufficient fields to enforce accountability (e.g., subject and object identity and type of access) and should be protected from tampering and compromise.

Note that the third reference monitor principle, verifiability, is conspicuously missing from this list. That is because most of the components we currently have available to us do not satisfy this principle, so we cannot compose a system out of them that would satisfy the principle. If you are designing a system from scratch, however, you are not off the hook! Proper attention to good software engineering practices, such as the use of layering and modularization, can help increase the assurance of the reference monitor mechanism. The higher evaluation levels of the Common Criteria (ISO/IEC 15408) are a good source for techniques that can be used when de-

Continued on page 30

Profiling in a Digital World

By Ray Yepes – ISSA member, South Texas Chapter

The author discusses the important role of profiling in digital cases by presenting a case study where criminal profiling played a vital role in the outcome of the investigation.

Abstract

Profiling is an investigative technique that has been used for many years. Although applied to many areas not associated with criminal justice, it is generally associated with law enforcement. Traditionally this technique has been applied to criminal activities that are visible to the naked eye; however, this technique has been evolving to fit into the cybercrime arena. This article discusses the important role of profiling in digital cases by presenting a case study where criminal profiling played a vital role in the outcome of the investigation.

Traditionally, the art of profiling applies to criminal activities that are visible to the eye and observation of the investigator. In fact, profiling involves the analysis of personal characteristics or behavioral patterns, which allows an investigator to make generalizations about a person or a crime scene. In other words, profiling employs analysis to determine whether a particular person may be engaged in a particular crime, as determined by evidence. However, unlike traditional crime scenes that are tangible and have observable evidence, cybercrimes are not as easily examined and observed— there are no physical weapons or visible signs that might contribute to the art of profiling.

With the evolution of cybercrimes, digital investigations have increased exponentially. A decade or two ago, criminals were primarily murderers, gangsters, bank robbers, burglars, and those who committed other "traditional" crimes, but this is no longer true. Now, cybercriminals are by far the most predominant type of criminal. In fact, computer crime is the fastest growing type of illegal activity in both the United States and abroad. According to the US Department of Justice, "Cybercrime is one of the greatest threats facing our country [today], and has enormous implications for our national security, economic prosperity, and public safety."[1] As technology continues to evolve, the range of threats and challenges will continue to grow.

Criminal profiling as an investigative tool for computer-related crimes

John Edward Douglas, former special agent with the Federal Bureau of Investigation (FBI), was one of the first to master and develop a criminal profiling methodology.[2] During his career, Douglas examined hundreds of crime scenes and interviewed dozens of serial killers with the intention of creating criminal profiles of the perpetrators. But how might a criminal profiler like Douglas apply a traditional profiling approach to cybercriminals? While difficult and different, it is possible to apply an alternative investigative approach to aid in the profiling of computer-related crimes.

Although Douglas trained me in the art of criminal profiling, many of the techniques and methodologies I learned from him could not be applied directly to computer-related crimes. However, his training did provide the foundation that allowed me to develop a methodology that could be applied in the cyber arena, and Douglas' methodology (Why + How = Who) is still applicable when profiling cybercriminals and cyber "crime scenes." Determining why and how the crime was committed will facilitate the discovery of who committed the crime.

As stated previously, the traditional approach to criminal profiling is largely based on tangible evidence and observation. When dealing with cybercrimes, evidence is much less tangible. For instance, in the trial of a murder case the prosecutor can hold and show jurors the actual murder weapon that was used to commit the crime. The question, then, is how does one relay the same presentation when dealing with digital evidence? By collecting and analyzing the details of digital crimes, an investigator can develop profiles of the perpetrators. To accurately do so, however, the examiner must possess a unique blend of knowledge in various disciplines including but not limited to profiling techniques, technology, cybersecurity, digital forensics, and interviewing and interrogation techniques. One can be an expert in profiling techniques (like Douglas), but without proficient knowledge of technology

1 Cybercrime. US Department of Justice, Offices of the United States Attorneys – http://www.justice.gov/usao/priority-areas/cyber-crime.

2 Criminal Profilers Investigate Murderous Minds, NPR – http://www.npr.org/templates/story/story.php?storyId=16117564.

and digital investigations, it would be nearly impossible to accurately profile the perpetrator of a computer-related crime.

The following case study, based on an actual case I assisted with, may better illustrate the importance of this unique blend of knowledge. A few years ago, a large oil company was the victim of a digital breach, and due to the magnitude and sensitivity of the occurrence the FBI was brought in to assist with the investigation of the cyber attack. I was brought in to assist with the profiling of the investigation.

Criminal profiling from crime scene analysis

My first task as the profiler was to narrow down the list of possible crime suspects. With this goal in mind, I needed to determine if the attack came from the outside (possible crime suspects numbering approximately six billion people, at the time) or from the inside (possible crime suspects numbering approximately 60,000 people), so I requested the blueprints of the entire information technology (IT) infrastructure for that particular location, including but not limited to switches, routers, security appliances, VPN appliances, NAS, and firewalls (containing firewall rules and filters in place). One might expect a Fortune 100 company to have this information readily available, but this was not the case. This process required numerous interviews and meetings with the different stakeholders (e.g., director of infrastructure, IT security director, enterprise administrator, and so forth).

It is highly recommended that this particular task be handled by an outside party (whether law enforcement or a vendor) to avoid omissions. Internal personnel may not disclose appliances that have not been properly patched, maintained, or updated for fear their employers may think they are not doing their jobs, and if appliances are omitted, the investigator will not be able to see the full spectrum of the IT infrastructure. Also, it is normally through these unmaintained appliances that hackers will gain unauthorized access to networks; without the analyses of them, the investigator will only have a partial picture of the crime. For this particular investigation, I performed a full network scan, which proved very helpful and uncovered about a dozen appliances that were unaccounted for, undocumented, and undisclosed.

It is important to note that if a crime was committed by an insider with intimate knowledge of the network infrastructure, it makes it more likely that the appliances that would prove useful for the investigation would not be disclosed. Therefore, it is recommended that investigators not rely on interviews and network blueprints, but conduct their own network scans to ensure all appliances are known.

My next task was to dissect the attack from a technological standpoint. This required studying and analyzing the propagation technique (if any), transmission protocol, communication port, payload (purpose of the attack), replication of the attack, and packet analysis of network traffic while the attack is propagating or occurring, among other factors. For this analysis, it is highly recommended to set up a virtual honeypot[3] in order to capture all network traffic and activity directed to and from the affected system(s) or host(s) and look for irregularities. For this task, a network traffic analyzer tool, which normally uses a statistical algorithm (or algorithms) to pinpoint local network irregularities, will help the investigator determine what is normal and what is not normal network traffic within the environment.

Although honeypots were initially designed as baits for hackers to monitor their activities and hacking techniques, I have used honeypots in conjunction with packet sniffers when analyzing cyber attacks such as viruses, Trojans, malware, ad-

3 "Virtual Honeypots," Infosec Institute – http://resources.infosecinstitute.com/virtual-honeypots/.

Every Secure System Wants to Be a Reference Monitor continued from page 28

veloping security requirements and requirements for secure development.[8]

The components and principles of the reference monitor are, separately, prominent features of many secure design guidelines and frameworks, and most experienced secure system designers intuitively incorporate them in their designs. The guidelines, frameworks, and intuition, nevertheless, lack a unifying order that ties all of the components together into a cohesive whole, and that could mean that some of the principles and components of the reference monitor concept may be missed, leaving gaps in the security of a system.

Furthermore, the reference monitor concept represents an idealized system, but systems in the real world are imperfect. We create systems, for example, out of components whose correctness we cannot verify, and reliably enforcing completeness and isolation of the reference validation mechanism is difficult, at best. The reference monitor concept, however, encapsulates what it means to say a system is "secure." Conscious application of the reference monitor concept during system design helps ensure that our systems will be as secure and as resistant to malware as available technology can make them.

8 N. Mead, "The Common Criteria," U.S. Department of Homeland Security, 5 July 2013. [Online] – https://buildsecurityin.us-cert.gov/articles/best-practices/requirements-engineering/the-common-criteria. [Accessed June 2015].

About the Author

Mark Heckman, PhD, CISSP, has worked in the field of information security for over 30 years as a researcher, developer, and practitioner. He currently is a senior lecturer in the Informatics Cybersecurity Engineering program at the University of Southern California. He may be reached at mark.r.heckman@gmail.com.