Recovery-based Resilient Cyber Ecosystem

Ajay Nagarajan, George Mason University
Arun Sood, George Mason University and SCIT Labs, Inc.

Abstract. Today's approach to security is largely based on perimeter defense and reactive strategies like Intrusion Detection and Intrusion Prevention Systems (IDS/IPS), firewalls and anti-virus products. Past experience has repeatedly shown us that this strategy is not complete and secure. Intrusion tolerance is an approach that treats intrusions as inevitable and shifts the focus from detection and prevention to containing losses and rapid recovery. We suggest that a complete security strategy is one that does defense in depth and involves both traditional security strategies and intrusion tolerance. Security Information and Event Management (SIEM) is a framework that consolidates the plethora of information available from all of the network and security devices into useful information. In this paper, we propose a stand-alone and a collaborative architecture that makes use of information provided by the SIEM framework to perform adaptive intrusion tolerance in unsupervised learning environments. Resilient systems need to be adaptive, and to achieve this goal we show how environmental information can be used to adaptively change system parameters.

1. Introduction

The variety and complexity of cyber attacks are ever increasing. Verizon's 2012 Business Data Breaches Investigation Report [1] shows that customized malware is difficult to detect and data ex-filtration often occurs over a period of days, weeks and months. The current IDS/IPS approaches are reactive in nature and depend on prior information that is inadequate to prevent all attacks. Events such as the VeriSign security breach [2] and the Playstation Network breach [3] reinforce two notions: 1) even the most sophisticated IDS/IPS systems fail to detect/prevent every intrusion and 2) once the system is compromised, the intruder stays in the system doing damage for extended periods of time.

In addition to the shortcomings of IDS/IPS systems, the costs of operating them are high and increasing. To illustrate the issue we take the example of an enterprise with an average of 1 million raw events occurring per day. About 10,000 alerts are generated by perimeter defense systems. Out of these, 100 alerts are correlated on the basis of severity and other considerations. Assuming it takes 1.5 man-hours to handle one alert, a total of 150 man-hours are required per day to handle alerts generated. The cyber security requires 365 days, 24 hours per day support and in general about 30 people are required to carry out this task. How many large companies can afford such an allocation of manpower? In companies we talk to, only two or three people perform this task. What is worse, 50 % of the alerts are false positives—a tremendous waste of resources. With ever increasing bandwidth and millions of new malware items created every day, these numbers are bound to increase.

Despite years of research and investment in developing such reactive security methodologies, our critical systems remain vulnerable to cyber attacks. The reactive perimeter defense approach relies heavily on threat modeling and vulnerability elimination. We suggest that additional attention should be given to the consequences of a successful attack. In our approach, we focus on limiting the consequences, like reducing the losses that are induced. We believe that we must make our cyber systems more proactive and resilient. Such systems will have the property of (1) supporting continuity of operations—working even in the presence of an intruder; (2) losses, if any, must be limited; (3) systems must resume full operations, i.e. system must be restored to a known good state; and (4) the resilient system operations should be independent of the threat.

To design such a system, we assume that intrusions are inevitable. Therefore, we shift our focus from modeling threats/vulnerabilities to developing methods that will minimize the consequences of an intrusion, increase the work effort of the adversary and increase the visibility of the adversary to the defenders. For this, we have developed a moving target defense approach to computer security. We focus on building mission resilient systems that are able to work through an attack. To ensure reliable operations, the system is restored to a pristine state once every short period of time known as the exposure time, thus negating any malicious action performed by the adversary and minimizing consequences. In addition to this, we use redundancy to provide uninterrupted service and increase overall system availability. The more frequent the computer restoration the less likely it is for the intruder to do damage. The restoration frequency can be random to confuse the adversary and increase his work effort. The shortest time between restorations is a trade-off between available system resources and the throughput of the computer. This intrusion tolerant technology is called Self Cleansing Intrusion Tolerance (SCIT) [4]. The recovery driven approach of SCIT is compared to the detection driven and other intrusion tolerance approaches [5].

Consistent with CrossTalk's theme for the September/October 2012 issue, in this paper, we propose a resilient cyber ecosystem in which every member is able to work together and learn from one another in near-real time to predict and prevent cyber attacks, limit propagation of attacks across participating entities, minimize losses occurring from successful attacks and rapidly recover to a pristine state. To build such a system that is resilient to a variety of sustained attacks, we propose a model that integrates tools and mechanisms that provide protection and detection as well as adaptive tolerance. The rest of the paper is organized as follows: Section 2 provides a brief overview of how SCIT works and motivates the rest of the paper by presenting the need for adaptive SCIT, Section 3 introduces SIEM solutions and presents our idea on how information from SIEM solutions can be used to build adaptive intrusion tolerance systems. We will review two scenarios—stand-alone adaptive intrusion tolerance architecture and a peer-to-peer collaborative intrusion tolerance architecture.

2. How SCIT Works

In [4] we presented SCIT, an intrusion tolerant technique that provides enhanced server security. SCIT research has focused on critical servers that are most prone to malicious attacks. The technique involves multiple virtual instances of servers that are rotated and self-cleansed periodically irrespective of the presence or absence of intrusions. Self-cleansing refers to loading a clean image of the server's OS and application into the Virtual Machine. Rotation here refers to the process of bringing an exposed virtual server off-line, killing it, restarting it and in the meanwhile, bringing another virtual server online to assure availability. By doing so, in the event of an intrusion, the intruder is denied prolonged residence on the server. Once the virtual server's exposure time to the Internet is completed, the virtual server instance is automatically rotated. This virtual instance of the server is what is referred to as virtual server throughout this paper.

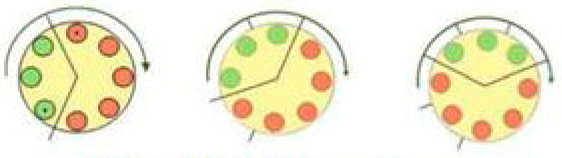

Figure 1: SCIT Server rotation

This illustrative example in Figure 1 shows 3 different time periods. At any given time, there are five servers online and three servers being wiped clean. In each case a different set of servers is being cleaned. Eventually every server will be taken offline, cleaned and restored to its pristine state. SCIT technology can be used to build a variety of servers that meet enhanced security requirements. It is best suited to servers that are designed to handle short transactions—the lower the exposure time the shorter the transaction.

2.1 Need For Adaptive SCIT

Resilient systems have to exhibit adaptive and recovery behavior. SCIT is recovery driven, and in this section we show how SCIT can be made more adaptive to the ongoing changes in the environment.

At any point of time, the resilience of a SCIT system is affected by (1) the current attacks; (2) the current workload; (3) the current data integrity level; (4) the current data availability level; and (5) the current behavior of the system [6]. The first four factors together make up the environment of the SCIT system. Two SCIT systems with different behaviors can yield different levels of resilience. This suggests that as the environment and the behavior of the system changes, the effectiveness of SCIT changes as well. To achieve the maximum amount of resilience, the SCIT system must adapt itself to its environment. Through an architecture for adaptive SCIT, we can (1) adapt SCIT to different application semantics; (2) significantly improve the cost-effectiveness of SCIT; (3) prevent dramatic performance degradation due to system environment changes; and (4) maintain trade-off between system security and system performance [6].

In the case of SCIT, the primary metric is exposure time. In [7], we illustrated the relationship between exposure time and security of a system in terms of data compromised. In [8], we discussed the SCIT approach from the perspectives of effectiveness, tunable parameters, performance impact, and integration to application systems. From the derived expression for Mean Time to Security Failures $MTTSF_{SCIT}$, we were able to conjecture mathematically that decreasing the exposure time window will improve the resilience of a SCIT-based system. To adapt SCIT we will need to adapt the exposure time in response to systems parameters. Increasing $MTTSF_{SCIT}$ would require decreasing the exposure window; hence the cycle that a SCIT server has to go through will become shorter. In this space, there is a tradeoff between system security, performance and cost. Adaptive SCIT could help balance this trade-off in real time with the use of a dynamic exposure time window given the current operating environment and system behavior.

3. Use of SIEM Solutions

"The term SIEM, describes the product capabilities of gathering, analyzing and presenting information from network and security devices; identity and access management applications; vulnerability management and policy compliance tools; operating system, database and application logs; and external threat data" [9].

In addition to receiving inputs from IDS/IPS systems, we will use a SIEM solution to collect and correlate data from all the other sources mentioned in Figure 2 to characterize overall network behavior. This behavioral pattern is then compared with a database of normal network behavior patterns to identify irregularities. Based on the findings of this comparison and the severity of the irregularities, the SCIT controller tunes the "exposure time" of the SCIT-ized system to adapt to the current environment. Similar iterative periodic comparisons will help guide the unsupervised learning and automatic adaption of the SCIT-ized system.

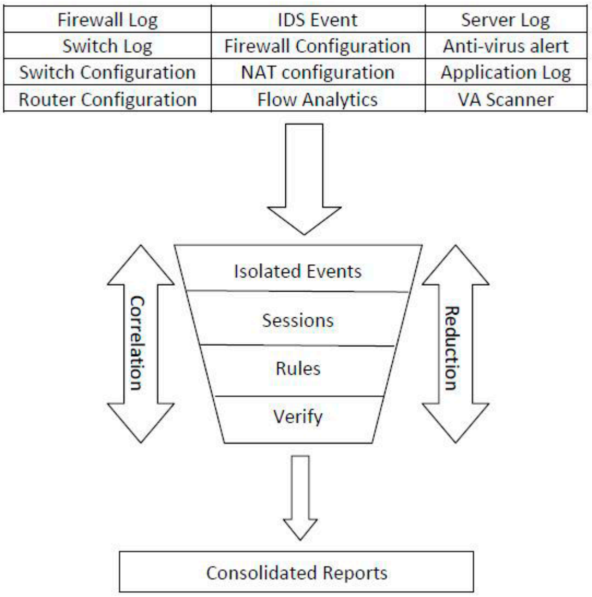

Firewall Log	IDS Event	Server Log
Switch Log	Firewall Configuration	Anti-virus alert
Switch Configuration	NAT configuration	Application Log
Router Configuration	Flow Analytics	VA Scanner

Figure 2: Security Information and Event Management Framework [10]

3.1 Use of Information from SIEM Solutions in Building Adaptive Intrusion Tolerant Systems:

In this section, we expand on the idea of using aggregated information from SIEM solutions to build adaptive intrusion tolerant systems. For the purposes of this paper, SCIT is the intrusion tolerance architecture of choice.

To address the needs outlined in section 2.1, an adaptive SCIT framework must do the following:

1. Employ a dynamic exposure time—the exposure window must keep changing with time as the SCIT environment and the system behavior changes.

2. Constantly receive input from the SIEM framework on the current SCIT environment and state of behavior to make informed alterations to the exposure window.

We present two adaptive SCIT architectures with a common assumption that SCIT is deployed at Enterprise level.

1. Stand-alone adaptive SCIT

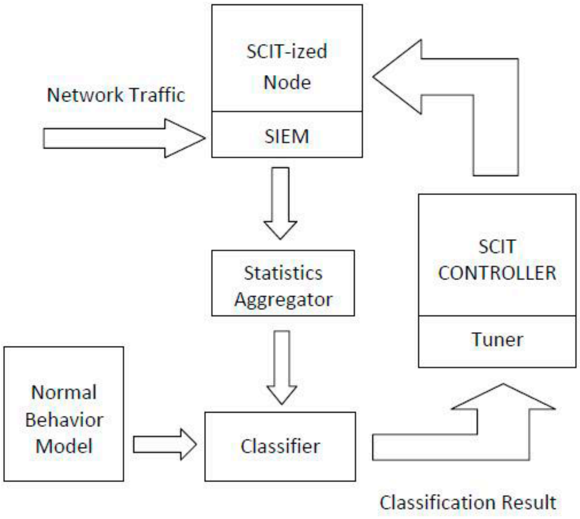

Figure 3: Stand-alone adaptive SCIT

In this architecture, SIEM is constantly monitoring the SCIT-ized node and periodically generates consolidated reports based on the information it has gathered and correlated from varying sources. These reports are fed into the Statistics Aggregator which converts massive information obtained from SIEM into meaningful metrics and their respective values. Further, the classifier compares pre-defined Normal Behavior Model (in terms of metrics and values) with the current values obtained from the Statistics Aggregator. The classifier then feeds the results of the comparison to the Tuner of the SCIT Controller. Based on this, the Tuner makes an informed decision on whether or not to alter the existing "exposure time."

For example, if the results from the classifier identify malicious behavior that points to a Distributed Denial of Service attack, then the SCIT Controller can now reduce the "exposure time" thereby hardening the system against such an attack.

2. Peer-to-peer collaborative SCIT

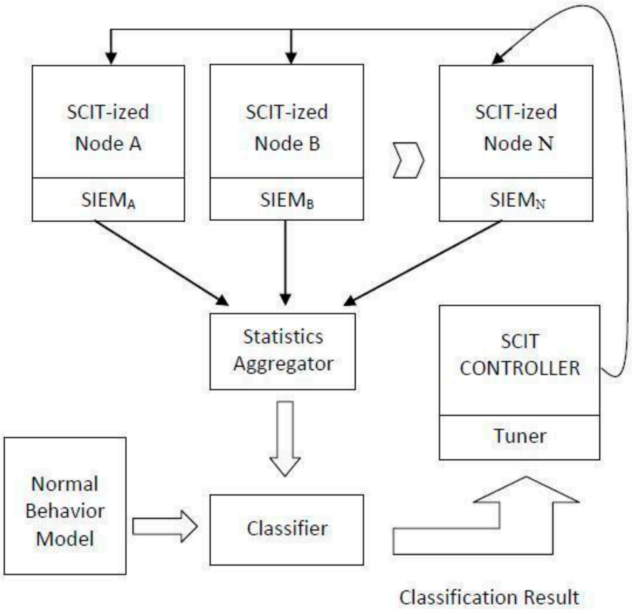

Figure 4: Peer-to-peer collaborative SCIT

This architecture is an extension of the stand-alone architecture. It is meant to mimic a cyber ecosystem with multiple participants in the community that offers recovery-based resilience. In this case, there are 'N' SCIT-ized nodes that are online concurrently. SIEM solutions of each individual node namely SIEM_A, SIEM_B so on till SIEM_N generate reports individually and keep forwarding them to the Statistics Aggregator periodically. The advantages of collaborative SCIT are straightforward:

1. There is more information to work with—the statistics aggregator is now fed with useful information from 'n' different SIEM solutions.

2. Acts as a pre-warning system: malicious behavior in any one of the nodes in the community can now be used to warn/harden the rest of the community.

3. Unsupervised Learning—malicious behavior in any one node in the community can help teach an attack pattern to the rest of the community.

4. Fewer chance of false positives since isolated events now carry less weightage.

4. Conclusion

Cyber attacks are becoming more widespread, sophisticated, and consequential with time. However, detecting, handling and identifying the consequences of an intrusion are still persistent problems. This is partly due to the lack of trust between the members of the cyber ecosystem that impedes information sharing and collaboration. If every entity of the cyber ecosystem were to collaborate with one another and took coordinated security decisions, it could lead to unsupervised learning systems that provide hardened proactive defense.

In this paper, we propose two such recovery based cyber resilient adaptive SCIT architectures. One is a stand-alone system and another is a collaborative system that encourages information sharing and promotes cyber health among communities. In addition to the periodic system self-cleansing done proactively, our system constantly partakes in unsupervised learning from other members of the ecosystem to adapt to the current environment and system behavior.

REFERENCES

1. Verizon Business Data Breach Investigation Report 2012
2. "Key Internet Operators VeriSign hit by hackers" Reuters 02/02/2012
3. "Security Experts: Playstation Network breach one of largest ever" USA Today, 04/27/2011
4. Yih Huang, David Arsenault, and Arun Sood, "Incorruptible System Self-Cleansing for Intrusion Tolerance", Proceedings Workshop on Information Assurance (WIA 2006), Phoenix, AZ, 2006
5. Quyen L. Nguyen and Arun Sood, "Comparative Analysis of Intrusion-Tolerant System Architectures", IEEE Security and Privacy, Volume 9 Issue 4, July-Aug 2011
6. Luenam P. and Peng Liu "The design of an adaptive intrusion tolerant database system" Foundations of Intrusion Tolerant Systems, 2003
7. Ajay Nagarajan and Arun Sood, "SCIT and IDS Architectures for Reduced Data Ex-filtration" 4th Workshop on Recent Advances in Intrusion-Tolerant Systems, Chicago, IL, USA, June 28 2010
8. Quyen Nguyen and Arun Sood, "Quantitative Approach to Tuning of a Time-Based Intrusion-Tolerant System Architecture", 3rd Workshop on Recent Advances in Intrusion Tolerant Systems, Portugal, June 29, 2009.
9. Security Information and Event Management – Wikipedia article
10. CISCO Security Monitoring, Analysis and Response System (MARS) Framework

ABOUT THE AUTHORS

Ajay Nagarajan is currently a Ph.D., candidate in Computer Science at George Mason University working under Dr. Arun Sood. He received his M.S. in Computer Science from George Mason University in 2010. He is affiliated with the SCIT Research group at GMU and his main research interests include Intrusion Tolerance, Survivability and Security Evaluation.

Volgenau School of Information Technology & Engineering
George Mason University, MS 4A5
4400 University Drive
Fairfax, Va. 22030
Phone: 540-687-0363
E-mail: anagara1@gmu.edu

Dr. Arun Sood is Professor of Computer Science in the Department of Computer Science, and Co-Director of the International Cyber Center (ICC) at George Mason University, Fairfax, VA. His research interests are in security architectures; image and multimedia computing; performance modeling and evaluation; simulation, modeling, and optimization.

He and his team of faculty and students have developed a new approach to server security, called Self Cleansing Intrusion Tolerance (SCIT). We convert static servers into dynamic servers and reduce the exposure of the servers, while maintaining uninterrupted service. This research has been supported by the U.S. Army, NIST through the Critical Infrastructure Program, SUN, Lockheed Martin, Commonwealth of Virginia CTRF (in partnership with Northrop Grumman).

Recently SCIT technology was the winner of the Global Security Challenge (GSC) sponsored Securities Technologies for Tomorrow Challenge. This technology has been awarded three patents and three additional patents are pending. SCIT Labs, a university spin-off, has been formed to commercialize SCIT technology. Dr. Sood is the founder and CEO of SCIT Labs.

Since 2009 Dr. Sood has directed an annual workshop on Cyber Security and Global Affairs with Office of Naval Research support – Oxford 2009, Zurich 2010 and Budapest 2011.

Dr. Sood has held academic positions at Wayne State University, Detroit, MI, Louisiana State University, Baton Rouge, and IIT, Delhi. His has been supported by the Office of Naval Research, NIMA (now NGA), National Science Foundation, U.S. Army Belvoir RD&E Center, U. S. Army TACOM, U.S. Department of Transportation, and private industry.

He was awarded grants from NATO to organize and direct advance study institutes in relational database machine architecture and active perception and robot vision.

Dr. Sood received the B.Tech degree from the Indian Institute of Technology (IIT), Delhi, in 1966, and the M.S. and Ph.D. degrees in Electrical Engineering from Carnegie Mellon University, Pittsburgh, PA, in 1967 and 1971, respectively.

His research has resulted in more than 160 publications, and his resume including publications list is available at <http://cs.gmu.edu/~asood>.

Volgenau School of Information Technology & Engineering
George Mason University, 4A5
4400 University Drive
Fairfax, Va. 22030
Phone: 703-993-1524
Fax: 703-993-1710
E-mail: asood@gmu.edu