REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
05-07-2017	Final Report	1-Oct-2011 - 30-Sep-2013

4. TITLE AND SUBTITLE
Final Report: A Laboratory for Cyber Situation Awareness Using Heterogeneous Virtual Machine Replication

5a. CONTRACT NUMBER
W911NF-11-1-0340

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611103

6. AUTHORS
Sushil Jajodia

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
George Mason University
4400 University Drive, MSN 4C6

Fairfax, VA 22030 -4422

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
59418-CS-RIP.2

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT
Under ARO funded MURI project entitled "Computer-Aided Human Centric Cyber Situation Awareness," we at George Mason University and our research partners are developing an integrated end-to-end (spanning the whole life cycle) cyber situation awareness solution to fill the gap between machine information processing and analysts' mental processes. A novel aspect of this project is the virtual machine (VM)-replication based damage assessment, recovery, and service regeneration.

15. SUBJECT TERMS
Cyber situation awareness, virtual machine replication

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Sushil Jajodia
UU	UU	UU	UU		19b. TELEPHONE NUMBER 703-993-1653

Report Title

Final Report: A Laboratory for Cyber Situation Awareness Using Heterogeneous Virtual Machine Replication

ABSTRACT

Under ARO funded MURI project entitled "Computer-Aided Human Centric Cyber Situation Awareness," we at George Mason University and our research partners are developing an integrated end-to-end (spanning the whole life cycle) cyber situation awareness solution to fill the gap between machine information processing and analysts' mental processes. A novel aspect of this project is the virtual machine (VM)-replication based damage assessment, recovery, and service regeneration.

With our existing resources, we were able to build a small test range consisting of a server and multiple client machines to demonstrate the feasibility of our approach. However, a larger test range was needed to get statistically significant measures of performance and security for current efforts, as well as to enable future research and development to demonstrate enterprise-wide scalability of our solutions. To this end, we proposed the acquisition and building of a laboratory for large-scale testing. With this DURIP award, we built as a multi-purpose re-configurable test range via virtualization to support a virtually limitless range of network topologies of clients, servers, and routers found in typical enterprises. The testbed enabled us to realistically assess the efficacy of our research against the type of enterprise systems they are expected to defend.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received	Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME	PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Names of Post Doctorates

NAME	PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Names of Faculty Supported

NAME	PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME PERCENT_SUPPORTED

FTE Equivalent:
Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Under ARO funded MURI project entitled “Computer-Aided Human Centric Cyber Situation Awareness,” we at George Mason University and our research partners are developing an integrated end-to-end (spanning the whole life cycle) cyber situation awareness solution to fill the gap between machine information processing and analysts’ mental processes. A novel aspect of this project is the virtual machine (VM)-replication based damage assessment, recovery, and service regeneration. Damage assessment is part of the comprehension stage while recovery and service regeneration are part of the response stage (see Figure 1). The novelty of our techniques lies in that they are built atop heterogeneous virtual machine replication (H-VM-R), an innovative technology being developed by us. Due to this innovation, recovery and service regeneration could be as fast as live synchronization/migration of VM images.

With our existing resources, we were able to build a small test range consisting of a server and multiple client machines to demonstrate the feasibility of our approach. However, a larger test range was needed to get statistically significant measures of performance and security for current efforts, as well as to enable future research and development to demonstrate enterprise-wide scalability of our solutions. To this end, we proposed the acquisition and building of a laboratory for large-scale testing. With this DURIP award, we built as a multi-purpose re-configurable test range via virtualization to support a virtually limitless range of network topologies of clients, servers, and routers found in typical enterprises. The testbed enabled us to realistically assess the efficacy of our research against the type of enterprise systems they are expected to defend.

Technology Transfer

Final Report
A Laboratory for Cyber Situation Awareness Using Heterogeneous Virtual Machine Replication
ARO Award No. W911NF-11-1-0340

Submitted by

Sushil Jajodia
Center for Secure Information Systems
George Mason University
Fairfax, VA 22030-4422
jajodia@gmu.edu

1. Introduction

Under ARO funded MURI project entitled "Computer-Aided Human Centric Cyber Situation Awareness," we at George Mason University and our research partners are developing an integrated end-to-end (spanning the whole life cycle) cyber situation awareness solution to fill the gap between machine information processing and analysts' mental processes. A novel aspect of this project is the virtual machine (VM)-replication based damage assessment, recovery, and service regeneration. Damage assessment is part of the comprehension stage while recovery and service regeneration are part of the response stage (see Figure 1). The novelty of our techniques lies in that they are built atop *heterogeneous virtual machine replication (H-VM-R),* an innovative technology being developed by us. Due to this innovation, recovery and service regeneration could be as fast as live synchronization/migration of VM images.

With our existing resources, we were able to build a small test range consisting of a server and multiple client machines to demonstrate the feasibility of our approach. However, a larger test range was needed to get statistically significant measures of performance and security for current efforts, as well as to enable future research and development to demonstrate enterprise-wide scalability of our solutions. To this end, we proposed the acquisition and building of a laboratory for large-scale testing. With this DURIP award, we built as a multi-purpose re-configurable test range via virtualization to support a virtually limitless range of network topologies of clients, servers, and routers found in typical enterprises. The testbed enabled us to realistically assess the efficacy of our research against the type of enterprise systems they are expected to defend.

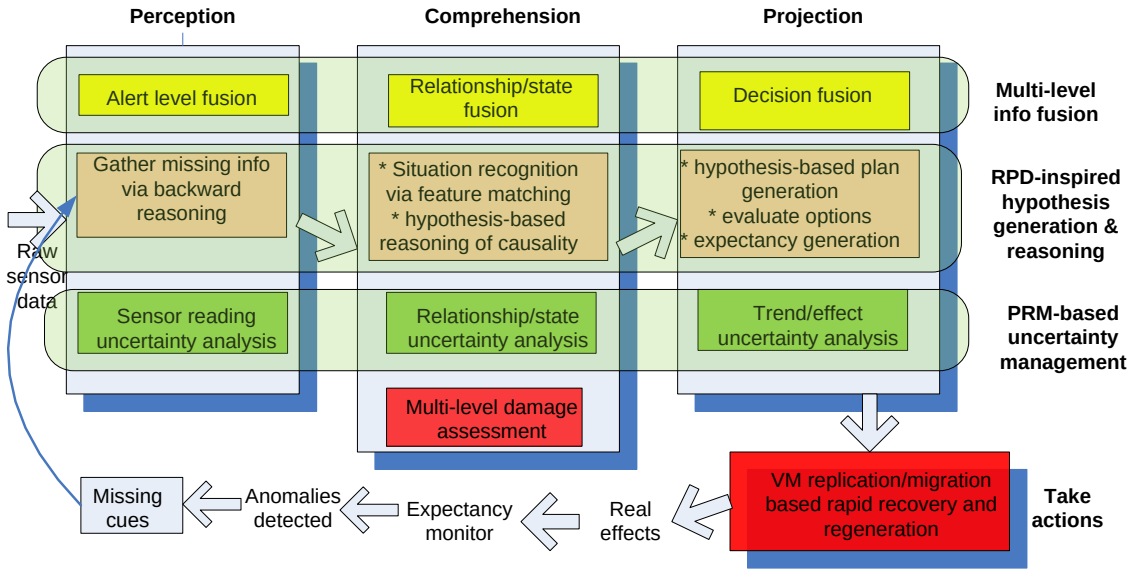

Figure 1: The life-cycle view of the proposed cyber SA framework

2. Overview of Our Approach

Figure 2 shows the basic idea behind the H-VM-R approach we developed under this grant. Unlike homogeneous VM replication, H-VM-R literally *replicates* one virtual machine execution (e.g., Xen) onto a totally different virtual machine (e.g., QEMU, Hyper-V, or VMware) without requiring that these two virtual machines are of the same type. H-VM-R does intrusion detection by comparing heterogeneous VM images resulted from the same execution history. H-VM-R does active response and recovery by proactively setting up standby VM replicas; migration from a compromised VM replica to a clean yet *heterogeneous* VM replica would be instantly performed.

The *objectives* we achieved through our H-VM-R approach are as follows:

- Making redundancy and high-availability practically affordable.
- Transforming microscopic intrusion analysis and detection from pure offline security operations to an *online* capability directly participating in active response.
- Developing an innovative intrusion detection technology based on cross-VM inconsistency checking.
- Achieving fine-grained intrusion detection, response, and recovery.
- Developing a new artificial diversity technology which is simpler, more robust, and less expensive.

The *innovative aspects* of our H-VM-R approach are as follows:

- H-VM-R is the *first* systematic approach that uses heterogeneous VM replication to do intrusion detection, active response, and recovery.
- The cross-VM inconsistency checking based intrusion detection technology is new and the first attempt to integrate replication with intrusion detection.
- H-VM-R is a new artificial diversity technology.
- H-VM-R is the *first* systematic approach that consolidates four areas of systems security research: redundancy, microscopic cyber situation awareness, autonomic response (including recovery), and unpredictability (via randomization and artificial diversity).

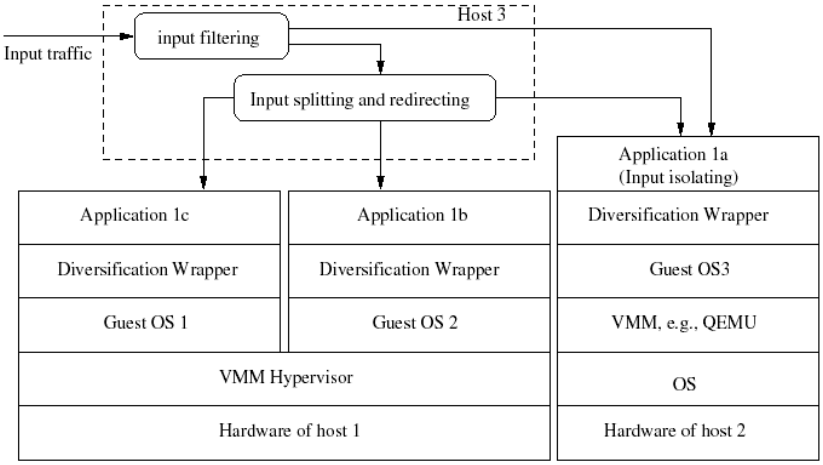

Figure 2: Overview of VM-Replication

3. Laboratory Setup

Under this grant, we built a realistic data center security test-bed. Specifically, we implemented the proposed H-VM-R technology atop existing open source virtual machine monitors such as Xen Hypervisor and QEMU. In order to evaluate the advantage of the proposed VM replication technology, we compared with the performance of the most up-to-date virtual machine technologies, which included VMware Virtual Infrastructure 3 (built upon VMware's highly robust, award-winning products, VMware ESX Server, Virtual Center and V-Motion), and Citrix XenServer software.

As for the evaluation, we used three categories of metrics: *business continuity metrics* (e.g., down time, response time delay), *cost metrics* (e.g., hardware cost, software cost, management cost), and *information assurance metrics* (e.g., system integrity levels). To make the evaluation results representative, we used widely-deployed open source services (e.g., Apache) to build the appropriate benchmarks.

The equipment purchased under this grant was needed to enable the proposed VM replication technology to be thoroughly evaluated in a practical data center setting. Without a realistic data center testing environment, even if the proposed technology is evaluated on real machines, the evaluation results may still fail to serve as an applicable reference or guideline for real world data center managers and administrators to assess the benefits of deploying the proposed VM replication technology.

Also, this equipment was key to provide two realistic VM server clusters that can be used to *comparatively* study the efficacy of the proposed VM replication technique in terms of the three families of metrics mentioned above, namely, *business continuity metrics, cost metrics,* and *information assurance metrics*. Figure 3 illustrates the overall architecture of the testbed we deployed, with two VM clusters.

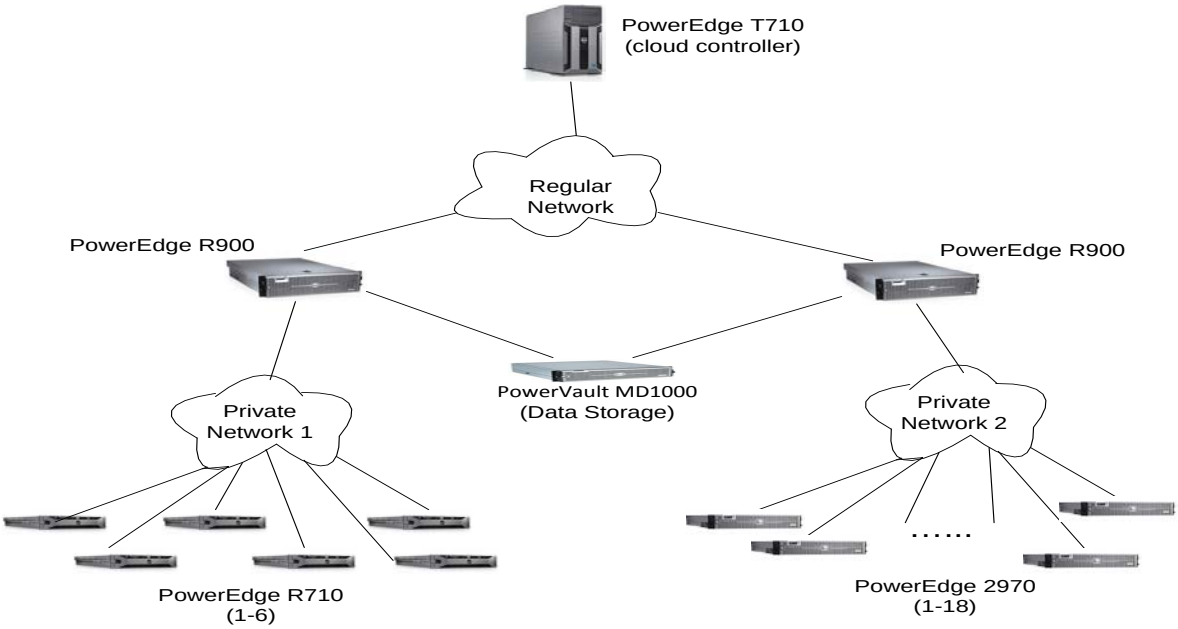

Figure 3. Architecture with Two VM clusters for comparative study of VM replication techniques

4. Publications

Our research resulted in the several publications, some of which are listed below:

- Zhan Wang, Kun Sun, Sushil Jajodia, Jiwu Jing, "Verification of data redundancy in cloud storage," *Proc. ASIACCS International Workshop on Security in Cloud Computing*, Hangzhou, China, May 8-10, 2013.

- William Nzoukou Tankou, Lingyu Wang, Sushil Jajodia and Anoop Singhal, "A unified framework for measuring a network's mean time-to-compromise," *Proc. 32nd Int'l. Symp. on Reliable Distributed Systems (SRDS)*, Braga, Portugal, September 30 - October 3, 2013 (Acceptance ratio 22/67).
- Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, Giuseppe Psaila, Pierangela Samarati, "Integrating trust management and access control in data-intesive web applications," *ACM Trans. on the Web (TWEB)*, Vol. 6, No. 2, Article 6, May 2012, 43 pages.
- Arun Natrajan, Peng Ning, Yao Liu, Sushil Jajodia, Steve E. Hutchinson, "NSDMine: Automated discovery of network service dependencies," *Proc. 31st Annual Int'l. Conf. on Computer Communications (INFOCOM 2012)*, Orlando, FL, March 25-30, 2012, pages 2507-2515 (Acceptance ratio 278/1547).
- Nelson Nazzicari, Javier Almillategui, Angelos Stavrou and Sushil Jajodia, "Switchwall: Automated topology fingerprinting & behavior deviation identification," *Proc. 8th International Workshop on Security and Trust Management (STM 2012), Springer Lecture Notes in Computer Science Vol. ??*, Pisa, Italy, September 12-14, 2012 (Acceptance ratio 20/57).
- Massimiliano Albanese, Alessandra De Benedictis, Sushil Jajodia, Paulo Shakarian, "A probabilistic framework for localization of attackers in MANETs," Proc. 17th European Symp. on Research in Computer Security (ESORICS 2012), Springer Lecture Notes in Computer Science, Vol. 7459, Sara Foresti, Moti Yung, Fabio Martinelli, eds., Pisa, Italy, September 10-14, 2012, pages 145-162 (Acceptance ratio 50/248).
- Zhan Wang, Kun Sun, Sushil Jajodia, Jiwu Jing, "Disk storage isolation and verification in cloud," *Proc. IEEE Global Communications Conference (GLOBECOM)*, Anaheim, CA, December 3-7, 2012.
- Pengsu Cheng, Lingyu Wang, Sushil Jajodia, Anoop Singhal, "Aggregating CVSS base scores for semantics-rich network security metrics,"*Proc. 31st International Symposium on Reliable Distributed Systems (SRDS 2012)*, Irvine, California, October 8-11, 2012.