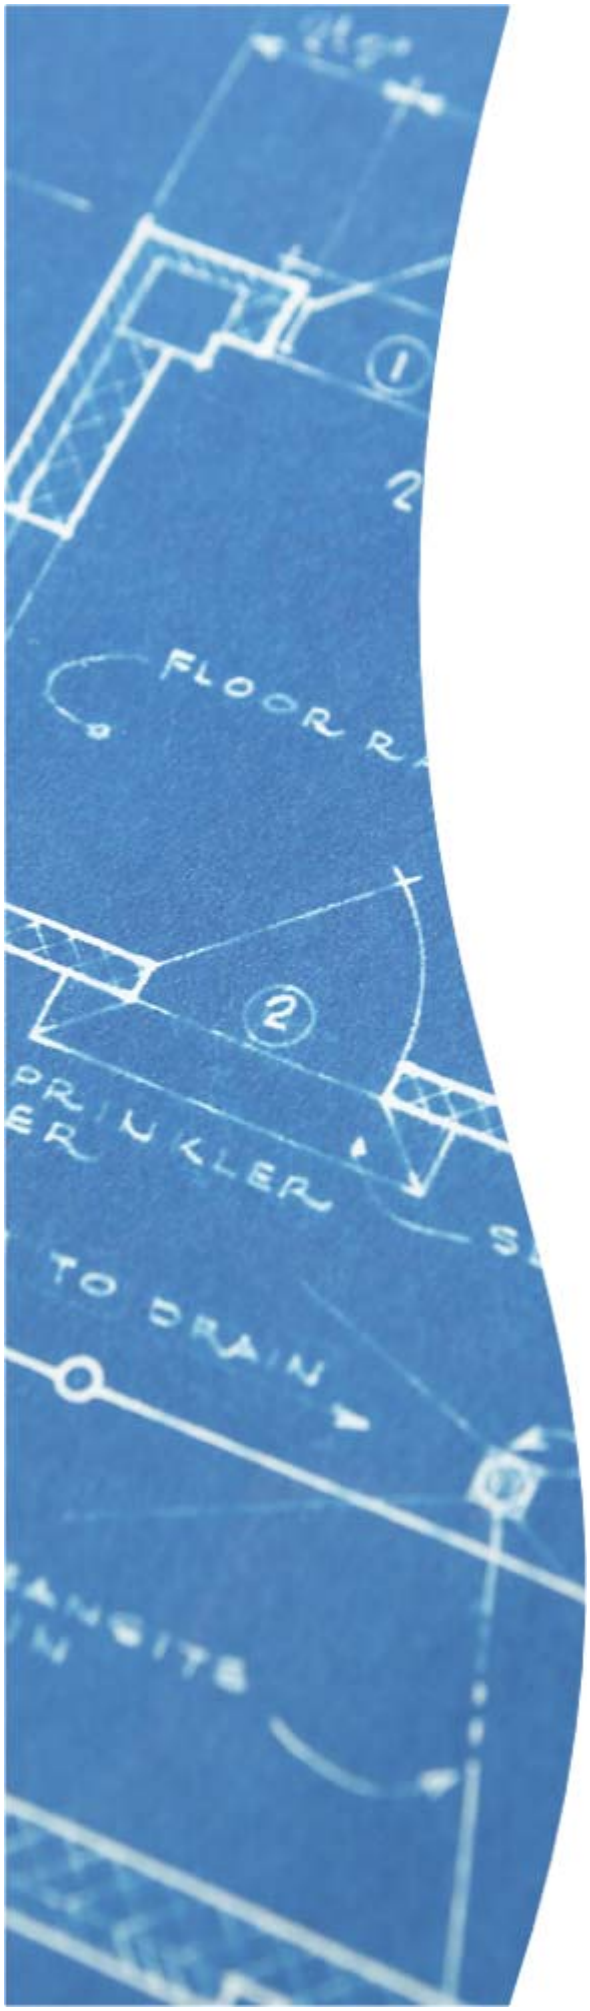

DESIGNING AN ADAPTIVE SECURITY ARCHITECTURE

Joel Weise, Sun Global Systems Engineering Security Office

Sun BluePrints™ Online

Part No 820-6825-10
Revision 1.0, 11/3/08

Table of Contents

Designing an Adaptive Security Architecture

Security threats can reduce the functionality, performance, availability, and integrity of IT systems. These systemic qualities are critical — so much so that they are typically instantiated formally into Service Level Agreements (SLAs). In IT environments, the desire is to reduce potential security threats at least to the degree by which SLAs can be satisfied. This article discusses the principles and characteristics of a new architectural approach — *adaptive security* — which is based in part on the concepts and principles of complex adaptive systems.

The task of security architecture is closely tied to the concept of risk management. Like any security model, adaptive security seeks to contain risk and meet SLA requirements. By countering an attack in a timely fashion, adaptive security strives to reduce the impact and magnitude of potential threats. This includes the possibility of responding to "zero-day" attacks in which a threat is so new that a patch or other countermeasure does not yet exist. Although adaptive security measures continue to evolve, an adaptive approach can be implemented at least in part using technologies available today.

This article introduces concepts associated with adaptive security and how the approach enhances system survivability. It explores the biological and ecological origins of adaptive security, discusses why the approach is useful, reviews its characteristics and principles, and introduces a recommended design approach. The article addresses the following topics:

- "Complex Adaptive Systems in Security Design" on page 2 discusses the impact of complexity and introduces biological complex adaptive systems.
- "Objectives of Adaptive Security" on page 3 describes the goals of implementing an adaptive security model.
- "An Architectural Approach Using Adaptive Security" on page 5 discusses parallels between complex adaptive systems in human ecosystems and in the IT infrastructure.
- "Design Approach to an Adaptive Security Model" on page 10 introduces fundamental steps to prepare for implementing an adaptive security model.

This Sun BluePrints article assumes a basic working knowledge of contemporary security issues. It introduces biological parallels with complex adaptive systems, and discusses common objectives of security architectures and responses to active threats.

Complex Adaptive Systems in Security Design

Complexity is the biggest roadblock in designing secure IT architectures. Dan Geer et al summarize the impact of complexity in deploying secure IT systems:

> "The central enemy of reliability is complexity. Complex systems tend to not be entirely understood by anyone. If no one can understand more than a fraction of a complex system, then, no one can predict all the ways that system could be compromised by an attacker. Prevention of insecure operating modes in complex systems is difficult to do well and impossible to do cheaply: The defender has to counter all possible attacks; the attacker only has to find one unblocked means of attack." [1]

A. Elkhodary et al agree that complexity is a major issue in effectively combatting security threats. The authors also note that "...one possible solution to the increased complexity of IT security infrastructure is adaptive security."[2] Adaptive security takes a twofold approach in addressing the issue of complexity: at the microscopic level, it uses autonomic systems that mimic biologic auto-immune systems, and at the macroscopic level, it uses ecosystem behaviors of disparate entities that resemble complex adaptive systems.

What is a complex adaptive system? There are multiple definitions. The term originates from interdisciplinary efforts at the Santa Fe Institute (SFI) by John H. Holland and others. Since this article focuses on IT systems, the following definition is based, in part, on Holland's work: [3]

> A complex adaptive system is a dynamic network of multiple dispersed and decentralized agents that constantly interact and learn from one another. Any coherent behavior in the system arises from agent interaction.

A security architecture that mimics a complex adaptive system can be effective in that it can adjust and respond to constantly emerging and changing security threats. An adaptive model is designed to counter these issues:

- As system complexity increases, security and integrity decrease.
- A monoculture of systems (quite common in today's IT infrastructures) allows a pandemic to spread quickly.
- Offensive viruses, worms, and adversarial attacks are created and spread faster than the development of effective defensive responses.

---

1. Daniel Geer, Daniel; Bace, Rebecca; Gutmann, Peter; Perry Metzger, Perry; Pfleeger, Charles P. ; Quarterman, John S. , and Schneier, Bruce. *CyberInsecurity: The Cost of Monopoly*, Computer & Communications Industry Association, September 24, 2003.
2. Elkhodary, A. Whittle, J. Inf. & Software Eng., *A Survey of Approaches to Adaptive Application Security*, Software Engineering for Adaptive and Self-Managing Systems, Minneapolis, MN: ICSE Workshops SEAMS '07, 2007.
3. Waldrop, M. Mitchell. C*omplexity: The Emerging Science at the Edge of Order and Chaos,* New York: Touchstone, 1992.

Objectives of Adaptive Security

For IT infrastructures, an adaptive security approach tries to contain active threats and to neutralize potential attack vectors. As with other security architectures, the model strives to:

- Reduce threat amplification — restrict the potential spread of a pandemic in a monoculture (i.e., reduce cascading failures)
- Shrink the attack surface — make the target of an attack smaller
- Decrease attack velocity — slow the rate of attack
- Reduce remediation time — respond to an attack quickly
- Facilitate the availability of data and processing resources — prevent or contain attacks that try to limit resources
- Promote the correctness of data and the reliability of processing resources — respond to attacks intended to compromise data or system integrity

Beyond supporting SLAs, the primary objectives of adaptive security and other security models are to maintain system and data integrity, promote trustworthiness, and provide assurance. Ultimately, a security model seeks to instill confidence in data and processing resources — namely, to ensure that they are trustworthy, reliable, available, and operating within acceptable parameters.

Integrity, Trustworthiness, and Assurance

Integrity, trustworthiness and assurance must work in conjunction in a synergistic manner. These cannot act independently or in standalone modes because of the level of inter-dependence that a comprehensive security architecture requires of such characteristics. To promote continuous business operations that meet SLAs, an adaptive security model tries to promote integrity, trustworthiness, and assurance in IT solutions:

- *Integrity.* Integrity is critical to the functionality of individual systems as well as to the IT infrastructure and entire enterprise. Integrity is a property that applies to both data and systems:
  - Data integrity is the property that data has not been altered in an unauthorized manner while in storage, during processing, or while in transit.
  - System integrity is the quality of performing an intended function in an unimpaired manner, free from unauthorized exposure or manipulation.

  Integrity as a systemic quality is something that enables confidence, i.e., some measurable degree of assurance that one's IT infrastructure is installed, configured, and operating correctly. Services are reliable and available; transactions and interactions with employees, partners, customers, etc., are timely, correct, and consistent.

Integrity is a characteristic of a well-designed, well- implemented, and well-managed infrastructure. In an adaptive security model, designing for integrity means providing applications, systems, networks, and IT infrastructures with the ability to self-configure, self-detect, self-protect, self-optimize, and self-heal to protect against corruption of data and processing resources.

- *Trustworthiness*. A primary tenet of adaptive security is trustworthiness, meaning that systems and the data that they process are reliable. In short, users that rely upon IT system results to make business and other decisions will only do so if the information is accurate and the systems are available when required. Systems and data that encourage integrity are considered trustworthy. An adaptive security architecture helps to ensure integrity, enhancing a system's trustworthiness.
- *Assurance*. When an infrastructure exhibits integrity, it also typically complies with some assurance criteria (i.e., things work as advertised and in accordance to agreed upon requirements). Assurance provides a basis for confidence that security measures, both technical and operational, are working as intended to protect the system and the information it processes. The security objectives of integrity, availability, confidentiality, accountability, and assurance are adequately met for an implementation when:
  - the required functionality is present and correctly implemented
  - there is sufficient protection against unintentional errors (by users or software)
  - there is sufficient resistance to intentional penetration or by-pass

  Assurance is essential — without it other security objectives are not met. However, assurance is a continuum, and the amount of assurance needed varies between systems. Assurance highlights the need for a system to not only provide intended functionality, but also to make sure that undesired actions do not occur, or can be detected if they do.

Collectively the objectives of integrity, trustworthiness, and assurance are addressed by automatic detection of attacks in progress, dynamic threat response, systemic quarantine of impacted IT elements and processes (for forensic purposes), and automatic recovery (e.g., bringing in new code instances to replace those quarantined to preserve functionality and availability).

Survivability

One of the primary tenets of adaptive security is survivability. In fact, survivability is really the ultimate goal of adaptive security. Survivability is defined as the "capability of a system to fulfill its mission in a timely manner in the presence of attacks, failures, or accidents".[4] To ensure a system's survival, it is necessary first to identify the essential system elements (what must survive) versus elements that are deemed sacrificial.

4. Taylor, C. and Alves-Foss, J. *Attack Recognition for System Survivability: A Low-level Approach,* Proceedings of the 36th Annual Hawaii International Conference on System Sciences, 2003.

This analysis ties into the notion of "self" and "non-self." (See "Determining 'Self' versus 'Non-Self'" on page 8.) For the sake of discussion, a system is considered to have survived if it continues to fulfill its business purpose within expected Service Level Agreements.

An Architectural Approach Using Adaptive Security

Biological and ecosystem metaphors provide interesting parallels for issues, threats, and countermeasures related to IT system security. Biological and ecological systems maintain integrity and survivability by reacting to known attacks, adapting to unknown threats, or dying. Responses can be at a microscopic biological level, (e.g., molecular or cellular), or at a macroscopic or ecosystem level (e.g., system or species).

As Darwin tells us, organisms must adapt or die. To extend the parallel between biological ecosystems and IT ecosystems, successful IT infrastructures must adapt or they will eventually fail due to predator attacks, viral infections, or the inability to adjust to environmental changes. According to the IBM System Journal:

> "...by enabling computing systems to make their decisions in consistent and reliable ways, autonomic techniques will engender an extremely adaptive and dynamic operational style...." [5]

It is this "adaptive and dynamic operational style" that describes robust adaptive security solutions that can help to address the integrity of data and system resources and survivability.

Over time, the behavior of human immune systems has evolved from reactive to adaptive:

> "The adaptive immune response provides the vertebrate immune system with the ability to recognize and remember specific pathogens (to generate immunity), and to mount stronger attacks each time the pathogen is encountered. It is adaptive immunity because the body's immune system prepares itself for future challenges." [6]

Consider an IT infrastructure as an ecosystem comprised of many elements (systems, storage, networks, applications, etc.), similar to a natural ecosystem (with components such as people, food, air, and environmental conditions like temperature, humidity, etc.). One can draw similarities in how a localized disease outbreak or even a pandemic affects these ecosystems, as well as how they respond against such threats. In a natural human ecosystem, when a pandemic strikes, not every human survives. Survival depends on a number of factors such as genetic make-up and environmental

5. Chess, D. M.; Palmer, C. C.; White, S. R. *Security in an Autonomic computing Environment*, IBM System Journal 2003, Volume 42, Number 1, 2003. (http://www.research.ibm.com/journal/sj/421/chess.pdf)
6. http://en.wikipedia.org/wiki/Adaptive_immune_system, October 9, 2008.

conditions. Some individuals have the genetic material and environment that enable them to survive, while others do not. In the natural ecosystem, although some individuals perish, the ecosystem as a whole survives.

Using an adaptive approach, the IT infrastructure can function as an autonomic system that effectively mimics both an organic immune system and a large-scale natural ecosystem. It is possible to design an architecture such that when a virus, worm, or other threat strikes, individual components can be sacrificed for the survival of the entire ecosystem. An IT infrastructure designed for survivability exhibits the following characteristics:

- The flexibility to respond adaptively to new and different threats
- The ability to be self-detecting, self-regulating, self-healing, and self-protecting
- A foundation on a standardized security model with enforcement mechanisms that compel security policy compliance
- The ability to learn about ecosystem norms, to detect unauthorized resource modifications (e.g., to data, files, file systems, operating systems, and configurations), and if necessary to take remedial actions such as:
  - quarantine of resources for forensic purposes so that the ecosystem can learn from the breach
  - provisioning of other resources to replace affected systems, enabling service continuity
  - application of corrective measures as needed

Immune Response Mediators

Part of the human immune system includes immune response mediators (e.g., T-cells). Implementing the functionality of immune response mediators is an important component in an adaptive security approach. The role that such mediators play in the IT infrastructure is almost identical to how they are used in the biological sense — they are guardian agents or sentinels that are deployed throughout the infrastructure that act as sensors to identify potential threats. In conjunction with threat response and feedback mechanisms, these sentinels moderate the immune response as they do in a biological system. When a threat is detected, they act as “threat triggers,” contacting the appropriate threat responders and identifying the potential threat. Such sentinel functionality can be directly incorporated into various application, network, and operating system components so that independent intrusion detection systems (IDS) or firewall systems are no longer necessary.

Determining “Self” versus “Non-Self”

For systems to identify threats effectively, they must understand a baseline of what is considered normal behavior and what is not. To continue the biological metaphor, the concepts of “self” and “non-self” are fundamental in auto-immune systems. Functional auto-immune systems discriminate properly between what is native to the organism

and what is not, and what is not native is treated as a threat and eliminated. Mimicking this notion, an IT system can autonomically protect itself by correctly identifying and countering threats and suspicious activity, and distinguishing these from acceptable components, protocols, and operational processes within the IT infrastructure.

Biological and IT Security Principles

Adaptive security leverages architectural and operational principles from a variety of disciplines. The following principles map different properties of biological and ecological systems that are applicable to information systems. De Castro et al have likewise identified these as beneficial characteristics that are useful in IT systems to reduce exposure to threats, contain the magnitude of threats, and counter them in a timely fashion.[7]

- *Pattern recognition.* In the biological world, cells recognize various proteins via surface pattern matching. Likewise, IT systems must be capable of sophisticated pattern matching techniques to identify normal and abnormal behavior in code, command/response dialogs, communication protocols, etc.
- *Uniqueness.* Biological organisms possess their own unique immune system that varies from individual to individual. This uniqueness is associated with different individual strengths and weaknesses. Expressed in IT systems, uniqueness discourages the existence of monocultures that can be susceptible to a common computer virus. Uniqueness endows an ecosystem of different IT systems with the necessary robustness to survive targeted threats.
- *Self identity.* The notion of self and non-self allows an organism to comprehend what is native and what is not, and triggers elimination of non-self entities that are considered a threat. The IT world mimics this concept by isolating and eliminating what does not belong according to baseline manifests and security policy. Part of implementing this concept includes support for intra/inter-systems communication and sharing information on threats, countermeasures, security policies, and trust relationships between systems and IT infrastructures.
- *Diversity.* In the biological world diversity refers to the different types of elements (proteins, cells, etc.) that together exhibit a wide range of defenses against different threats, including innate and adaptive immune responses. In IT systems diversity manifests itself through different control mechanisms such as compartmentalization via operating system virtualization or Trusted Platform Module (TPM)-based hardware trust anchors.

7. de Castro, Leandro Nunes, and Timmis, Jonathan. *Artificial Immune Systems: A New Computational Intelligence Approach*, London: Springer-Verlag, 2002.

- *Disposability.* With respect to the immune system, disposability refers to the notion that no single cell or molecule is essential to the entire organism. A sacrificial IT system — a system or virtual machine instance that can be eliminated if necessary — represents the concept of disposability in an IT infrastructure. Disposability enables flexibility that contributes to the overall robustness of the infrastructure.
- *Autonomy.* Autonomy in biological systems means that there is no single element that controls the immune system — different elements in the immune system function autonomously to counter threats. In the IT infrastructure, it is likewise desirable for security and integrity control mechanisms to function in an autonomous fashion to address threats.
- *Multilayered*. In biological organisms, molecular, cellular, and other elements cooperate to provide a comprehensive threat response capability. This is similar to the notion of "defense-in-depth" that a well-designed security architecture maintains, implementing multiple security measures to overcome the risk of a single measure's compromise.
- *No secure layer.* In the biological world, any and all cells in an organism are at risk of being attacked at any point in time. This reality has an exact parallel in the IT world and is the underlying assumption in any security policy. This assumption is instantiated via a "deny all" security policy that grants access on only a need-to-know basis, and is the foundation of a "defense-in-depth" approach in which multiple security layers are implemented to provide a comprehensive defensive strategy.
- *Anomaly detection*. In biological immune systems the notion of non-self enables the immune system to recognize and respond to external entities. Likewise, an IT system must support the capability to recognize and respond automatically to abnormal behavior or known threats. The intention of using an adaptive approach to security design is to anticipate threats before they manifest themselves.
- *Dynamically changing coverage*. Biological immune systems have limitations on the number and type of cells and molecules that can detect and respond to pathogens. As such, they maintain a dynamically changing set of these cells and molecules in the hope that the correct mix exists to respond to potential threats. In an IT infrastructure, there are similar limitations for the number of threat signatures and threat response mechanisms. Thus a security design must include a means to intelligently predict and anticipate what threat response mechanisms must be deployed at any point in time.
- *Distributivity.* Different elements of biological immune systems are widely distributed throughout an organism and are not under any centralized control. In an IT infrastructure this distributivity reduces the attack surface, narrowing exposure.

- *Noise tolerance*. Biological immune systems do not require an absolute match to recognize pathogens. In the IT world it is also desirable to recognize threats without an absolute match to a threat signature.
- *Resilience.* Although various conditions can reduce the effectiveness of a biological immune system, the system maintains a level of resilience that allows it to continue recognizing and countering pathogens. An IT system must similarly be resilient so that it continues to function in spite of reduced capacity.
- *Fault tolerance.* Biological immune systems are composed of redundant elements that function in a complementary fashion. In addition, elements can be modified to respond to pathogens to which they normally would not respond. In an IT infrastructure, fault tolerance enables some threat response mechanisms to be retooled or modified to respond to threats to which they normally otherwise would not respond.
- *Robustness*. In the biological world robustness is really the aggregate benefit of diversity and distributivity. In the IT world, it makes sense that the infrastructure and systems also exhibit robustness.
- *Immune learning and memory*. In the biological world the immune system is by definition adaptive in nature. This adaptiveness allows for faster and more effective responses that improve over time as the immune system learns and retains pathogen memories. In a security architecture, it is desirable to mimic this adaptiveness and, in particular, to learn and remember threats over time. Many malware and virus detection vendors update signature databases over time for this reason.
- *Predator-prey pattern of response*. Biological immune systems respond to pathogens via a mediated response mechanism. This allows them to scale up the response as the number of pathogens increases. Such a mediated response mechanism is likewise necessary in IT environment so that the appropriate level of threat response controls can be brought to bear. Triggering and feedback mechanisms are often used to provide such mediation.
- *Self organization*. A biological system does not predetermine how it will respond to a challenge, but remembers how it responds. It determines the most effective response, and then keeps the elements that contributed to that response while shedding other elements. In an adaptive approach, all threat response controls must be capable of adjusting their behavior in a similar fashion so that they can employ the most effective countermeasures.
- *Integration with other systems*. Biological organisms are made up of many systems that can be used independently or in concert within a larger ecosystem. In a security architecture, a “defense-in-depth” strategy helps IT systems to achieve a similar level of flexibility.

Adhering to the above set of architectural principles can help in developing an effective security design. The parallels between adaptive immune systems can help to shape the goals and characteristics of an adaptive IT security model.

Design Approach to an Adaptive Security Model

Considering the overall success of most organic ecosystems, mimicking an autonomic system that incorporates an immune response capability is a reasonable design approach in creating a secure infrastructure. The primary difficulty is designing a system that does not replicate the complexities of biological systems but supports system diversity and resists the use of monocultures.

A simple example of applying some adaptive principles is a “defense-in-depth” security architecture that implements several different strategies. Diversity can be achieved using mechanisms such as clustering, redundant hardware, or multiple types of firewall appliances from multiple vendors. In this way, if one component succumbs to a particular threat, it is likely that the others will not, so the survivability of the system is maintained. Likewise, the property of resilience can be supported through virtualization techniques. Using virtualization technologies, infrastructure systems can compartmentalize different system services in secure execution containers. Such containers can be used to isolate service instances. If a threat affects a service in one container, it does not impact the execution of services running in other containers, so services within the IT infrastructure can continue. At the same time, response mechanisms can isolate the affected container and localize the attack’s impact.

The primary differentiator in an adaptive security architecture from current state-of-the-art practices is that adaptive security measures are designed to protect not only against known threats, but to anticipate unknown pathogens or threats in a manner similar to the human immune response system. Essentially, emulating the biodiversity found in nature is the underlying means of developing effective attack responses in the IT world.

How can one go about implementing an adaptive security architecture? What follows is an outline of one possible approach. Note that this approach must be integrated into a larger fabric of the overall security architecture. It must occur within the context of other security features such as application, system, and network design, quality assurance, and configuration validation to ensure that all components and design entities comply with overall security policy.

The following is a list of steps to begin designing an adaptive security model:

- Define threats and threat characteristics that are desirable to avoid or destroy. A threat characteristic may be an attribute of a known threat but may not include the entire threat structure. It may also be particular behavior exhibited by some entity or process. (This concept is linked to the notion of self and non-self discussed above.)

- Identify acceptable behavior, trusted components, and actions that must not be mistaken for a threat. (This step is critical to prevent a Denial-of-Service (DOS) attack, and also relates to the notion of self and non-self discussed above.)
- Define triggers to monitor for threats and, as necessary, to invoke an auto-immune system response. These "immune response mediators" are the threat detection sensors that alert the larger IT infrastructure of potential threats and prime the threat response mechanisms.
- Implement redundancy for critical functions. Note that there should not be any critical "trusted" elements that if compromised could cause the entire system to fail. This instantiates the notions of survivability, diversity, and redundancy.
- Define threat response mechanisms that are focused and that do not result in killing the host.
- Define a recovery process whereby systems are capable of adaptively reconfiguring and restarting themselves. Part of this process also includes a learning and knowledge distribution mechanism so the infrastructure learns how to avoid similar threats in the future.
- Define feedback capabilities that allow the threat response mechanisms to validate threats so that they only respond to legitimate and realistic threats. These feedback mechanisms help to ensure that the triggers and threat response mechanisms understand the security context in which they operate. This enables the desired adaptive behavior.

Not every system must have every threat or threat characteristics defined. The goal is to create a diverse population of systems that each have different threat response capabilities and are pooled together in an ecosystem that can facilitate the overall ecosystem's survival. By populating the larger ecosystem with the underlying building blocks of threats and threat responses, individual systems can adapt to threats, respond to these, and then, when response is successful, share that knowledge with other trusted systems that have not been exposed to the original threat.

It is assumed that sacrificial elements are designed into the overall IT ecosystem. Therefore a threshold of acceptable losses must be established and monitored.

Adaptive Security Processing

Adaptive security relies upon the basic notion of "self" and "non-self" in that a system must be capable of understanding and recognizing what is normal behavior and what is not, and determining if abnormal behavior is a potential threat. Thus the fundamental approach of adaptive security involves threat detection, analysis, and response, following these three steps:

- Telemetry — the gathering and monitoring of information
- Correlation — the evaluation of telemetry data

- Response — one or more actions taken to react to a perceived threat

This approach is not unlike that proposed by Taylor et al, who suggest a process comprised of four elements — attack resistance, attack recognition, system recovery, and system adaptation.[8]

Telemetry

Telemetry is the gathering and monitoring of information about a system, networks, data flows, file management, command issuance, manual and automated operational activities, and other activities that can affect the IT infrastructure. Telemetry must be gathered in real time to anticipate threats effectively (i.e., reviewing historical logs does not have the same benefit as evaluating real-time data). This is not to say that historical data is not useful since it is necessary to build a body of information that can be used for trend analysis and predictive analytics.

Correlation

Correlation is the evaluation of real-time telemetry data in conjunction with historical information. Correlating active and real time data with historical information takes into account the possibility of long latency periods between related attacks. If only current activity is evaluated, then it is possible to miss a likely correlation to some previous activity. Correlation must be performed in conjunction with a well-defined security policy and a rules engine. These help to flag significant activity and to reduce false positives (and false negatives).

The ability to correlate information into actionable responses is the crux of adaptive security. Although current state-of-the-art technologies do not allow security mechanisms to actually predict or anticipate future threats, this continues to be a goal, and there are some areas of research that are promising in this regard. Predictive analytics (such as that used in fraud detection) is one such area. Predictive analytics can sift through large quantities of data looking for repetitive patterns of activity as well as anomalous behavior. A similar approach may be useful in analyzing system information to see if there are trends pointing to a potential threat. Likewise, an adaptive system security model seeks to detect patterns as well as unusual behavior since these are tell-tale signs of potential threats.

Response

Response mechanisms act to prevent the introduction or spread of a threat within the IT infrastructure — they perform the concrete actions that cause the ecosystem to adapt. As with the correlation process, response mechanisms take specific actions according to a well-defined security policy and set of rules. Responses often include the modification of system configurations, system characteristics, and behavior, as well as

8. Taylor, C. and Alves-Foss, J. *Attack Recognition for System Survivability: A Low-level Approach,* Proceedings of the 36th Annual Hawaii International Conference on System Sciences, 2003.

halting systems if necessary. The goal of response mechanisms is to limit exposure and impacts that might adversely affect service levels. It is assumed that some infrastructure components will succumb to attack or may be sacrificed. What is critical is that the response mechanisms contain the threat before it can fully manifest into a catastrophic event.

Integral to response is an orchestration mechanism used to alert other IT system components that a successful threat response strategy has been used. To alert other components, a communications capability broadcasts the set of threat characteristics (so it can be determined if other IT components are susceptible), along with any specific code, configuration, patch, or other information that may be needed to prevent the spread of the attack.

These three basic steps — telemetry, correlation, and response — must be automated to enable rapid deployment of patches and security updates. Automating these steps and removing the human element (where possible) has the benefit of reducing error, which can sometimes be a roadblock in remediation efforts. By responding to threats before they fully manifest themselves, an adaptive security model aims to increase operational efficiency, limit the attack surface within the IT infrastructure, reduce the attack velocity of active threats, and decrease remediation time to recover from hostile actions.

Summary

Biological organisms and ecosystems provide both interesting and useful parallels in designing an adaptive security architecture for IT infrastructures. Mimicking an autonomic system that includes immune response mechanisms can be an effective design approach in building an infrastructure that follows the principles of adaptive security.

The approach described in this article provides a starting point in creating an adaptive infrastructure that can adjust and respond to potential security threats. Establishing security policies and defining baselines for “self” and “non-self” analyses are some of the first steps. Applying telemetry, correlation, and response mechanisms to process threats are key to the effort to protect infrastructure systems and remediate effectively against attack.

Ultimately adaptive security measures seek to protect against not only known threats, but also to anticipate unknown pathogens. Although many elements of adaptive security constructs exist today, not all function together in a coherent manner to fully realize this goal — namely being able to anticipate threats. The primary barrier to a completely adaptive solution is the inability to evaluate data in a timely manner against security policy rules that are actually anticipatory in nature. Adaptive security

measures continue to evolve over time, and there will no doubt be improvement in the ability of the IT ecosystem to adapt and respond to a new and changing generation of threats.

About the Author

Joel Weise has worked in the field of data security for over 29 years. As a Principal Engineer and Chief Technologist for Sun's Global Systems Engineering Security Office, he designs system and application security solutions for a range of different enterprises.

Joel is the Sun Microsystems ANSI representative, an original member of the Information Systems Security Association, the chairman of the ISSA Journal Editorial Advisory Board, a member of the Sun Patent Review committee, a member of the Sun Blueprints Merit Review Board, and author of various Sun Blueprints as well as other external publications. Joel also is a member of and a Subject Matter Expert for the American Bar Association Science and Technology working committee.

References

Brunette, Glenn. "Toward Systemically Secure IT Architectures," *Sun BluePrints Online*, February 2006. To access this article online, go to `http://www.sun.com/blueprints/0206/819-5605.pdf`

Chao, Dennis L., and Forrest, Stephanie. *Information Immune Systems*, Proceedings of the First International Conference on Artificial Immune Systems. Springer Netherlands, 2002.

Chess, D. M.; Palmer, C. C.; White, S. R. *Security in an Autonomic Computing Environment*, IBM System Journal 2003, Volume 42, Number 1, 2003. (http://www.research.ibm.com/journal/sj/421/chess.pdf)

de Castro, Leandro Nunes, and Timmis, Jonathan. *Artificial Immune Systems: A New Computational Intelligence Approach*, London: Springer-Verlag, 2002.

Elkhodary, A. Whittle, J. Inf. & Software Eng., *A Survey of Approaches to Adaptive Application Security*, Software Engineering for Adaptive and Self-Managing Systems, Minneapolis, MN: ICSE Workshops SEAMS '07, 2007.

Geer, Dan. *Monoculture on the Back of the Envelope*, ;login: The USENIX Magazine, Volume 30, Number 6, December 2005. (http://www.usenix.org/publications/login/2005-12/openpdfs/geer.pdf)

Geer, Daniel; Bace, Rebecca; Gutmann, Peter; Perry Metzger, Perry; Pfleeger, Charles P. ; Quarterman, John S., and Schneier, Bruce. *CyberInsecurity: The Cost of Monopoly, How the Dominance of Microsoft's Products Poses a Risk to Security*. Computer & Communications Industry Association, September 24, 2003. (http://www.ccianet.org/papers/cyberinsecurity.pdf)

Liang, Gang. *An Immunity-Based Dynamic Multilayer Intrusion Detection System,* Lecture Notes In Computer Science. Heidelberg: Springer Berlin 2006.

Mazhar, Nauman and Farooq, Muddassar. *BeeAIS: Artificial Immune Systems Security for Nature Inspired, MANET Routing Protocol, BeeAdHoc,* Lecture Notes In Computer Science. Heidelberg: Springer Berlin, 2007.

Sante Fe Institute, http://www.santafe.edu/research/topics-information-processing-computation.php.

Saxena, Anshuman; Lacoste, Marc; Jarboui, Tahar; Lücking, Ulf; and Steinke, Bernd. *A Software Framework for Autonomic Security in Pervasive Environments*, Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Security and Defense Applications. Heidelberg: Springer Berlin, 2007.

Stevens, M. and Williams, P. D. *Use of Trust Vectors for CyberCraft and the Limits of Usable Data History for Trust Vectors*, Computational Intelligence in Security and Defense Applications, 2007.

Taylor, C. and Alves-Foss, J. *Attack Recognition for System Survivability: A Low-level Approach,* Proceedings of the 36th Annual Hawaii International Conference on System Sciences, 2003.

Ulieru, Mihaela, and Worthington, Paul. *Autonomic Risk Management for Critical Infrastructure Protection*, Integrated Computer Aided Engineering. Amsterdam: IOS Press, 2006.

Waldrop, M. Mitchell. C*omplexity: The Emerging Science at the Edge of Order and Chaos,* New York: Touchstone, 1992.

Weise, Joel. *Security Architecture and Adaptive Security*, Information Systems Security Association (ISSA) Journal, July 2008.

Ordering Sun Documents

The SunDocs℠ program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is

`http://docs.sun.com/`

To reference Sun BluePrints Online articles, visit the Sun BluePrints Online Web site at:

`http://www.sun.com/blueprints/online.html`

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** sun.com

Printed in USA 11/08