FEDERAL CYBERSECURITY BEST PRACTICES STUDY: INFORMATION SECURITY CONTINUOUS MONITORING

October 2011

Bruce Levinson
Center for Regulatory Effectiveness
1601 Connecticut Avenue, NW
Washington, DC 20009
www.TheCRE.com/fisma

FEDERAL CYBERSECURITY BEST PRACTICES STUDY: INFORMATION SECURITY CONTINUOUS MONITORING

1.0 Introduction

This study documents the successful work by NASA's Earth Observing System (EOS) Security Team in thwarting the cybersecurity challenges posed by an Advanced Persistent Threat (APT). Through a combination of initiative and creativity by the NASA EOS Security Team and their use of sophisticated software for continuous monitoring which could adapt to changing needs on-the-fly, the team prevented the agency's information system security from being breached following the highly publicized hack of RSA which compromised a key component of the agency's protocol for authenticating users.

In recognition of NASA's cyberdefense success, the NASA EOS Security Team's use of Splunk for Information Security Continuous Monitoring (ISCM) is recognized by the Center for Regulatory Effectiveness as a Federal Cybersecurity Best Practice.

1.1 FISMA and Continuous Monitoring

The Federal Information Security Management Act (FISMA, Title III of the E-Government Act (Pub. Law 107-347)) provides a framework for "for ensuring the effectiveness of information security controls over information resources that support Federal operations and assets...." One of the law's requirements is for monitoring. Specifically, FISMA amended the Paperwork Reduction Act (PRA) to include as one of OMB's Information Policy responsibilities the "monitoring, testing, and evaluation of information security controls...." [44 U.S.C. § 3505 (c)(3)(C)(iii)]

The National Institute of Standards and Technology (NIST) created the Risk Management Framework (RMF) as a risk-based paradigm to help guide their FISMA implementation work. The final step in the RMF, as discussed in Chapter 3 and Appendix G of NIST Special Publication 800-37 Rev. 1, is Monitor Security Controls/Continuous Monitoring. Additional discussion and guidance specific to continuous monitoring may be found in NIST's Special Publication (SP) 800-53 Rev. 3, SP-800-53A, SP 800-137, and draft NIST/DHS Interagency Report 7756, *CAESARS Framework Extension: An Enterprise Continuous Monitoring Technical Reference Architecture*.

ISCM should be viewed as the capstone of an effective security control program. NIST's SP 800-37 Rev. 1, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach* discusses the central role of continuous monitoring in IT security in stating that,

> *A well-designed and well-managed continuous monitoring program can effectively transform an otherwise static security control assessment and risk determination process into a dynamic process that provides essential, near real-time security status-related information to organizational officials in order to take appropriate risk mitigation actions and make cost-effective, risk-based decisions regarding the operation of the information system.*

The dynamic role of ISCM and its close integration into an organization's IT security defenses is explained in NIST's continuous monitoring guidance document, SP 800-137 which states that

> *the process of implementing ISCM is recursive. ISCM informs and is informed by distinct organizational security processes and associated requirements for input and output of security-related information.*[1]

1.2 About the Center for Regulatory Effectiveness

The Center for Regulatory Effectiveness (CRE) is a non-partisan regulatory watchdog founded by former senior career officials from the Office of Management and Budget. As a watchdog, CRE works to ensure agency compliance with the "good government" laws that regulate the regulators including the Data Quality Act, the Paperwork Reduction Act, the Regulatory Flexibility Act, and Executive Order 13563 on regulatory review.

CRE has worked extensively on federal cybersecurity issues. In 2005, as part of its CyberSecurity Policy Project, CRE wrote a widely reprinted article, *What Will You Do When the Cyber-Levee Breaks?* which called for the creation of an interactive cybersecurity forum.[2]

In 2010, CRE established FISMA Focus (www.thecre.com/fisma/), an Interactive Public Docket (IPD)[3] dedicated to enhancing the transparency and effectiveness of federal cybersecurity policy. It is inevitable that federal cybersecurity regulations will be increasingly applied to at least some private sector networks. FISMA Focus seeks to ensure that any such regulation is transparent, meets stringent benefit-cost tests, and complies with other good government protections including those specific to small businesses.

Application of federal cybersecurity requirements to the private sector is not merely speculative. On June 29, 2011, the Department of Defense published a proposed rule to amend the Defense Federal Acquisition Regulation Supplement (DFARS) "to add a new subpart and associated contract clauses to address requirements for safeguarding unclassified DoD information."[4] The proposed rule would apply security controls from NIST SP 800-53 *Recommended Security Controls for Federal Information Systems and Organizations* to defense contractors.

Additional application of FISMA standards and practices to the private sector should be expected.

---

[1] K. Dempsey, N.S. Chawla, et al., "NIST Special Publication 800-137: Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations," Computer Security Division, Information Technology Laboratory, NIST, September 2011, p. 6.

[2] B. Levinson, "What will you do when the cyber-levee breaks?" CIO, September 21, 2005 *available at* http://thecre.com/pdf/What_will_you_do_Cyber-Levee.pdf.

[3] See, Wikipedia, *available at* http://en.wikipedia.org/wiki/Interactive_Public_Docket.

[4] 76 Fed Reg. 38089-95, Wednesday, June 29, 2011, *available at* http://www.gpo.gov/fdsys/pkg/FR-2011-06-29/pdf/2011-16399.pdf.

2.0 NASA's Leadership in Continuous Monitoring

NASA is among the federal agencies that have taken a leadership role in implementing ISCM. The agency moved from taking a reactive approach to addressing security breaches to implementing a proactive, automation-aided, risk-based approach to confronting IT security challenges.

Key to NASA's new approach was their focusing on cost-effective cybersecurity which meant shifting away from viewing monitoring requirements as a paper-based, checklist exercise of limited value. OMB played a crucial role in agencies moving to a more effective security stance by issuing a Memorandum (M-10-15) on FISMA reporting requirements that emphasized automated security tools, flexibility, and the need for outcome-based metrics.[5]

2.1 OMB Memorandum M-10-15 of April 21, 2010

The OMB Memorandum provided instruction to agencies on their FY 2010 FISMA reporting requirements and emphasized that "Agencies need to be able to continuously monitor security-related information from across the enterprise in a manageable and actionable way." The document went on to explain that "agencies need to automate security-related activities, to the extent possible, and acquire tools that correlate and analyze security-related information."

The Memorandum reflected an advance in how agencies are directed to implement their FISMA reporting responsibilities. Of particular note is the Memorandum's reporting requirement for data feeds to come directly from security management tools. Specifically, OMB directed, that "reporting should be a by-product of agencies' continuous monitoring programs and security management tools."

The OMB direct data feed requirement provides the path for agencies to reduce paperwork burdens as they become able to take advantage of paperless reporting through the CyberScope platform. The longer term significance of the memo, however, is the philosophical shift it represents in how senior IT officials should approach FISMA compliance.

Instead of viewing FISMA reporting as a make-work, checklist exercise that is undertaken largely to satisfy OMB and statutory requirements, the White House's new approach encourages CIOs and CISOs to think of the continuous monitoring data as actionable intelligence and to focus on using it to improve security.

2.1.1 About CyberScope

CyberScope is a reporting system developed under the auspices of the Federal CIO who described it as an "interactive data collection tool" allowing "agencies to fulfill their FISMA reporting requirements through a modern digital platform. The broad range of meaningful information collected, the use of secure

---

[5] J. Zients, V. Kundra, H.A. Schmidt, "Memorandum for Heads of Executive Departments and Agencies," M-10-15, April 21, 2010 *available at* http://thecre.com/pdf/OMB.M-10-15.pdf.

two-factor authentication, and the online access to data provides for a more efficient and effective reporting process."[6]

The Center for Strategic and International Studies (CSIS) stated that CyberScope was intended "to replace the existing insecure paper or e-mail based reporting. In addition to improving the security of the reports, CyberScope streamlines the process by providing a standard format for reporting, allowing for greater insight into the data and negating the need to combine reports submitted in various formats. Ultimately CyberScope will result in a 'cybersecurity dashboard,' not unlike the IT Dashboard (it.usaspending.gov) that currently tracks federal spending on IT projects."[7]

CyberScope has been a controversial project and was greeted with significant resistance by much of the federal IT security community. As one trade publication noted, "According to a survey of chief information officers from 34-Cabinet level departments and other agencies by MeriTalk, 15 percent of CIOs had tried CyberScope, with the large majority of those who had not used it doubtful of its purpose and 'suspicious of its effectiveness'...."[8]

Thus, the decision by senior NASA IT security officials, discussed below, to welcome CyberScope and to take advantage of the opportunities presented by M-15-10 stand in contrast the views of much of the federal IT establishment.

2.2 NASA Deputy CIO Memorandum of May 18, 2010

NASA's IT security leadership used the OMB Memorandum as the launching pad for transforming the agency's approach to FISMA. Specially, NASA's Deputy CIO for IT Security distributed a Memorandum that recognized and capitalized on the revamped approach to FISMA.[9]

The NASA memo explained that the meaning of the OMB Memorandum "is clear regarding a shift away from cumbersome and expensive C&A [Certification and Accreditation] paperwork processes, in favor of a value-driven, risk-based approach to system security." NASA emphasized the new IT security flexibility by further explaining that "Per M-10-15, NIST recommendations inherently 'allow agencies

---

[6] Statement of Vivek Kundra, Federal Chief Information Officer, before the Senate Homeland Security and Governmental Affairs Subcommittee on Federal Financial Management, Government Information, Federal Services, and International Security, October 29, 2009 *available at* http://www.cio.gov/Documents/Vivek_Kundra_Federal_Cyber_Defense_Testimony_10-29-2009.pdf.

[7] P. Kimmey, "FISMA, Cyberscope and Federal IT Security," February 26, 2010 *available at* http://csis.org/blog/fisma-cyberscope-and-federal-it-security.

[8] B. Kalish, "CIOs Not Into CyberScope," NextGov, October 4, 2010 *available at* http://techinsider.nextgov.com/2010/10/fed_cios_not_using_cyberscope_despite_upcoming_deadline.php.

[9] J. L. Davis, "Suspension of Certification and Accreditation Activity," May 18, 2010 *available at* http://www.nasa.gov/pdf/501521main_Suspension%20of%20C%26A%20Activities.pdf.

latitude in their application [of security solutions ...]. Consequently, the application of NIST guidelines by agencies can result in different security solutions that are equally acceptable and compliant.'"

NASA's OCIO took quick advantage of OMB memo to streamline FISMA compliance in a number of ways including not requiring "Information System Owners (ISO) to recertify their systems in FY 2010 to satisfy OMB requirements" and "In lieu of C&A activities in FY 2010, AOs [Authorizing Officials] must extend current Authorizations to Operate (ATO) for a period not to exceed one year...."

One of the most important aspects of the NASA memo was its embrace of continuous monitoring. The memo stated that NASA's IT Security Division

> *is creating a more streamlined system security authorization process with a focus on continuous monitoring, automated tools, and significant paperwork reduction. These developing processes will eventually enable near real-time risk management and ongoing security authorizations that reflect the true intent of NIST guidance, and fall in line with the objectives of DHS, DOJ, the Whitehouse, recently proposed amendments to federal security legislation, and new OMB mandated tools.*

The memo put NASA in the forefront of complying with FISMA implementation requirements and also in using continuous monitoring tools to improve IT security. In June 2011, DHS published the initial set of Reporting Metrics for use with CyberScope.[10] In commenting on the development, three Editor's Notes on the SANS Institute's news blog encapsulate the crucial change in cybersecurity perspective embodied in the document, a perspective that NASA had championed and leveraged over a year earlier and that some agencies were still wrestling with.

> *[Editor's Note (Hoelzer): This is an extremely important step. Federal CIOs and others have known for a long time that the 'Report Card' method just doesn't work since it completely fails to address the real risks that a particular agency faces. A Continuous Monitoring focus means that FISMA compliance is starting to align with what much of the FISMA constituency has been saying: Government agencies must have the correct monitoring systems deployed, they must be monitoring the correct things and they must be providing meaningful information to inform the defenders about events and trends. It is heartening to see FISMA compliance coming closer into line with the 20 Critical Security Controls.*

---

[10] U.S. Department of Homeland Security, National Cyber Security Division, "FY 2011 Chief Information Officer Federal Information Security Management Act Reporting Metrics, Version 1.0," June 1, 2011 *available at* http://www.thecre.com/fisma/wp-content/uploads/2011/06/DHS_FISMA-ReportingMetrics.v-1.0.pdf.

> *(Pescatore): To most federal agencies, the reporting requirements are increasing much faster than security budgets are increasing.*
>
> *(Paller): The agencies do not have to continue wasting money on the old reporting - they continue only because it makes the FISMA contractors money and because of the Stockholm syndrome (the CIOs and CISOs have been captives of the paper-compliance fanatics for so long that the victims cannot believe they are free to use the money to do the right thing (continuous, automated, daily monitoring).]*[11]

2.3 NASA's Presentation to the Information Security and Privacy Advisory Board (ISPAB) of August 4, 2010

ISPAB is a Congressionally-mandated federal advisory committee originally created in 1998 as the Computer System Security and Privacy Advisory Board. The Board's duties include advising NIST and OMB "on information security and privacy issues pertaining to Federal Government information systems, including thorough review of proposed standards and guidelines developed by NIST."

ISPAB meetings often include presentations from agencies on information security issues and developments. ISPAB's August 2010 meeting included a presentation from the NASA Deputy CIO who drafted the May 18th memo. In his presentation, Mr. Davis discussed the NASA memo and the shift it represented in how his office approached IT security.[12]

At the start of the presentation, NASA acknowledged that their previous risk management strategy had been "to 'wait' for the incident to occur and then, if detected, respond (highly reactive) and then repeat" which meant that the agency's response "was generally slow and almost always after the incursion has taken place and the data or system is completely compromised." Moreover, "Continuous risk management did not take place so root cause is generally unknown and thus data, information and systems remain at risk of further compromise."

The May 18th NASA memo was described by the security official as "a shift in direction" in the way the agency's IT leadership viewed security. It should be noted that NASA's shift was completely in keeping with SP 800-37's principle of "transform[ing] an otherwise static security control assessment and risk determination process into a dynamic process...."

---

[11] SANS NewsBites - Volume: XIII, Issue: 45, "FISMA Compliance Metrics Focus on Continuous Monitoring (June 6, 2011)" available at http://www.sans.org/newsletters/newsbites/newsbites.php?vol=13&issue=45&rss=Y#sID200.

[12] NASA Office of the Chief Information Officer, "NASA Information System Security: The Path Forward with Automated Continuous Monitoring," August 04, 2010 *available at* http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2010-08/NASA-Continous-Monitoring-Program.pdf

Operating under the philosophy that "what get's measured, get's improved," the NASA presentation outlined the agency's approach to utilizing continuous monitoring to improve security. Three of the points made in the presentation are of particular importance for understanding how the agency's revamped security environment underpinned the EOS Security Team's continuous monitoring success story.

2.3.1 Continuous Monitoring

The need for automating security controls was the first key point made in the agency's presentation. The NASA official highlighted to the ISPAB that the agency "must move away from sporadic paperwork exercises to effective continuous monitoring."

In discussing NASA's Tools and Reference Architecture for continuous monitoring, the presentation discussed a point which will have broader significance for the work by the NASA EOS Security Team. Specifically, NASA noted that "Antivirus (AV) logs can also provide really good information on malware vectors into the environment."

Log data is worthless, except for forensic/post-mortem purposes, unless it can be analyzed in near real-time. As will be discussed, the ability to quickly analyze and integrate log data from multiple sources in novel and unexpected configurations was critical to the EOS Security Team in defending their systems following the RSA hack.

2.3.2 Risk Score Cards

The NASA presentation highlighted the agency's use of Risk Score Cards as a mechanism for providing actionable data about a NASA facility's security performance. NASA's IT security chief discussed the ongoing development of the Risk Score Cards in an interview with a trade publication prior to the ISPAB presentation. In that interview, the official explained,

> *Those risk scorecards will have a drill down capability that will let a center know why they have a particular score for their center; they can drill down to a single system. Let's say if they got a C for the week and they can drill down and see why they got a C; it may be because there is a particular system that needs a critical patch and our policy says you have got to patch it in X number of days and it has now gone past those number of days and that has brought their score down and then they have the opportunity to bring their scores up by applying those patches.*[13]

It should be noted that NASA made use of one of the agency's special resources, the Jet Propulsion Laboratory, managed by the California Institute of Technology, in developing the mathematics behind the automated scorecards. NASA also worked with the State Department which uses their own scorecard

---

[13] E. Chabrow, "Switch to Continuous Monitoring Requires New Skills," GovInfoSecurity.com, June 17, 2010, *available at* http://www.govinfosecurity.com/articles.php?art_id=2660&opg=1.

system and has been recognized as a leader in continuous monitoring, as discussed in Section 3.0.1[14] The NASA official discussed their interagency cooperation in ISCM activities when noting that,

> *We work very closely with (Deputy CIO/Security) John Streufert from the State Department on the scorecards that they use and the mathematics behind that. We are working with our folks at Jet Propulsion Laboratories on the math side of it again. It is all automated, what I call middleware or this engine that does correlation. It takes all of this information from the various tools and information about the systems and it crunches it and does some magic and out comes the score.*

Risk Scoring is an important aspect of ISCM data analysis and is an integral component of the Department of Homeland Security's (DHS') CAESARS [Continuous Asset Evaluation, Situational Awareness and Risk Scoring] Reference Architecture and the NIST/DHS CAESARS Framework Extension discussed in Section 3.0.1.

2.3.3 Attack Tree

The third point that should be noted from the ISPAB presentation is the agency's use of an "Attack Tree" diagram to illustrate key security vulnerabilities and how they interrelate – a visualization of the results from the agency's risk assessment process. Attack trees are diagrams of IT threats which "provide a formal, methodical way of describing the security of systems, based on varying attacks. Basically, you represent attacks against a system in a tree structure, with the goal as the root node and different ways of achieving that goal as leaf nodes."[15]

The Attack Tree in Figure 1 on the next page is an extract from the NASA IPSAB presentation. It shows five stages of cyberattack: Reconnaissance; Targeted Attack; Compromise + Network Intrusion; Installation of Tools/Utilities; and Malicious Endeavors.

For each stage of attack, the chart displays potential methods (nodes on the tree) for an adversary to accomplish the attack along with the nodes' connections to other nodes and other stages of attack. The nodes which NASA determined have a high frequency of incidence at the agency are highlighted in yellow. The presentation noted that "these nodes represent the areas where NASA should focus attention in order to ensure the greatest measurable improvement in the overall Agency security posture."

---

[14] Additional information about the State Department's continuous monitoring work may be found in their Continuous Monitoring Case Study Update presentation to the ISPAB *available at* http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2011-03/ISPAB-FISMA-Continuous-Monitoring-JStreufert.pdf

[15] B. Schneier, "Attack Trees," Dr. Dobb's Journal, December 1999 *available at* http://www.schneier.com/paper-attacktrees-ddj-ft.html.

Thus, the chart shows that Phishing Campaigns[16] and Social Engineering[17] are the type of malevolent reconnaissance which have been most successful at NASA as well as the pathways that attacks initiated through these means could take.

Figure 1: NASA OCIO Attack Tree

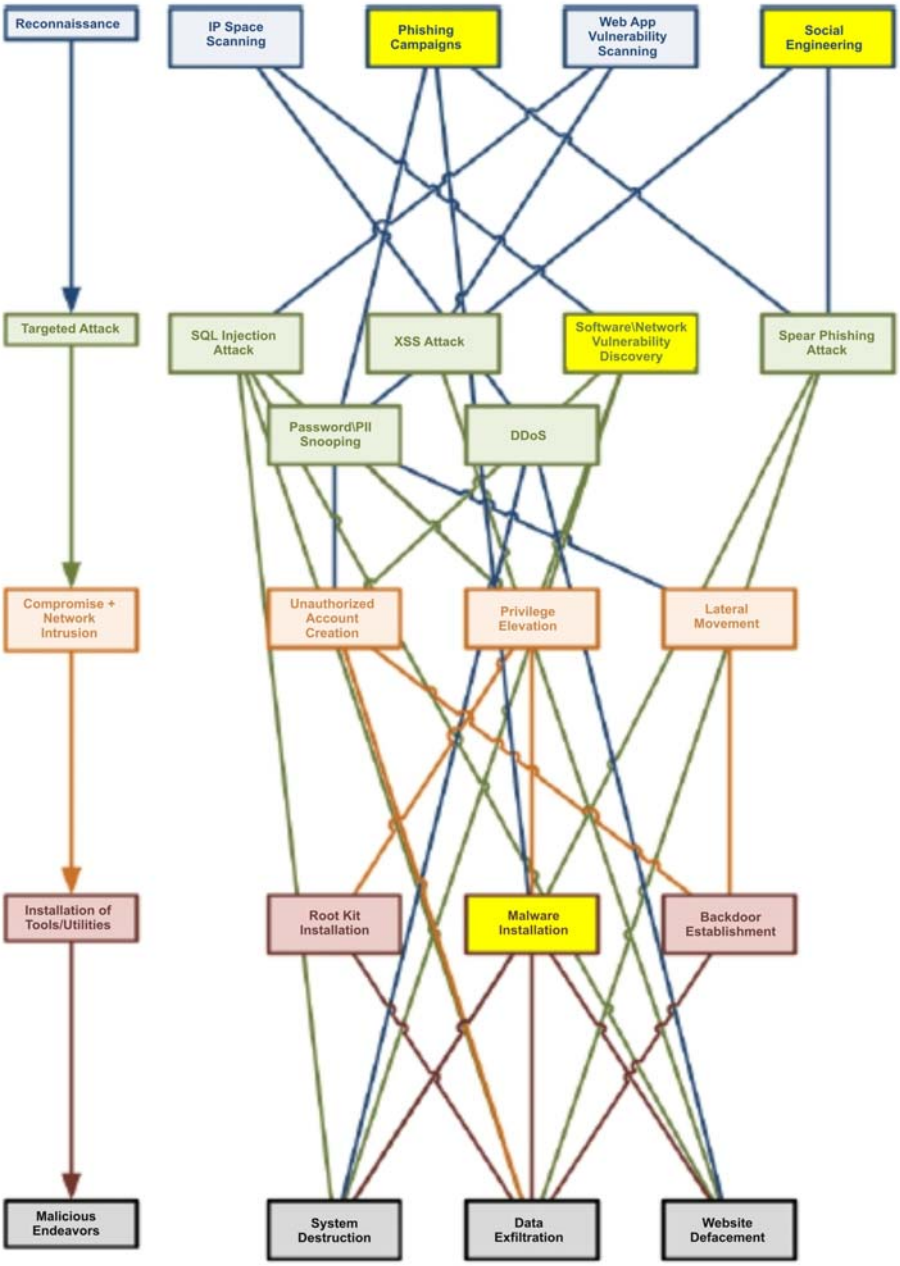

[16] See, Wikipedia, *available at* http://en.wikipedia.org/wiki/Phishing.

[17] See, Wikipedia, *available at* http://en.wikipedia.org/wiki/Social_engineering_(security).

NASA summarized the purpose of the attack tree by stating that it helps "Identify the attacker's modus operandi from end-to-end and then implementing controls that shunt their capabilities. From there it's just continuous monitoring of those controls."

3.0 Best Continuous Monitoring Practices: NASA EOS Security Team Use of Splunk

Several agencies have made significant contributions to advancing federal continuous monitoring practices. The CAESARS Framework Extension, discussed below, is itself a best practices document based on the work of multiple agencies. Instead of focusing on the important continuous monitoring work already codified, this CRE Best Practices study analyzes a specific instance in which an agency successfully used continuous monitoring best practices to address an unexpected development. The study also compares the agency actions with a set of best practice principles for continuous monitoring derived from the NIST/DHS CAEASARS Framework Extension.

3.0.1 CAESARS Framework Extension – Continuous Monitoring Best Practices

The draft NIST/DHS CAESARS Framework Extension is an essential reference work for understanding ISCM. The Framework Extension built on DHS' 2010 *CAESARS Reference Architecture Report*, version 1.8 which was developed in response to an OMB directive to the Departments of State, Justice and Treasury/IRS to "coordinate with the Department of Homeland Security (DHS) to evaluate their continuous monitoring (CM) best practices and scale them across the government."[18]

One of the purposes of the Framework Extension was to expand the applicability of CAESARS to the entire US government as well as to industry and State and Tribal governments.[19] The Framework Extension explains that,

> *The end goal of CAESARS FE is to enable enterprise CM by presenting a technical reference architecture that allows organizations to aggregate collected data from across a diverse set of security tools, analyze that data, perform scoring, enable user queries, and provide overall situational awareness.*

Thus, the NIST/DHS best practices document provides practical guidance as to the capabilities a continuous monitoring regime should provide. Based on the CAESARS FE, there are five principles to guide development of an organization's computational and human continuous monitoring capabilities. Based on the NIST/DHS document, an agencies IT staff should engage in the following continuous monitoring best practices:

---

[18] NIST Interagency Report 7756 (Draft), "CAESARS Framework Extension: An Enterprise Continuous Monitoring Technical Reference Architecture (Draft)" February 2011, p. 1 *available at* http://csrc.nist.gov/publications/drafts/nistir-7756/Draft-nistir-7756_feb2011.pdf.

[19] Id.

- Aggregate data from "across a diverse set" of security tool sources;
- Analyze the multi-source data;
- Engage in explorations of the data based on changing needs;
- Make quantitative use of the data for security (not just reporting) purposes including the development and use of risk scores; and
- Maintain actionable awareness of the changing security situation on a real-time basis.

The above CAESARS FE-derived principles define what CRE means by Continuous Monitoring Best Practices. The Principles are stated in a somewhat more concise form in Section 3.5.

3.0.2 SP 800-53 Technical Controls – Continuous Monitoring Prerequisites

Implementing the above Best Practices principles requires that security staffs have and are properly trained in use of a continuous monitoring tool set that has all of the requisite technical capabilities. The complete list of security control baselines for low, moderate and high-impact information systems is located in Appendix D of SP 800-53. The controls are divided in categories, such as Access Controls which includes 22 specific controls.

A chart listing the categories of technical controls relevant to continuous monitoring is located in Appendix B of CAESARS v. 1.8. As the Appendix notes, the chart "provides a template for mapping tools needed to conduct continuous risk analysis and scoring as described in this reference architecture."

The categories of Technical Controls listed in Appendix B are:

- Identification & Authentication (IA);
- Access Control (AC);
- Audit and Accountability (AU); and
- System & Communications Protection (SC).

Implementing the continuous monitoring best practice principles described in Section 3.0.1 requires that IT security staff have the technical control tools from SP 800-53 in place.

3.1 NASA EOS Security Team

NASA's Earth Observing System (EOS) is a component of the agency's Earth Science Division located in their Science Mission Directorate. The goal of NASA EOS is to enable "an improved understanding of the Earth as an integrated system." NASA EOS consists of "a coordinated series of polar-orbiting and

low inclination satellites for long-term global observations of the land surface, biosphere, solid Earth, atmosphere, and oceans."

Thus, NASA EOS's duties include collecting, analyzing and disseminating massive quantities of data which need to be protected from unauthorized use and tampering.

The NASA EOS Security Team was composed of a contractor from L-3 Communications and an additional person. Additional information about NASA EOS may be found on their website, http://eospso.gsfc.nasa.gov/.

3.2 Splunk Software for Continuous Monitoring

Splunk has developed a FISMA Continuous Monitoring App. The Splunk App "builds upon the core capabilities of Splunk Enterprise software to index and provide visibility into the machine data generated by agency IT systems and infrastructure - whether physical, virtual or in the cloud - to align agency security operations to FISMA controls, including real-time views of NIST 800-53 controls." The company stated that the "Splunk App for FISMA Continuous Monitoring helps federal security teams meet compliance challenges while supporting timely security incident response."[20]

3.3 The RSA Hack

RSA is major vendor of IT security products. RSA's SecurID® is a widely-used two-factor authentication system that requires users provide data displayed on a device "token" in the user's possession along with their user name and password. The pseudo-random set of digits displayed on the RSA token changes every 60 seconds. Accessing a protected system requires that the user enters their user name, matching password, and the current set of digits from the SecurID® token. The two-factor system reduces the possibility of unauthorized access through such common security breaches as password-sharing or lost/stolen tokens, since neither the user name/password nor the token on its own can be used to gain system access. The two-factor system is also intended to significantly increase the difficulty of a brute-force attack.

In March 2011, a senior RSA official posted notice on the company website that they had been hacked. The communication stated,

> *Recently, our security systems identified an extremely sophisticated cyber attack in progress being mounted against RSA. . . . Our investigation has led us to believe that the attack is in the category of an Advanced Persistent Threat (APT). Our investigation also revealed that the attack resulted in certain information being extracted from RSA's systems. Some of that information is specifically related to RSA's SecurID two-factor*

[20] "Splunk Develops App for FISMA Continuous Monitoring" *available at* http://www.splunk.com/view/SP-CAAAGCZ.

> *authentication products. While at this time we are confident that the information extracted does not enable a successful direct attack on any of our RSA SecurID customers, this information could potentially be used to reduce the effectiveness of a current two-factor authentication implementation as part of a broader attack. . . . we strongly urge you to follow the steps we've outlined in our SecurCare Online Note.*[21] [Emphasis added]

RSA later reported that the hack was launched via a phishing email using a "zero day" exploit of Adobe Flash contained in an Excel file. Zero day exploits refer to attacks using unreported software vulnerabilities. As was noted, NASA's IT security leadership had identified phishing campaigns as a high frequency attack vector. RSA explained that "two emails were sent to two small groups of employees; you wouldn't consider these users particularly high profile or high value targets. The email subject line read '2011 Recruitment Plan.'"[22]

The spreadsheet contained "a zero-day exploit that installs a backdoor through an Adobe Flash vulnerability (CVE-2011-0609). . . . Adobe has released a patch for the zero-day, so it can no longer be used to inject malware onto patched machines." After exploring their capabilities gained through the exploit, the attackers "went into the servers of interest, removed data and moved it to internal staging servers where the data was aggregated, compressed and encrypted for extraction." The attacker then "used FTP to transfer many password protected RAR files from the RSA file server to an outside staging server at an external, compromised machine at a hosting provider" where they were obtained by the attacker.

Based on the information from RSA, the extent to which customer systems could be at increased risk was not clear. RSA did not publicly disclose what information had been stolen. Cyber-prudence and RSA dictated that customers take increased security measures. As one trade publication noted,

> *If attackers were able to access the seeds for a specific company, they might be able to generate the pseudo-random numbers of one of its tokens, allowing them to clear a crucial hurdle in breaching the company's security.*
>
> *Other possibilities include the theft of source code that gives attackers a blueprint of vulnerabilities to exploit, or the theft of private cryptographic keys that might allow them to imitate RSA servers or register new employee tokens.*[23]

---

[21] A. Coviello, "Open Letter to RSA Customers" *available at* http://www.rsa.com/node.aspx?id=3872.

[22] U. Rivner, "Anatomy of an Attack," Speaking of Security: The Official RSA Blog and Podcaast," April 1, 2011, *available at* http://blogs.rsa.com/rivner/anatomy-of-an-attack/

[23] "RSA breach leaks data for hacking SecurID tokens," The Register, March 18, 2011, available at http://www.theregister.co.uk/2011/03/18/rsa_breach_leaks_securid_data/.

3.3.1 Advanced Persistent Threats

The term "Advanced Persistent Threats" (APTs) refers to a wide variety of long-term, sophisticated cyberattacks. Government Computer News (GCN) explained that APT "is a descriptive rather than technical term that describes a broad class of attacks."[24] GCN quoted an RSA official, prior to the attack, explaining that one of the distinguishing characteristics of an APT is that it "is targeted, going after high-value assets, such as intellectual property, that can provide a return on the expense of sophisticated, possibly one-of-a-kind attacks."

NIST SP 800-137, *Information Security Continuous Monitoring (ICSM) for Federal Information Systems and Organizations*, defines APTs as being:

> *An adversary with sophisticated levels of expertise and significant resources, allowing it through the use of multiple different attack vectors (e.g., cyber, physical, and deception) to generate opportunities to achieve its objectives, which are typically to establish and extend footholds within the information technology infrastructure of organizations for purposes of continually exfiltrating information and/or to undermine or impede critical aspects of a mission, program, or organization, or place itself in a position to do so in the future; moreover, the advanced persistent threat pursues its objectives repeatedly over an extended period of time, adapting to a defender's efforts to resist it, and with determination to maintain the level of interaction needed to execute its objectives.*[25]

As RSA explained, "APTs do not 'defeat' security products. They just find ways to fly below the existing technology."[26] Thus, an essential attribute of an effective continuous monitoring system is that it can adapt as needed to successfully defend against an APT. The NASA EOS Security Team's ability to quickly and successfully defend against the unexpected APT breach of RSA was an important factor in their being selected for this Best Practices study.

3.4 NASA EOS Defends Against an APT

Little information has been publicly disclosed about the RSA hack other than it was an APT which achieved some measure of success. The extent to which the RSA hack would allow a hostile person or organization an advantage in attacking one or more of RSA's government or industry customers is not part of the public record and may or may not be known to RSA itself.

---

[24] W. Jackson, "Advanced persistent threats are a new way of life," Government Computer News, April 1, 2011, available at http://gcn.com/articles/2011/04/04/cybereye-apt-advanced-persistent-threats-rsa.aspx.

[25] NIST, Special Publication 800-137, September 2011, p. B-1.

[26] Id., Anatomy of an Attack.

In that the RSA attack at least partially succeeded, their clients had to consider that the SecurID® authentication system was at least potentially compromised and constituted a potential vector for an APT attack against their own systems. This is a situation where an organization's continuous monitoring capabilities, in terms of both tools and ability to conceptualize, initiate and follow-through on a response strategy, are crucial.

If the NASA EOS Security Team was not able develop an effective response to the unexpected threat posed by the RSA hack or if their continuous monitoring tools could not be rapidly adapted to implement the new strategy, the agency's IT security was at an unacceptable risk of failure.

Thus, maintaining cybersecurity depends on:

1. Human capabilities. Senior security staff needs the creativity, flexibility and authority to act against a new, poorly understood threat. Similarly, all IT security professionals need the technical skills, resources, and work ethic necessary to effectively execute the response strategy.

2. Software capability. Agencies need the continuous monitoring capabilities specified in SP 800-53 and discussed in Section 3.0.2. It is only because NASA's continuous monitoring tools were sufficiently flexible, powerful and user-friendly to be employed in a novel configuration on-the-fly that the Security Team's strategy could be conducted.

It is because NASA EOS' human and software security capabilities were both up to the task that the Continuous Monitoring Best Practice is the EOS Security Team's use of Splunk rather than the personnel or software in isolation. No effective security program can rely only on human or automated capabilities.

3.4.1 The Baseline Situation

The NASA EOS Security Team stated that their baseline situation prior to the hack was working "quite well" and providing "premier security."[27]

3.4.1.1 NASA EOS Use of RSA

NASA EOS uses RSA SecurID® as the basis for their two-factor authentication system. One of the advantages of the RSA security tool is that it compatible with NASA EOS' RADIUS (Remote Authentication Dial In User Service) configuration. RADIUS is a protocol for communications between a Network Access Server (NAS) and a RADIUS server. The RADIUS protocol supports authentication,

[27] Information about NASA EOS security response to the RSA hack was taken from a presentation by T. Meader at the Splunk>Live! conference held in Washington, DC on May 12, 2011.

authorization and accounting functions and is intended as a means by which remote users can securely access the NAS.[28]

The RSA token system is also compatible with a variety of devices and third-party software packages, including NASA EOS' continuous monitoring software.

3.4.1.2 NASA EOS Use of Splunk

NASA EOS was using Splunk v. 4.2 for continuous monitoring. The Security Team had significant experience with the product having started with version 2.0 over three years ago. The product provided easy access to and usable information concerning NASA EOS' "voluminous" firewall logs which were generated by virtually all major firewall vendors.

The continuous monitoring software was being used to process a substantial number of daily events. About half the data was from firewalls and the rest from other security and system logs. NASA was in the process of re-architecting their use of Splunk to provide for longer data retention and performance improvements in preparation for expected larger growth in systems usage.

The Security Team noted that the continuous monitoring product easily correlated IDS/IPS (Intrusion Detection System/Intrusion Prevention System) log and host/device logs, along with firewall log data, to "track the path of activity throughout the system." It should be noted that IDS and IPS are each distinct approaches to security which may be layered to provide enhanced security.

In brief, IDS is "the art of detecting inappropriate, incorrect, or anomalous activity."[29] IPS, by contrast, "is used to actively drop packets of data or disconnect connections that contain unauthorised data." Not surprisingly, an IPS can be configured to act on IDS data. The two approaches to system security can be combined in a single device.

In addition to correlating IDS/IPS data, NASA EOS also used the software to perform HIDS (Host-based Intrusion Detection System) analysis. HIDS analysis is based on data collected from a specific computer system and allows analysis of "activities to determine exactly which processes and users are involved in an attack on a particular system or host. HIDS can see the outcome of an attempted attack, as they can directly access and monitor the data files and operating system processes targeted by the attack."[30] The

---

[28] CISCO, "How Does RADIUS Work?" *available at* http://www.cisco.com/en/US/tech/tk59/technologies_tech_note09186a00800945cc.shtml.

[29] T. Holland, SANS Institute, "Understanding IPS and IDS: Using IPS and IDS together for Defense in Depth," February 23, 2004, *available at* http://www.sans.org/reading_room/whitepapers/detection/understanding-ips-ids-ips-ids-defense-in-depth_1381

[30] State of Missouri, "Compliance Component" *available at* http://oa.mo.gov/itsd/cio/architecture/domains/security/CC-HostBasedIDS040303.pdf.

functions and output of HIDS may includes "log analysis, integrity checking, Windows registry monitoring, rootkit detection, real-time alerting and active response."[31]

NASA EOS managed and tracked changes in their use of the software product through an OSSEC (Open Source SECurity) interface obtained from the vendor.

3.4.2 The EOS Security Team Response

NASA received an email from RSA notifying them of the security breach. Agency security officials quickly decided that simply scrapping the RSA system and moving to another vendor was not a feasible option.

As was noted, customers were not told what RSA information was stolen and thus did not know the extent to which their security was undermined. NASA, and presumably other RSA customers, were told to protect the serial number on the tokens which suggested to the agency that the map linking tokens to authorizes users had been exfiltrated from RSA.

If an attacker knew which tokens went to which authorized users, the two-factor authentication system was potentially compromised. If an attacker had the map (a distinct possibility) along with the token's "seed" used with an algorithm to generate the token's pseudo-random number (whether this data was compromised is unknown), it would give them a significant advantage in any effort to gain unlawful access to the NASA EOS systems.

As NASA's EOS Security Team explained, a successful exploitation would require knowing the token serial number (assuming the pseudo-random number was compromised), the user it went to, and the user's password. It should be noted that an attacker could use brute force (multiple attempts to penetrate the system) to make up for data, such as a user password, they lacked. A key function of two-factor authentication is to render most brute force attacks unfeasible by requiring too much information not known to the attacker.

The issue then for the NASA EOS Security was to devise a response strategy given the Rumsfeldian situation where there were known knowns (that RSA had been hacked, potentially compromising the tokens), known unknowns (whether the attacker would be able to replicate a token's pseudo-random number and map it to a user), unknown unknowns (there are always unrecognized information gaps) and the team was limited to the resources they had in place.

3.4.2.1 Developing a Response Strategy

The simplest and most effective security option, using a blanket "deny" at the firewall, was not consistent with the agency's mission of maintaining public access to EOS data. Since the RSA hack made it

---

[31] OSSECWiki *available at* http://www.ossec.net/wiki/Faq:Whatis#1.01_-_What_is_an_HIDS_.28Host-based_Intrusion_Detection_System.29.3F

potentially easier for an attack to succeed, the Security Team focused their response activities on looking for unusual traffic or other behavior of person(s) attempting to gain system access.

In short, the response strategy NASA EOS developed was a continuous monitoring strategy aimed at:

1. Identifying unusual traffic/user behavior;

2. Linking unusual behavior to the exact source(s) of the behavior; and

3. Terminating traffic from the attacker(s).

3.4.2.2 Data Needed for the Response Strategy

Executing the response strategy required data from multiple sources. What made the situation particularly difficult was that not all of the data was in NASA's possession. The two-factor authentication process meant that the authentication data from login attempts passed through RSA's servers; only RSA's computers are able to verify whether or not the numeric sequence a user enters from their SecurID token is indeed the correct set of numbers. Thus, it was RSA, not NASA, that had the logs of attempts to use the token. NASA EOS firewall logs show attempts to enter a user name and password to gain access to the server, but not the RSA portion of the authentication process.

Moreover, NASA EOS also needed data from certain remote access points that passed through their Cisco Virtual Private Network (VPN).

The NASA EOS Security Team, therefore, needed data from three distinct sources (not merely from three different vendor's products):

- NASA EOS firewall logs;
- RSA logs; and
- Logs from the Cisco VPN.

Obtaining the raw data from the various sources was not difficult, making use of the data was another matter. NASA EOS needed to extract the data certain from the logs (obtain the relevant fields) and compare the data obtained from completely different sources.

3.4.2.3 Executing the Response Strategy

Splunk's field extractor graphical user interface (GUI) was used for extracting some of the data. The EOS Security Team found that using a regexe (regular expressions) process was a better way to proceed. Regexes are a formal methodology for pattern recognition. A regular expression can be thought of as a "powerful way to select data that matches a pattern, as well as to manipulate, rearrange, and change that

data."[32] NASA used Expresso, a regular expression development tool, for extracting much of the data from the logs.

As was noted, NASA EOS needed to extract data from the RSA logs as well as their own firewalls. Manipulation of the Cisco VPN data was also required.

A key issue the team had to get perfectly was the timestamp from each set of logs. The timestamps were essential to identifying events occurring simultaneously on different systems at different locations of different organizations.

Once the log data was extracted, NASA EOS used Splunk to create searches to identify anomalous login behavior. The security team was also able to use the tool to schedule searches and to set alerts depending on the search results.

NASA EOS created "at a glance" views of recent RSA traffic and other "dashboards" to provide a quick visual display of what was happening on a real time basis.

One critical step that the NASA EOS EOS undertook was to link both the firewall logs data and the VPN log data with the RSA authentication traffic. It is this crucial step which allowed the team to identify the internet protocol (IP) address from which attacks were originating. This allowed staff to cut-off access from the suspect IPs while leaving routine access unaffected.

Once the Security Team developed their methodology for combating the threat, they created a continuous monitoring app to automate the tasks on an ongoing basis including providing alerts as needed and showing real time status updates on a dashboard. It is important to note that no attacker breached NASA EOS IT security.

3.5 Concordance Between ISCM Best Practice Principles and NASA EOS Actions

In developing and executing their response strategy, the NASA EOS Security Team adhered to the continuous monitoring best practices principles derived from the CAESARS Framework Extension. As was noted, using these principles required the team having, and knowing how to use, the continuous monitoring technical tools specified in SP 800-53.

The following chart illustrates NASA's use of the five Federal Cybersecurity Continuous Monitoring Best Practice Principles:

---

[32] Oracle Regular Expressions Pocket Reference, available at http://regexlib.com/?AspxAutoDetectCookieSupport=1.

Best Practices: Continuous Monitoring

- **Principle 1: Aggregate Diverse Data**
  The EOS Security Team had the tools in place and had the skills to combine data from multiple sources generated by different products/ vendors and organizations in real time.

- **Principle 2: Maintain Real-Time Actionable Awareness**
  The NASA EOS security staff developed real time dashboards to allow them to see the attacks-related metrics and set real time alerts to detect anomalous changes in various systems status.

- **Principle 3: Analyze Multi-Source Machine Generated Data**
  Comparison of large data sources from multiple systems and applications was undertaken and accomplished.

- **Principle 4: Create Real-Time Data Searches**
  The IT security staff developed and automated "Google-style" searches across unrelated data sets to identify the IP addresses from which attacks were originating.

- **Principle 5: Transform Data Into Actionable Intelligence**
  IT security staff analyzed the data to identify specific IP addresses from which attacks originated and terminated hostile traffic.

4.0 Lessons Learned

1. **Leadership**. Leadership from OMB and senior NASA IT security officials was crucial in empowering the NASA EOS Security Team.

2. **Human and Software Capabilities**. Effective continuous monitoring requires skilled human resources and software with the technical controls specified by NIST SP 800-53.

3. **Real-Time Monitoring and Analysis**. There is no substitute for IT security staff being able to monitor and analyze diverse security-related data on a continuous basis.