ShadowNet: An Active Defense Infrastructure for Insider Cyber Attack Prevention

Xiaohui Cui[1], Wade Gasior[2], Justin Beaver[1], and Jim Treadwell[1]

[1] Oak Ridge National Laboratory
Oak Ridge, TN, USA
{cuix,beaverjm,treadwelljn}@ornl.gov
http://cda.ornl.gov
[2] University of Tennessee at Chattanooga
Chattanooga, TN, USA
wadegasior@gmail.com

Abstract. The ShadowNet infrastructure for insider cyber attack prevention is comprised of a tiered server system that is able to dynamically redirect dangerous/suspicious network traffic away from production servers that provide web, ftp, database and other vital services to cloned virtual machines in a quarantined environment. This is done transparently from the point of view of both the attacker and normal users. Existing connections, such as SSH sessions, are not interrupted. Any malicious activity performed by the attacker on a quarantined server is not reflected on the production server. The attacker is provided services from the quarantined server, which creates the impression that the attacks performed are successful. The activities of the attacker on the quarantined system are able to be recorded much like a honeypot system for forensic analysis.

1 Introduction

Cyber security has become a national priority. Despite the number of recent news reports about hacker attacks and external network intrusions, trusted employees and business partners with authorized access to network still pose the greatest security risk to the government and private companies [1]. It is usually assumed that users who are given access to network resources can be trusted. However, the eighth annual CSI/FBI 2003 report [2] found that insider abuse of network access was the most cited form of attack or abuse. It is reported that 80 percent of respondents were concerned about insider abuse, although 92 percent of the responding organizations employed some form of access control and insider prevention mechanism. There are also large amount of insiders committing espionage cases have caused tremendous damage to U.S. national security. Two infamous insider threat cases, one is the case of former FBI agent R. P. Hanssen, who was convicted for spying for Russia. The another case is the United States diplomatic cables leak. In the FBI Hanssen case, over a span of more than 15 years, Hanssen provided his Russian contacts with highly classified

B. Murgante et al. (Eds.): ICCSA 2012, Part IV, LNCS 7336, pp. 646–653, 2012.

documents and details about U.S. intelligence sources and electronic surveillance taken directly from his employer, the FBI.

Detecting and preventing insider user misuse involves many challenges because insiders understand their organization's computer system and how the computer network system works. Inside users typically also have greater knowledge than outsiders do about system vulnerabilities. Therefore, the chances of a successful attack can be greater for an insider attack than for an outsider attack [4]. For instance, the knowledge that a malicious insider has about the sensitivity of information gives him/her a better chance to breach information confidentiality. Insider user misuse is different from outsider misuse with respect to the nature of the threats that both cause. However, because of lacking of understanding the differences in implementation of detection and prevention techniques between insider misuse and outsider misuse, most institutions intent to apply existing cyber security techniques to both threats [5].

2 Related work

A Honeypot [6,7] is an internet-attached server that acts as a decoy, luring in potential hackers in order to study their activities and monitor how they are able to break into a system [8]. As shown in Fig.1, the traditional Honeypots are designed to mimic systems that an intruder would like to break into, but limit the intruder from having access to an entire network. The most widely used honeypots is the honeyds [9] that run on a honeyd server and represent unused IP addresses in the organization's network. They function through emulating operating systems and services, thus allowing them to interact with the attacker. Any attempted connection to one of the honeypot servers is assumed to be unauthorized (malicious activity). An outside attacker most likely has little knowledge about the enterprise network structure and the network location (Internal IP address) of the sensitive information stored. Those outsiders have to depend on NMAP [10] or other network scanning software for mapping the target network structure and finding out the most vulnerable system for hacking. It is possible that an intruder might spend days hacking into an old Windows system that is only used as printer server and contains no valuable information. This kind of situation gives the traditional honeypot technology a chance for acting as a decoy vulnerable system and attracting intruder attacks. By using honeypot in a network, it can reduce the change for intruder find out the real valuable target.

However, the traditional honeypot may have difficulty in preventing attacker who has some insider information about the network. Different from outsider, the insider has more information about the enterprize network architecture and the computer system he wants to attack. The malicious insider most likely knows the IP address or machine name where the sensitive information is stored. As an employee of the enterprize, the intruder can easily access the enterprize internal network without passing through the enterprize firewall. So, an insider doesn't need to use NMAP or other network scan software to randomly discover

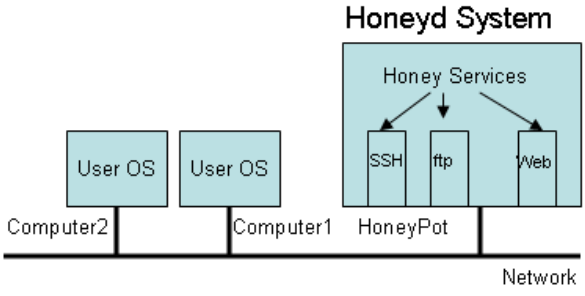

Fig. 1. Traditional HoneyPot Infrastructure

the vulnerable system in the network. Instead, the valuable system location, IP address event low privilege user account can be easily discovered by insider by using social engineering. To protect from being exposed, some insiders might not steal valuable data that he can legally access. Instead, they will use hacking tools to acquire information that he can not legally access. Before the insider launch an attack, he already knows the IP address of the systems that may have valuable data and he may even have low-level privileges to legally access the system. In this case, deploying the traditional honeypot system will not help detect a malicious insider.

3 An Active Defence Infrastructure

3.1 Overview

In this research, we developed an active defense infrastructure, called ShadowNet, for insider attack prevention and forensic analysis. The diagram of the infrastructure is shown in Fig.2 and Fig.3. The purpose of this ShadowNet is to help mitigate risk in an organization by actively preventing the malicious insider to harm a real computer or application and preventing him from spreading his attack to other computing resources. At the same time, the system provides a mechanism for real time collecting forensic data without the risk of shutting down the attacked system or leaking any stored sensitive information.

This infrastructure includes two elements: the ShadowNet client and the ShadowNet server. Different from the traditional honeypot strategy where static honeypot servers are placed in the network to lure attackers, the developed infrastructure will actively deploy a decoy host system only when host is under attack or suspicious behaviors are noticed. This capability is achieved by engaging the attacker with a virtual live clone [13] of the host when the suspicious behaviors are detected.

3.2 Cyber Attack Prevention Description

As shown in Fig.4, when a suspicious insider conducts suspicious behaviors on the network, such as logging in to a protected system where he or she has no

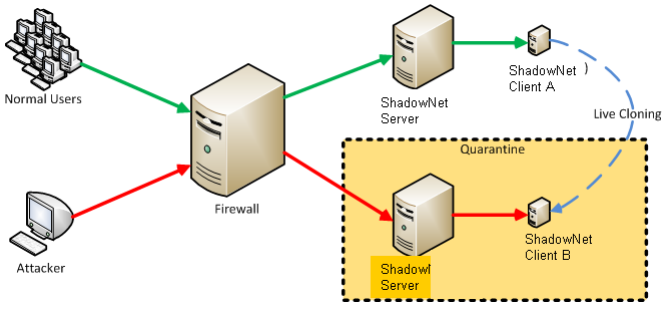

Fig. 2. ShadowNet Active Defense Network Topology

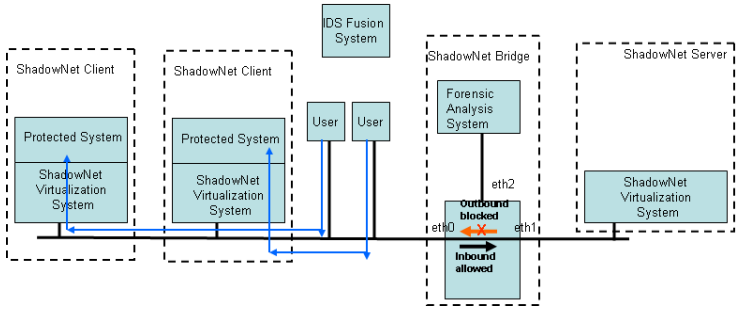

Fig. 3. ShadowNet Active Defense Infrastructure

legitimate access rights or logging into the system at an abnormal time, his behaviors will trigger the intrusion detection sensors installed on the network, which in turn sends out alarm messages to the IDS fusion system. The fusion system sends control messages to the "ShadowNet Client" and "Shadow Server" for system live clone action. As shown in Fig.5, a live clone [11] of the system that is breached by the suspicious insider will be prepared in real-time by the ShadowNet Client. This clone system will have exactly same status, file system structure, and network availability as the original system, but will not contain the original host's sensitive data. All these sensitive data are eliminated and replaced with fraud data that contain useless information. The system clone along with the suspicious network connection will be migrated to another physical server (ShadowNet Server) without noticed by the attacker. Using the live system migration technology, the total time for migrating the attacker to the Server will be within 100 ms. As shown in Fig.5, after the migration, the ShadowNet Client will automatically re-route all network connection package from the suspicious insider to the migrated system clone in ShadowNet Server behind the "ShadowNet Bridge".

The migrated clone has same environment of the original host that the insider breached. Thus, the suspicious user or insider will not aware that his/her

connection has been migrated to a closely monitored forensic analysis platform. A ShadowNet Bridge can be deployed to quarantine the ShadowNet server by preventing the potential of information leak out. The ShadowNet Bridge is special designed ISO second level system that are transparent in the network and allow inbound network connection in but stop all outbound connection initial from behind ShadowNet Bridge. This prevents the protected system clone from being used by the attacker as an ad-hoc attacking machine for attacking other machine. All network transactions and communications to the system clone can be collected by ShadowNet bridge for future forensic analysis.

3.3 Core Technologies

The innovative technology behind the ShadowNet system is a real-time system live clone and migration technology called the ShadowNet live clone. The technology enables the ShadowNet system construct a system clone, and to move the system clone along with the suspicious network connection to another physical server (ShadowNet Server) without being noticed by the attacker. At the same time, the original system and other users of the system are not impacted.

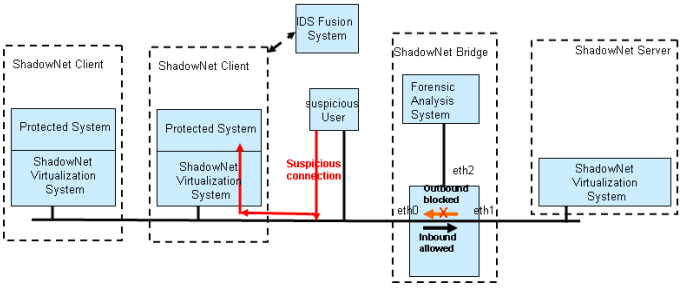

Fig. 4. Suspicious Connections from Insider User's Computer

Different from the traditional visualization technology [12], our system needs the capability to both migrate operating system instances across distinct physical hosts, and also to keep the current operating system continually running to support the normal business. This requires the change of MAC address and IP address of the live virtual clone. Most of the original system network connections are maintained and only the suspicious network connection ports are disabled or disconnected. By modifying Xen Virtual machine platform, we built an experiment implementation to demonstrate the ShadowNet system. As shown in Fig.6. In the implementation, we use a Virtual Machine (VM) Descriptor as a condensed VMimage that allows swift VM replication to a separate physical host. Construction of a VM descriptor starts by spawning a thread in the VM kernel issues a hypercall suspending the VM's execution. When the hypercall succeeds, a privileged process in domain0 maps the suspended VM memory to populate

the descriptor. The descriptor contains: (1) metadata describing the VM and its virtual devices, (2) a few memory pages shared between the VM and the VM hypervisor, (3) the registers of the main VCPU, (4) the Global Descriptor Tables (GDT) used by the x86 segmentation hardware for memory protection, and (5) the page tables of the VM. A monitor and control tool, SXMaster, is developed and is capable of receiving the alert from the existing IDS enclave to trigger the clone migration process that moves the suspicious user's system and network connection to quarantined ShadowNet server. SXMaster uses socket-based communications to execute commands on the ShadowNet infrastructure machines. Each instance of the application listens by default on port 4445 for incoming connections from other instances of the application. The software uses iptables commands to configure NAT, and uses conntrack-tools commands to assist in live session migration.

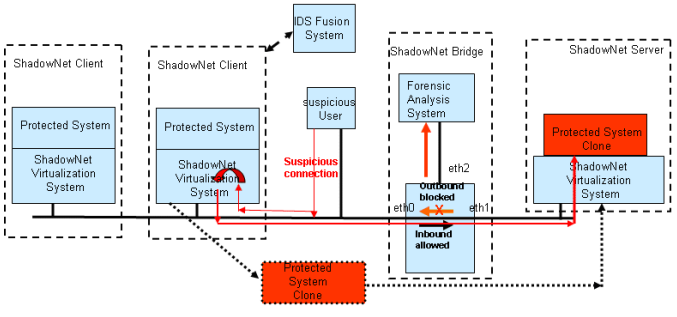

Fig. 5. Active defense for Attack Prevention and Forensics Collections

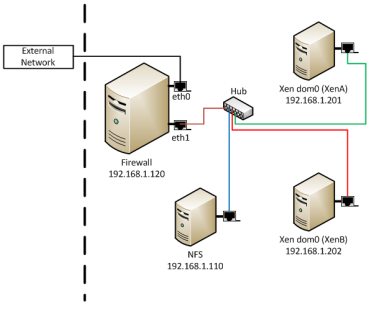

Fig. 6. Experiment Network Topology

4 Conclusion

According to our research, there are currently no COTS systems that can provide effective mechanisms to prevent insider attack. Most of the work associated with

preventing the insider attack focuses on studying security policies and policy enforcement. The ShadowNet technology provides a system that will be able to prevent the attack from suspicious insider who has knowledge about the network structure and the location of the sensitive information. The ShadowNet system also provides a real-time forensics data gathering capability to support large, geographically dispersed networks without disturbing the system's operations. This capability will aid in real-time attack analysis, countermeasures development, and legal prosecution.

Existing information security technologies such as firewalls or Intrusion Detection Systems (IDS) cannot provide an adequate defense against insider threats because they are oriented towards attacks originated from outside the enterprise. Insider attacks may begin from any of numerous potential attack points in the enterprise and have too many parameters to be monitored that existing systems cannot handle. By implementing the ShadowNet infrastructure developed in this research, the whole network becomes a distributed IDS grid. Any machine in the network is a potential honeypot to quarantine a malicious insider.

Acknowledgment. This work was supported in part by the Oak Ridge National Laboratory LDRD program. The views and conclusions contained in this document are those of the authors. This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

1. Salem, M.B., Hershkop, S., Stolfo, S.J.: A Survey of Insider Attack Detection Research. Advances in Information Security 39, 69–90 (2008)
2. The eighth annual CSI/FBI 2003 report: Computer Crime and Security Survey (2003)
3. Stone, C.: Information Sharing in the Era of WikiLeaks: Balancing Security and Collaboration, Office of The Director of National Intelligence, Washington, DC (March 2011)
4. Bellovin, S.: The Insider Attack Problem Nature and Scope. Advances in Information Security 39, 69–90 (2008)
5. Braz, F.A., Fernandez, E.B., VanHilst, M.: Eliciting Security Requirements through Misuse Activities. In: Proceedings of the 2008 19th International Conference on Database and Expert Systems Application (DEXA), pp. 328–333 (2008)
6. Bellovin, S.: There Be Dragons. In: Proc. of the Third Usenix Security Symposium, Baltimore MD (September 1992)
7. Bellovin, S.M.: Packets Found on an Internet. Computer Communications Review 23(3), 26–31 (July)
8. Spitzner, L.: Honeypots: Catching the Insider Threat. In: 19th Annual Computer Security Applications Conference (ACSAC 2003), p. 170 (2003)

9. Spitzner, L.: Honeypots: Tracking Hackers. Addison-Wesley Longman Publishing Co., Inc., Boston (2002)
10. Lyon, G.: Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security. Insecure Publisher, USA (2009) ISBN 9780979958717
11. Sun, Y., Luo, Y., Wang, X., Wang, Z., Zhang, B., Chen, H., Li, X.: Fast Live Cloning of Virtual Machine Based on Xen. In: 2009 11th IEEE International Conference on High Performance Computing and Communications, HPCC 2009, pp. 392–399 (2009)
12. Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., Warfield, A.: Xen and the art of virtualization. In: Proceedings of the ACM Symposium on Operating Systems Principles (October 2003)
13. Clark, C., Fraser, K., Hand, S., Hansen, J.G., Jul, E., Limpach, C., Pratt, I., Warfield, A.: Live migration of virtual machines. In: Proceedings of the 2nd ACM/USENIX Symposium on Networked Systems Design and Implementation (NSDI), Boston, MA, pp. 273–286 (May 2005)