

Design of Future Navy Warships: A Method For The Development of Security Requirements Within An Advanced Computing System

Kevin R. Smith
Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943-1037
831-656-4679
krsmith@nps.navy.mil

Peter A. Wu
Department of Computer Science
Naval Postgraduate School
Monterey, CA 93943-1174
831-656-4679
pawu@nps.navy.mil

ABSTRACT

Taking new technologies, which are incorporated in advanced computers, advanced networks, and quality-of-service middleware, designers have focused the larger percentage of their time and energy into improving performance in newly designed shipboard automated information and weapon systems. To meet mandated Navy performance-level goals in support of combat readiness, some research and development teams have oversimplified security considerations. Consequently, research test beds have become unrealistic in regard to the real world security threats. Hence from a security perspective, improvements in the architecture of a shipboard heterogeneous computing system need to address threats, which impact the availability, integrity and confidentiality of system resources. The position statement that follows recommends the need for the creation of a method for development of security requirements for the Total Shipboard Computing Environment envisioned in future Naval Warship Design.

1. POSITION STATEMENT

On a brisk October morning in the year 2008, the USS Zumwalt (DD 21), the first 21st Century Land Attack Destroyer of its class, got underway from a pier at the Bath Ironworks shipyard for the first in a series of sea trials. After reaching open ocean during the transit out to sea, the Commanding Officer (CO) ordered the Officer-of-the-Deck (OOD) to gradually increase the ship's speed in increments from the current 1/3-bell (5 knots) to a flank-bell (30 knots) in order to begin a set of full-power trials. As the ship's speed through the water increased with each successive speed change, the bridge crew felt a rush of pride and excitement - this was a special day that many onboard had been anticipating for months. With breakthroughs in technology, such as the design of the electric propulsion drive, the integrated ship's power grid, the large collection of weapons with unimaginable capabilities, and a total shipboard-computing environment tying everything together, the ship seemed invulnerable to every known external threat.

Suddenly, there was a loud "boom" followed by a winding down noise. The entire ship went dark. The generators and gas turbine engines spun offline disabling the propulsion drive, and within minutes the ship found itself dead in the water.

The CO's first expectation was to receive a call from the Chief Engineer explaining what casualties occurred in the engineering plant, but instead he received a call from the Combat Systems Officer (CSO). "What happened?" asked the CO.

"Sir," the CSO replied in a worried tone, "we've just been hacked!"

From the scenario above, one may wonder how much time the research and development teams spent on security threat analysis and how involved they were in finding ways to protect the network against all known attacks. Quite frequently a lack of emphasis in the area of security has led to the deployment of functional but vulnerable computing system solutions. The vulnerabilities within these solutions usually are not a major concern for the implementers until an unauthorized source exploits these vulnerabilities and the exploitation becomes known. When exploitation by an unauthorized source is identified, security becomes a high visibility organizational concern. Hence, the requirements engineering for information security within computing systems should not be something that is left for discovery at a later date or briefly discussed and planned during a project wrap-up. Instead, security requirements need to be thoroughly defined in parallel with other requirement analyses.

To meet performance level goals in their design, some research and development teams may oversimplify the security environment making their research test beds unrealistic with regard to the real world. One example of interest to the authors of this paper is the High Performance Distributed Computing (HiPer-D) project, which is being supported by both the Defense Advanced Research Projects Agency (DARPA) and the Navy's AEGIS program team at the Naval Surface Warfare Center, Dahlgren, Virginia. Taking technologies, which are incorporated in advanced computers, advanced networks, and quality-of-service middleware, the engineers involved with HiPer-D have focused the larger percentage of their time and energy into improving performance characteristics [4][5][7]. These improvements include new technological capabilities such as large-scale throughput of continuously refreshed data, meeting hard real-time deadlines, enabling soft real-time processing, and maintaining high availability and survivability of data. With

regard to the system architecture, which is built upon a pool of heterogeneous resources, concepts such as dynamic resource management and the allocation of distributed processing tasks have become new paradigms of the HiPer-D program [1][2][3][6].

The driving force behind these technological advances lies within the Navy's requirement to maintain a high state of combat readiness at all times. Therefore, acquiring a flexible, low-cost system that provides continuous, real-time, sensor data and reliable continuity to fire control and navigation systems is a goal that benefits all Navy platforms. From a performance point-of-view, the test bed designed to evaluate the new incorporated technologies thus far has proven effective. However, from a security perspective improvements need to be made in terms of protecting the system from the vast array of viruses, Trojan horses, worms, casual hackers, and other attacks introduced by the real world. Although the test bed is somewhat mature, it is still viewed as a prototype; therefore, because the system is still in the prototyping phase, any desired security related changes to the system architecture might be incorporated during system implementation. Regardless, an architectural approach needs to be applied, which integrates security requirements at the earliest possible development phase. Although security is presently viewed as a non-functional property of a system, it is nonetheless, an important requirement that must be implemented if the customer organization's goals associated with confidentiality, availability, and integrity are to be met by the operational system.

To meet the goals of a highly effective and secure system, the authors plan to initiate a study of the security requirements associated with the HiPer-D test bed simulated systems. The approach needed for analyzing, specifying and testing requirements will entail traditional engineering techniques. For example, our efforts will include the following: understanding the system security policy, performing a vulnerabilities assessment, initiating a common criteria analysis, conducting a threat assessment, developing and performing a user needs assessment for the operational system, and the generation of a security requirements document that reflects and appreciates overall system requirements to include the consideration of functional requirements as a constraint on security.

It is important to note that security requirements exist as a management tool. In essence, this tool assists in the identification and prioritization of project implementation decisions, which impact the evolution of a reasonably secure computing system. The effectiveness of this tool may be limited by the importance placed upon security in relation to the computing system and the data, which is generated and residing on the system.

When weighing performance requirements with system security requirements a balance must be defined and preserved. It is of little value to a user of a system that has an extremely high performance rating, if resource availability is sporadic due to control or manipulation of the system by an unauthorized source. It is also of little value to a system user if resource availability is significantly degraded due to over tasking in support of system security requirements.

In summary, security requirements within computing systems should not be left for discovery at a later date or briefly discussed

and planned during a project wrap-up. They need to be thoroughly defined in parallel with other requirement analyses in order to reduce cost and produce a system, which possesses a proper balance between other functionality and security.

2. REFERENCES

- [1] D. A. Hensgen, T. Kidd, D. St. John, M. C. Schnaidt, H. J. Siegel, T. D. Braun, M. Maheswaran, S. Ali, J. Kim, C. E. Irvine, T. Levin, R. F. Freund, M. Kussow, M. Godfrey, A. Duman, P. Cariff, S. Kidd, V. Prasanna, P. Bhat, and A. Alhusaini. An Overview of MSHN: The Management System for Heterogeneous Networks. In *Proceedings of the Eighth Heterogeneous Computing Workshop (HCW '99)*, pages 184-198, San Juan, Puerto Rico, April 1999.
- [2] R. Wright, D. J. Shifflett, and C. E. Irvine. Security Architecture for a Virtual Heterogeneous Machine. In *Proceedings of the Fourteenth Computer Security Applications Conference*, pages 167-177, Phoenix, AZ, December 1998.
- [3] T. Levin and C. E. Irvine. Quality of Security Service in a Resource Management System Benefit Function. Technical Report NPS-CS-00-02, Naval Postgraduate School, Monterey, CA, November 1999.
- [4] J. Kim, D. A. Hensgen, T. Kidd, H. J. Siegel, D. St. John, C. E. Irvine, T. Levin, N. W. Porter, V. K. Prasanna, and R. F. Freund. A QoS Performance Measure Framework for Distributed Heterogeneous Networks. In *Proceedings of the Eighth Euromicro Workshop on Parallel and Distributed Processing*, pages 18-27, Rhodes, Greece, January 2000.
- [5] P. M. Irely IV, B. L. Chappell, R. W. Hott, D. T. Marlow, K. F. Donoghue, and T. R. Plunkett. Metrics, Methodologies, and Tools for Analyzing Network Fault Recovery Performance in Real-Time Distributed Systems. In *Parallel and Distributed Processing Proceedings, 15 IPDPS 2000 Workshops*, LNCS 1800, pages 1248-1257, Cancun, Mexico, May 2000.
- [6] L. R. Welch, M. W. Masters, L. A. Madden, D. T. Marlow, P. M. Irely IV, P. V. Werme, and B. A. Shiazi. A Distributed System Reference Architecture for Adaptive QoS and Resource Management. In *Parallel and Distributed Processing Proceedings, 11 IPPS/SPDP '99 Workshops*, LNCS 1586, pages 1316-1326, San Juan, Puerto Rico, April 1999.
- [7] B. L. Chappell, D. T. Marlow, P. M. Irely IV, and Karen O'Donoghue. An Approach for Measuring IP Security Performance in a Distributed Environment. In *Parallel and Distributed Processing Proceedings, 11 IPPS/SPDP '99 Workshops*, LNCS 1586, pages 389-394, San Juan, Puerto Rico, April 1999.