Security Content Automation Protocol (SCAP) Purpose, Measurements, and Expected Outcomes

What is SCAP?

The Security Content Automation Protocol (SCAP) is a synthesis of interoperable specifications derived from community ideas. More than 3,000 people from diverse parts of industry, academia, and government participated in workshops and webinars around the country. NIST received hundreds of detailed suggestions and comments in response to requests for information (RFI) and feedback on several draft versions of the Framework. Comments from private sector stakeholders can be found at the RFI and Preliminary Cybersecurity Framework Comments(link is external) Web pages of NIST's Cybersecurity Framework Web site. Community participation is a great strength for SCAP, because the security automation community ensures the broadest possible range of use cases is reflected in SCAP functionality. The outcome of this collaboration is commonly referred to as the NIST Standards. The outcome of this collaborative effort is the practice referred to as Standards-based System Hardening. A properly hardened host inhibits threat actors from performing the core functions associated with exploitation such as Lateral Movement, Privilege Escalation, and Data Compromise. When combined with effective monitoring and proactive security patching, these practices can have a significant positive effect in reducing the host overall threat surface and exploitability.

Measuring SCAP Guidelines

SCAP articulates these standards using two general methods. Non-Automated and Automated Checklists. Non-Automated Checklists are designed to allow for manual measurement and implementation of the standards. These checklists can contain hundreds of pages with several hundred controls; a potentially time-consuming and cumbersome measurement task. NIST developed a standards-based XML implementation for consumption and measurement of standards via automated methods; referred to as SCAP. It's important to understand that not all of the controls can be measured via automated means; requiring some manual review for a full determination. NIST established the United Stated Government Computing Baseline (USCGB) as its baseline compliance set. A variance of the USCGB is the DOD implementation. Requiring a higher level of surety, the DOD developed a more refined derivative using the NIST USCGB as a baseline, referred to as the Security Technical Implementation Guide (STIG). STIGS differed from USCGB in two distinct areas. First, the DOD has a long-standing practice of standardizing tools and software via an Approved Products List (APL). This allowed the DOD to uniquely specify granular settings for mandated controls (e.g. Host-Based Intrusion Detection). This also allows the DOD to apply an additional level of hardening, designed to better withstand focused advocacy exploitation through the implementation of a more risk-adverse configuration. As an example, the DOD typically operates at a SCAP score range of 97% while the our organization, external to the larger managed DOD environment, operates within the 90% range.

Understanding NIST Guidelines Terminology regarding "Not Applicable", "Not Defined", and "Not Configured" Standards

With regards to the NIST Standards, some controls are addressed for completeness but do not specify a requirement. These are defined in the guidelines as "Not …". The "Not Applicable" remark signifies that the setting is not available in that version of Windows. For example, there are many new settings in Windows Vista that will have no effect on computers running Windows XP (e.g. settings for the Windows Firewall with Advanced Security). The "Not Defined" and "Not Configured" remarks are functionally equivalent and mean that the USGCB does not require a specific value for that setting; agencies are free to configure those settings however they wish. An important clarification regarding these values is with regard to their implementation. The "Not …" standard represents USCGB's position regarding the controls measurement for compliance and not the value set in the measured hosts Windows registry; this is an important differentiation which is detailed later in this paper.

Security Content Automation Protocol (SCAP) Purpose, Measurements, and Expected Outcomes

Ending Support for USCGB

A series of shifts within government began to occur. In March 2014, the DOD issued DODI 8510.01 that mandated the end of the DOD's FISMA (a competing compliance standard to NIST) and ordered the full adoption of NIST 800-53. Following this consolidation of government standards, NIST issued Special Publication 800-70 revision 3, further consolidating the methodology for determining the appropriate checklist for measuring compliance. About the same time, NIST shuttered most of its efforts regarding the issuing of new USCGB SCAP templates. When inquiring about this again, as last as October 2017 seeking guidance, NIST USCGB referred government users to the more secure STIG Templates which are available at "IASE.disa.mil/stigs" as the preferred Government template.

Using SCAP as a Guideline

Every operating environment is different with regards to the level and manner it processes non-public content, the risk tolerance of the organization, and the operational needs within the enterprise. SCAP provides a best practice approach to the analysis of the hosts role. Different roles will require differing levels of hardening. For instance, public Kiosks will likely have significant hardening, whereas air-gapped hosts running legacy software might have a different standard. It is up to the data owner to assess those recommendations against the factors above to determine how / if those recommendations are appropriate to the compute environment. The SCAP score is more of a general comparison of overall hardening with regards to the NIST Best Business Practice and a statement regarding the level of cybersecurity maturity currently implemented within the organization. This is best represented with regards to how organizations approach baseline security controls for their Windows operating systems. Organizations who deploy Microsoft Windows operating systems are familiar with Microsoft's general practice regarding the initial OS deployment Out-of-Box Experience (OOBE). This experience favors functionality over security. The intent of the OOBE is to get the host up and running quickly, leaving the process of layering additional host hardening to the IT management team as a follow-on task. For comparison purposes, Windows legacy OS's (prior to Windows 10 / Windows Server 2012) typically achieve an OOBE SCAP score in the low to mid 30's while modern OS's (Windows 10 and later) achieve initial OOBE scores in the lower 40's. If a modern Windows OS achieves a SCAP score under 40, its typically the result of the application of group policies that deprecate security below the OOBE baseline.

Interpreting SCAP Findings

Data owners should evaluate their SCAP scores and the associated recommendations to determine which are appropriate within their enterprise. In some cases, a finding may not be relevant to the operating environment. This is especially true where the SCAP template used is a STIG template and the required value is DOD unique. A common example of this would be a failed logon banner finding. The DOD legal team requires the same logon warning banner on all hosts and devices. While this would show as a finding, the larger question the data owners should be asking is does the enterprise have a warning banner standard and is it applied to the measured host? Data owners should routinely challenge themselves to reevaluate their host hardening and seek methods to further protect the environment. While we have provided an example of a DOD unique setting, the value of the setting from an overall compliance perspective is relevant and valuable. The data owner should seek to standardize on these settings or comparable ones as applicable, where possible to raise their overall organizational cybersecurity maturity. Once an acceptable baseline and unique settings standard is agreed upon and documented by the applicable organizations compliance entity, SCAP templates can be manually altered to mirror unique organizational settings. While the process is time consuming, it can allow for a more focused measurement and identification of host compliance drift.

Security Content Automation Protocol (SCAP) Purpose, Measurements, and Expected Outcomes

Understanding Indeterminate Registry Values and their Impacts

A critical concept to understand when evaluating findings is the notion of defined registry values. This is best clarified through the discussion of the differences between configured and unconfigured. Standards-based organizations (e.g. NIST, CIS, etc…) require measured controls which are applicable to the standard (e.g. control defined other than "not …" within their respective standard) to be set to a specific measured value. For instance, if the compliance value expected is "Disabled", then that is the acceptable measured value. If the organization implements other compensating controls that cause the value to deviate or has not correctly implemented the required standard, then the metric will be measured as "Fail". This condition differs regarding the value of "Not Configured", which is considered an indeterminate value for compliance. This value represents the absence of a determination by the Data Owner with regards to how the setting should be controlled, regardless of the OS default outcome. When an enterprise GPO is set to "Not Configured" and a less restrictive value is set in local group policy, the host will implement the specified setting over the unspecified GPO setting. Threat actors can leverage "Not Configured" or uncontrolled entries to cause downgrade attacks through implementation of less restrictive local policy settings. Therefore, it is critical the data owner resolve all indeterminate findings via enterprise controlled GPO to prevent this attack surface from being exercised.

Implementing Harding Measures:

Host hardening is a complex task that takes critical analysis and extensive testing prior to a successful implementation. We recommend that testing be accomplished using a separate Organization Unit (OU) Test Container. It is a given that IT Operations and the relevant Business Unit may have differing perspectives on acceptable settings. This is where a layered testing approach that considers both perspectives is required for a successful deployment. A host hardened to the point of non-functionality is both secure and unusable; an unacceptable operational and business state. IT Operations must work to find ways to support IT Security requirements or develop compensating controls. It is the responsibility of C-Level management to deconflict these often-competing perspectives, ensuring the continued confidentiality, integrity, and availability of the enterprise is maintained. Consider these following practices when developing enterprise group policies:

1. Create Multiple GPO to Address Diversity within the Enterprise:
    a. Create and apply a Baseline Hardening GPO that apply to all hosts regardless of business unit.
        i. Places all common setting in this GPO

    b. One-Off and unique Business Unit Requirements should be placed in separate GPO's and applied to the applicable business units OU host container. If the organization does not manage hosts in this matter, then alternatively filter for the business unit host names using WMI filtering

2. When implementing logical isolation for hosts running applications required deprecated security, ensure GPO's implement a unique membership within the local administrator context on those hosts. Further protect those hosts with logical Access Control Lists (ACL's) that restrict traffic to the host to only those specified in the security plan of the host.