SEADE

Countering the Futility of Network Security

Mr. Frank Konieczny
Lt Col Eric Trias, PhD, USAF
Col Nevin J. Taylor, USAFR

> *We cannot solve our problems with the same thinking we used when we created them.*
>
> —Albert Einstein

Today's media is flooded with stories of cyber attacks prompting a loss of public confidence, resignations by senior officials, and a significant near- and long-term impact on our nation. Most of these breaches stem from known vulnerabilities in existing network security architecture, presenting a distinct danger to our vital national interests. These vulnerabilities, which vary in sophistication, could be as simple as using weak passwords (e.g., default value, simple number strings, or the word *password* itself). Slightly more sophisticated attacks leverage phishing attempts through e-mail or social engineering, designed to elicit unsafe action or information that would allow adversaries unauthorized access.

The notion of "defense in depth" has been touted by leading security organizations (which rely on the National Institute of Standards) as the basis upon which a security framework can be developed to safeguard our networks. The depth includes both physical security protections (walls, gates, locks, guards, and computer cages) and logical security measures (network firewall and intrusion detection). However, no matter how many layers of network perimeter protection are employed, adversaries continue to overcome defenses through using a variety of countermoves or by exploiting poor cybersecurity practices.

Furthermore, successful cyber attacks highlight the fact that disciplined cyber hygiene is necessary but not sufficient to prevent all potential attacks. Systems are simply too complex to defer application and data security to the supporting network's defense appliances and infrastructure. Therefore, we propose that, from their inception, applications must be designed to protect themselves as stand-alone entities with security built-in and with minimal security dependence on network security appliances (e.g., firewalls).

As Secretary of Defense Ashton Carter proclaimed during a speech at Stanford University, to keep systems secure, we must build "a single security architecture that's more easily defendable and able to adapt and evolve to mitigate current and future cyber threats."[1] We propose that this next evolution be a "designer" security package at the application level: the security-encapsulated application and data enclave (SEADE) architecture composed of a virtual application data center (VADC) and enterprise-level security (ELS). SEADE will redirect the responsibility for an enterprise-level network security perimeter to each application. It will act as a separately secured virtual container that offers users enhanced data access and produces an application package that is exceedingly difficult to penetrate and easy to port; furthermore, SEADE requires little maintenance.

Insufficient Network Perimeter Defense

In the past, strategic endeavors in this area have focused on safeguarding the information that resides within our networks by building higher and thicker walls around our *crown jewels*, posting gate guards that interrogate everyone entering or leaving, and establishing multiple checkpoints. These efforts attempt to mitigate accessibility, the very capability our modern networks have been designed to provide. Clearly, this has been a losing proposition because the cost to safeguard these networks far exceeds that associated with attacking and penetrating them. Critically, it also impedes unobstructed and timely access by our forces to the information they so critically need.

The current network enclave defense model parallels these classic perimeter defenses by restricting accessibility to apparently valid users or transactions. However, it does little to define the purpose behind the effort. Thus, without a clear understanding of what is to be defended, we are left with the daunting task of defending everything in our "house/fort" without having any opportunity to prioritize a specific effort, such as those that will likely have the greatest impact on our ability to accomplish the mission.

It is imperative to note that our traditional approach to protection using only network boundaries is rendered useless when an adversary is already inside the network. Based

on recent events and given current levels of network complexity, it is unlikely that adversaries will appear via concentrated denial-of-service attacks as was once the case. Rather, we would be well advised to conclude that such enemies already exist within our networks. More realistically, they are striving to hide their presence in order to harvest information that represents the lifeblood of our companies, plans, and/or intellectual property. Consequently, the three core considerations that must be governed by security measures are (1) accessibility, (2) confidentiality (including the determination that data is correct and has not been altered), and (3) integrity (which relates to the essence of our trust in and reliance on information used in the decision-making process). The complexity of recent cyber attacks has indeed increased. Although they were once focused on pilfering or manipulating data, such attacks now seek not only to steal critical data but also to undermine its use within operational command and control centers. Indeed, threats that have remained dormant until triggered by a specific event (e.g., zero-day attacks) can have devastating consequences at the most inopportune times during military operations. Therefore, we must elevate our awareness of such threats and manage the associated risk by determining what must be defended, how such defenses will be carried out, what objective will be fulfilled, and why it is important. Ultimately, networks that continue to offer unfettered accessibility (albeit a worthwhile quality) will fail to secure the intellectual property that populates today's information environment. Clearly, then, we must take a step back and ask ourselves what we should defend. Should we protect the roads and highways (i.e., the network) leveraged by users and adversaries alike? Or should we protect the data and intellectual property inside?

Current State of Enterprise Defense

Today's perimeter defenses are instrumented for network-traffic-based analysis that assumes nothing bad will happen to applications/data if those defenses prevent malware transactions at the entrance. The solution—based on consistent, quick recognition of these rogue transactions—works well if one knows and understands *all* of the acceptable transactions so that the complement can be characterized as unacceptable (i.e., blacklisting undesirable network traffic).

Another defensive approach entails isolating the application from external access channels, but business requirements mandate access to areas inside the perimeter for collaboration (data sharing), interaction (web services), mobile/remote access (virtual private network), and business-to-business links. Hence, it is extremely difficult to determine which traffic to block because of multiple exceptions that must be accommodated for the business to function. Blacklisting has become slow and unwieldy to maintain and does not scale well, especially with the increasing adoption of IPv6.[2] Whitelisting at the perimeter level has become unmanageable due to the thousands of entries to maintain. The fact that the *walls* have to allow a superset of all of these exceptions creates a porous perimeter. Moreover, adding new or removing existing exceptions may cause unintended effects on other applications, typically discovered only after implementation. Further complicating the situation is the continuing maintenance requirement—for example, obsolete exceptions persist in configurations because of a failure to notify administrators to make the updates.

Compounding the situation is the scaling of network defenses to billions of transactions. The usual response to keeping pace with performance demands has been to increase the sophistication and scale of network defense appliances. Unfortunately, these "improvements" exert more overhead and cause greater latency (despite appearing faster or more robust) and do not always produce more effective systems.

There has to be a better way. To better defend our information, not only do we need to recognize that fact and account for the adversaries among us, but also we must continue to operate within this contested environment. Since our cyber adversaries have made their presence known, we must find novel ways to defend the vital information (today's crown jewels) that enables us to maintain our competitive edge, all the while accepting the idea that we will be operating in a contested environment. As we focus on protecting our property and establishing tighter security perimeters, we will also develop the ability to scale our approaches quickly and overcome continually increasing threats.

In the past, isolated enclave architecture was the initial design of the network—each group had its own enclave with no outside connectivity. The desire to share information led to connecting these enclaves, which generated some concern, but a trust agreement existed between them. As enclaves became increasingly interconnected, the level of trust degraded further, especially when control was lost and anonymity became pervasive within the World Wide Web. Regaining this trust involved employing enterprise perimeter defenses to control access to information and restricting data availability to maintain some degree of confidentiality.

Although this problem has long been recognized and many alternatives have been proposed, only a modicum of success has been achieved in safeguarding intellectual property. The obvious alternative is to construct multiple layers of network perimeter defenses that provide adequate confidentiality of strategic data. However, this approach requires that different settings, configurations, or tool sets be established at each point in the layered defense. Ultimately, such an action increases the maintenance burden and produces delays in transaction flow, the combination of which impedes timely dissemination of vital information.

Incident Identification/Reaction

Considering that network perimeter defenses are generating logs/alerts to billions of transactions in a large organization, how does one analyze these into a coherent picture? Even more desirable, how can one detect in "real time" that malware is present and that an incident can be prevented? This problem is difficult because little information exists to determine which application a specific transaction belongs to unless additional network defenses are placed in multiple locations in the enterprise, usually near data centers, to record and analyze all network traffic. Of course, this scenario generates even more data for analysis, and one winds up looking for the proverbial needle in a stack of needles. An obvious solution involves using special-purpose "big data" analysis tools such as predictive analysis techniques, cross-correlation analysis, and so forth, with plenty of storage for historical transactions. Obviously, this analysis overhead further adds costs and resources to defense efforts. There has to be a better way.

A Better Way

Since attacks continue despite our best network perimeter defenses, what if we begin with the assumption that adversaries are already on our networks? Consequently, we must adjust our threat model and think differently to protect our data and intellectual properties. What if we decrease the attack surface down to the application or data level with the same security capabilities currently used for perimeter defense but specialized for the particular application or data? This vision lies at the heart of the SEADE concept, which defuses the overall attack surface from gateways guarding the enterprise network perimeter to thousands of individual, specialized security enclaves. The multitude of enclaves, consisting of multiple products and specialized configurations, will force the attacker to increase his effort to penetrate a single application. Since each security enclave is specialized to a specific application, the attacker must customize attacks per application rather than focus on penetrating the perimeter to expose the entire network. Thus, it will no longer be possible for adversaries to exist unchallenged inside our networks.

SEADE—Virtual Application Data Center

Virtualization technology, available in the *cloud* or virtual data centers (VDC), has made possible the virtual application data center concept. A VDC is a software-defined data center that supports "infrastructure as a service" for applications. It is a commodity readily available in many commercial and government cloud data centers. We utilize a VDC to define a VADC. Essentially, one VADC is dedicated to only one application, which is supported by a platform as a service (PaaS). It consists of virtualized network monitoring and defense capabilities like firewalls and deep-packet inspection along with its associated web access point, database firewall, and traditional PaaS components of web servers, application servers, and database servers. SEADE-VADC extends this concept for each application.

A significant security benefit of this architecture is that network traffic can remain encrypted until it enters the VADC. Only after packets enter the VADC are they decrypted and inspected. Within each VADC, the application developer has tailored the network inspection defenses, which were "baked in" from the design phase, to the specific ports/protocols, transaction size/format, parameter range, and so forth, for that single application.[3] For instance, some applications may be tuned to support deep-packet inspection with abnormalities reported to the appropriate computer network defense service provider (CNDSP). Individual application risk management will drive the tailoring requirements. The VADC will improve the levels of *accessibility* and *confidentiality* by recognizing specific threats immediately and preventing an incident from occurring.

SEADE—Enterprise-Level Security

ELS is a dynamic attribute-based access-control system developed to reduce overall security risks by automating the access process, based on authoritative, related attribute information.[4] Today, each application has a uniquely configured access-control scheme maintained by system administrators, primarily based on users and groups,

which can be quite labor intensive. In the Air Force, the process is further burdened by a form-based, administrative-access approval process. As a new paradigm, ELS automates the authorization maintenance process; validates preconditions for access, such as training, security clearance, rank, and so forth; and allows a person access when an application-owner-defined set of conditions is met.

Accessibility to data is controlled by *claims*, based on a person's (or an entity's) attributes, dynamically generated and propagated when attributes change.[5] Claims can be additions, deprecations, or modifications to existing access rights. They are transmitted via encrypted channels, based on user-access requests in a security assertion markup language (SAML) token. A standard handler evaluates and validates the token (content, timing, and authentication) and passes the claim for access to the application. Logging occurs for every access request, and erroneous access information is sent to the appropriate CNDSP. A standard handler ensures that SAML validation and access logging are performed correctly, further freeing the application developer from producing similar capability.

ELS will improve the levels of *integrity* and *confidentiality* by preventing unauthorized data access. As shown in the figure below, SEADE combines both concepts (VADC and ELS) and is delivered as two VDCs—one for the application (VADC) and the other for the ELS claims engine (which includes the secure token service, enterprise attribute store, and generated SAML claims).
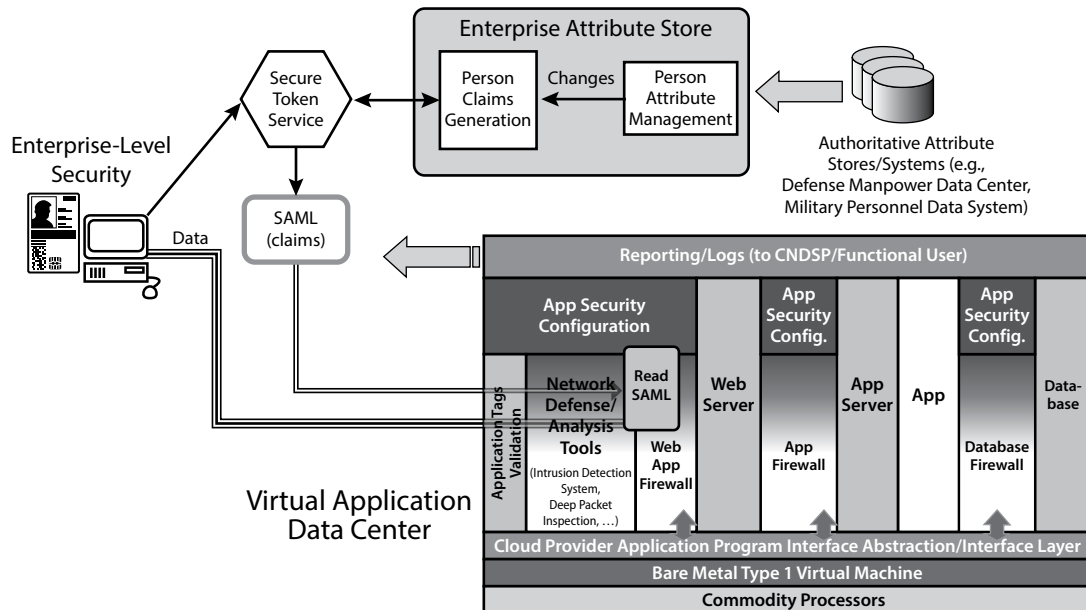

Figure. SEADE diagram

Benefits of SEADE

Employing SEADE throughout a large enterprise-level operation generates the following benefits:

- *Enables application portability*. SEADE promotes such portability by enabling applications to be hosted in any virtualized environment. Thus, owners have the freedom to maneuver applications where they are needed to meet operational and resiliency requirements.

- *Expedites application deployment*. Multiple SEADEs employed throughout the enterprise will significantly decrease the manpower associated with developing and fielding an application. Since network and application defenses are included in the standard PaaS environment, the application itself remains just the logic of the program as it inherits all of the security controls of the PaaS. This architecture has demonstrably decreased the time to production from months to weeks. Since a standard ELS handler may be used for the SAML token, the application developer need only code to the ELS handler's application program interface, further decreasing deployment time.

- *Facilitates accreditation*. Since applications are encapsulated with their own security functions, porting them into new hosting environments will be minimal, including justification of security measures to meet the accreditation process.

- *Eliminates individual access requests*. Dependence on form-based administrative processes will be eliminated, and system administrators' access-management burden will be significantly reduced. There will no longer be user and group permissions to maintain per application, drastically reducing the man-hours required to perform this basic system-administration function.

- *Provides immediate user access*. Users will have immediate access to applications and data, based on their attributes (e.g., position, training, duty location, and so forth). As soon as the authoritative data source is updated with their personnel information—say, to a new assignment—then users will be granted access accordingly.

- *Includes "baked-in" security*. Application development will change fundamentally by baking in security from the start. Developers will integrate network defense configurations (e.g., whitelisting) into their VADC. Further, they will have more options and stronger security-related capabilities by having various network appliances at their disposal. Developers must now think holistically and produce applications to respond to and interact only with defined, valid, and recognized inputs.

- *Focuses incident reports*. Instead of having cyber war fighters *look* at streams of network transactions, trying to determine an abnormality, incident reporting is narrowed to the actual application with detailed information, based on the application's tailored security profile. The CNDSP will be alerted only when thresholds are triggered.

- *Reduces the number of network administrators*. Network security operators will no longer have to make network appliance configuration changes (e.g., firewalls, proxies, and intrusion detection systems) to "allow only" legitimate traffic and block known, bad traffic. Additionally, less time will be spent on configuration-management meetings to approve mundane changes to network appliances.

- *Provides operational resiliency*. Since the VADC is composed solely of virtual components, if an abnormality is detected, the application can be dynamically reloaded from a previously known good image, or snapshot, to continue processing. As an added resiliency measure, SEADE instances can be spawned at multiple locations and numerous environments to attain heightened redundancy and increased mission assurance.

- *Enables continuity of operations (COOP) and agility*. By leveraging virtualization, one can provision applications in multiple environments, as well as COOP to another data center, provided that data has been streamed to the COOP site. This capability of provisioning anywhere further decreases the time for provisioning and provides significant mission agility.

- *Reduces insider threat*. This new paradigm enables creative approaches to data protection. Vulnerability to an insider threat will be reduced since ELS will block unauthorized access and track all access to applications or data. This information can be used to detect or predict abnormal activities. With appropriate data-access tagging, exfiltrated data will be unreadable outside an environment without SEADE.

- *Improves confidentiality, integrity, and availability*. The SEADE combination of ELS and VADC capabilities significantly increases the *confidentiality* and *integrity* of the data by preventing unwarranted access and *availability* of the application (and data) by dynamic analysis and elimination of threats to the application itself.

- *Maintains CNDSP*. The current CNDSP framework does not have to change. Alerts within each SEADE can be sent to the appropriate CNDSP unit, which will continue to triage alerts accordingly.

Trade-Offs

The primary trade-off with employing SEADE is that instead of relying on and deferring to network perimeter security, application developers now will be responsible for considering application security and ELS controls during design, test, and development. The developers must become intimately familiar with their application to address issues for both expected and unknown stimuli. This will undoubtedly increase the initial cost of system development, but it will ultimately save innumerable man-hours and will improve data protection. Developers will be responsible for ensuring that security is incorporated from the onset rather than waiting for operators to address the need retroactively.

Another trade-off is the building of a supporting environment for SEADE services. Application and functional owners must define and govern attributes re-

quired to provide the granularity necessary for applications to have the correct level of access-control fidelity. These attributes must come from known, authoritative data sources that have to be identified and integrated into enterprise attribute store for ELS's use.

Air Force Consolidated Enterprise Information Technology Baselines

Today, technology moves so quickly that one will never reach a 100 percent best solution in a reasonable amount of time. Agile solution delivery is the best approach to a problem via focused sprints and spiral development so one can adjust as the available technology changes. This affords the ability to capitalize on and garner strategic advantage from nimble actions and innovative solutions. Unfortunately, this paradigm shift unsettles many people who expect predefined requirements with predestined end points. However, this traditional approach only wastes resources as the environment and requirement change in their midst. As the *cheese* constantly moves in technology and cyberspace, we must be adaptable and decide to venture out to embrace the changes—lest we risk starvation.[6] We must harness and guide this spirit of innovation and provide a framework for inserting new technology—methodically and expediently—into our environment.

Accordingly, it is in this vein that the Air Force chief technology officer established and manages the Consolidated Enterprise Information Technology Baselines (CEIT-B) framework to purposely shape, adopt, and deliver a standard information technology environment. This disciplined effort conforms to the agile paradigm as the future target baseline is developed.[7] SEADE is a substantial component of CEIT-B that addresses security, portability, and efficiency requirements. Additionally, the Air Force, through CEIT-B, is addressing and informing the joint information environment (JIE) requirements for Department of Defense–level enterprise requirements.

Conclusion

The Air Force, as a service, emerged from technology. We must continue to harness the same innovative spirit for cyberspace that has enabled us to dominate air and space. Innovation is the fuel for future success, and we must keep striving to embrace new ways of solving our difficult problems. SEADE, comprised of a VADC and ELS, is a fundamentally different paradigm that will change the way systems are developed, deployed, and defended. By providing a separate security enclave for applications in a VADC, enabled by ELS dynamic access control, we can protect our most important treasure—the data within—as we continue to operate in a contested environment. The SEADE architecture will increase the speed of both user access and application delivery to the mission, decrease day-to-day management of the network and applications, and counter the futility of network perimeter security. ✪

Notes

1. Cheryl Pellerin, "Carter Unveils New DoD Cyber Strategy in Silicon Valley," US Department of Defense, 23 April 2015, http://preview.defenselink.mil/news/newsarticle.aspx?id=128659.

2. IPv6 (Internet Protocol version 6) is the latest Internet standard protocol that uses 128 bits versus the current IPv4's 32 bits. The new version has capacity for every person on Earth to have billions of Internet addresses personally allocated. Therefore, blocking by individual address or range of addresses will no longer be effective or efficient.

3. "Baked in" refers to integrating desired security features at the initial stage of design and development as opposed to adding them on (i.e., "bolted on") after the product has been released.

4. Vincent Hu, Adam Schnitzer, and Ken Sandlin, "Attribute Based Access Control Definition and Considerations," National Institute of Standards and Technology Special Publication 800-162, n.d., http://csrc.nist.gov/projects/abac/july2013_workshop/july2013_abac_workshop_abac-sp.pdf.

5. Coimbatore S. Chandersekaran and William R. Simpson, "A Uniform Claims-Based Access Control for the Enterprise," *International Journal of Scientific Computing* 6, no. 2 (December 2012): 1–23.

6. Spencer Johnson, *Who Moved My Cheese? An Amazing Way to Deal with Change in Your Work* (New York: G. P. Putnam's Sons, 1998).

7. SAF/CIO A6 CTO, *CIET-B, Target Baseline 2.0*, 2015, https://intelshare.intelink.gov/sites/afceit/TB/default.aspx.

Mr. Frank Konieczny

Mr. Konieczny (BS, MS, University of Illinois–Chicago; MAS, University of Alabama–Huntsville), a senior-level executive, is the Air Force chief technology officer, Office of Information Dominance, and chief information officer, Office of the Secretary of the Air Force, Pentagon, Washington DC. Prior to assuming his current responsibilities, he acquired extensive experience in industry, where he worked as a systems analyst, chief programmer, project manager, and business unit manager, including positions as chief scientist and chief technology officer.

Lt Col Eric D. Trias, PhD, USAF

Lieutenant Colonel Trias (BS, University of California–Davis; MS, Air Force Institute of Technology [AFIT]; PhD, University of New Mexico) is the acting chief, Air Force Enterprise Architecture Division, Cyberspace Strategy and Policy Directorate, Secretary of the Air Force, Office of Information Dominance and Chief Information Officer, Pentagon, Washington, DC. Charged with governing, developing, and maintaining the Air Force enterprise architecture, he also serves as the Air Force deputy chief technology officer, helping evaluate and define future information technology standards and implementation constraints for the Air Force information technology enterprise infrastructure. Lieutenant Colonel Trias has served as an assistant professor at AFIT, commander of a large detachment, and deployed squadron deputy commander. He has held various leadership positions in a base communication squadron, exercise control squadron, and combat communications squadron.

Col Nevin J. Taylor, USAFR

Colonel Taylor (BS, University of the State of New York; MS, Capella University) is the individual mobilization augmentee (IMA) to the director of cyberspace strategy and policy, deputy chief technology officer for special programs, and chair of the Cyber Task Force's Strategic Advisory Board in the Office of Information Dominance and Chief Information Officer, Office of the Secretary of the Air Force, Pentagon, Washington, DC. He is a 10-year space and 20-year cyber professional with over a decade of command experience and a plethora of unique, diverse operational expertise, including combat, fixed and space communications, mission support, acquisitions, policy, strategy, planning, cyber, and space. Colonel Taylor's joint assignments include director of Component Reserves, Joint Functional Component Command for Space, US Strategic Command; senior military assistant to the deputy undersecretary of defense for policy integration; and chief of staff as well as IMA to the undersecretary of Department of Defense policy in the Office of the Secretary of Defense.

Let us know what you think! Leave a comment!

Distribution A: Approved for public release; distribution unlimited.

http://www.airpower.au.af.mil