FedRAMP

Threat-Based Risk Profiling Methodology

Developed by: GSA FedRAMP PMO

Version 2.0

2/15/2022

GSA

info@fedramp.gov

fedramp.gov

DOCUMENT REVISION HISTORY

Date	Version	Page(s)	Description	Author
02/2021	1.0	All	Initial Publication	FedRAMP PMO
02/2022	2.0	All	Update to Methodology: Scored controls against the MITRE ATT&CK threat framework	FedRAMP PMO

TABLE OF CONTENTS

Acknowledgements

This publication was developed by the Federal Risk and Authorization Management Program (FedRAMP) with representatives from the Department of Homeland Security (DHS) Cybersecurity Infrastructure Security Agency (CISA) in an ongoing effort to produce a threat-based approach to risk management for the federal government. The FedRAMP team, Ashley Mahan (Acting Assistant Commissioner for Solutions), Brian Conrad (Acting Director of FedRAMP), and Zachary Baldwin (FedRAMP Program Manager for Strategy, Innovation, and Technology), wishes to acknowledge and thank their partners from the CISA .govCAR team, the Chief Information Officers (CIO) Council, General Service Administration's (GSA) 10x program, and members of the Volpe Information Technology Group, who provided support services as part of the research for this publication.

Organizational Affiliations

- **General Services Administration (GSA)**
    - FedRAMP PMO
    - 10x Program
    - Contractor Support - The Volpe Information Technology Group, Inc.
- **Department of Homeland Security (DHS) Cybersecurity Infrastructure Security Agency (CISA)**
    - .gov Cybersecurity Architecture Review Program (.govCAR Program)
    - Contractor Support - Johns Hopkins Applied Physics Laboratory (APL); MITRE Corporation
- **CIO Council**
- **Chief Information Security Officers (CISO) Council**

Scoring Teams

In addition to the above acknowledgments, a special note of thanks goes to the scoring team participants for their superb technical contributions. These scoring teams included the following individuals:

Organization	Scoring Members	
Department of Homeland Security (DHS) Cybersecurity	• Branko Bokan • David Otto	• Greg Bastien • Jim Quinn

Infrastructure Security Agency (CISA)	• Jody Patilla (Johns Hopkins APL) • Pete Dinsmore (Johns Hopkins APL) • Michael Smeltzer (Johns Hopkins APL)	• Edward Sweitzer (MITRE) • Kurt Beernink (MITRE)
Department of Interior (DOI)	• Min Oh	
General Services Administration (GSA)	• Scott Boger (Noblis) • Scott Williams (Noblis) • Ashley Taylor (Noblis)	• Tom Volpe Sr. (VITG) • Tom Volpe Jr. (VITG)

Executive Summary

FedRAMP promotes the adoption of secure cloud technology across the federal government by providing a standardized approach to security and risk assessment. FedRAMP aims to empower agencies to modernize operations using secure cloud solutions to improve agencies' information technology (IT) security. FedRAMP successfully made the authorization process more efficient by standardizing the security control requirements for cloud systems which enables security authorization package re-use.

In 2017, the Office of American Innovation (OAI) sponsored a feasibility study, coordinated by the Office of Management and Budget (OMB) and managed by the GSA FedRAMP Program Management Office (PMO). The objective of the study was to determine the feasibility of an agile approach to authorizations. It was determined that an agile approach to authorizations was feasible if a defensible methodology was established to prioritize controls.

FedRAMP, in collaboration with the DHS CISA .govCAR team, developed a methodology for scoring each National Institute of Standards and Technology (NIST) Special Publication (SP) 800-53 security control against threat frameworks to determine which security controls and capabilities are most effective to protect, detect, and respond to current prevalent threats.

From July 2019 until June 2020, the govCAR team worked with GSA to score the NIST 800-53 Rev 4 control baseline against the National Security Agency's (NSA)/CSS Technical Cyber Threat Framework v2 (NTCTF). In September 2020, NIST 800-53 Rev 5 was released, and the .govCAR team migrated to the MITRE ATT&CK Framework version 8.2 as the NTCTF was discontinued. In February of 2021, the govCAR team worked with GSA to update scoring to align with NIST 800-53 Rev 5 control baseline against the MITRE ATT&CK Framework.

The goal of this initiative is to enable agencies, Cloud Service Providers (CSPs), and other industry partners to prioritize security controls that are relevant and effective against the current threat environment. This leads to informed, quantitative-based risk management decisions in authorizing information systems for government use.

This white paper outlines the methodology behind the threat-based scoring approach and informs stakeholders of potential applications.

The prioritization of controls, based on protection values scored against real world threats, will help shift the cybersecurity paradigm from compliance to informed risk management.

Introduction

Cybersecurity is an essential part of the federal government's IT infrastructure and operations. FedRAMP established uniform security baselines (High, Moderate, Low, and Tailored) and standardized a repeatable authorization process for government officials when authorizing cloud systems. As many organizations have limited resources to combat a vast environment of dynamic threats, there may be an inherent acceptance of more risk, presenting the opportunity to prioritize inherent risks based on efficacy against the most prevalent real world threats.

Organizations need to prioritize their cybersecurity investments to utilize resources effectively and reduce the greatest amount of risk. Standards such as the NIST Cybersecurity Framework (CSF) and the Risk Management Framework (RMF) provide the foundation for achieving additional levels of security. When these frameworks are combined with real cybersecurity threat intelligence, a structured methodology for risk profiling and risk mitigation emerges.

The FedRAMP PMO, in partnership with the DHS CISA .govCAR Team, developed a threat-based framework and scoring methodology to prioritize NIST SP 800-53 security controls. The scoring methodology was adopted from the Department of Defense (DOD) Cybersecurity Analysis and Review (DoDCAR) and .govCAR. FedRAMP applied this scoring methodology using the following frameworks against NIST's baselines.

- FedRAMP analyzed each NIST SP 800-53, rev 4. control within the FedRAMP moderate baseline on its ability to protect, detect, and/or respond to each of the threat actions outlined in the NSA/CSS Technical Cyber Threat Framework.
- FedRAMP analyzed each NIST SP 800-53, rev. 5 control within the FedRAMP High baseline on their ability to protect, detect, and/or respond to each of the techniques outlined in the MITRE ATT&CK Framework version 8.2.

Application of the threat-based scoring methodology enabled the prioritization of controls and controls items (i.e., specific countermeasures/protection capabilities) based on their efficacy to protect against real world threats.

.govCAR Scoring Methodology

The .govCAR scoring methodology provides an end-to-end holistic assessment of cybersecurity capabilities provided by DHS CISA and representative cybersecurity architectures of federal agencies. The results of the iterative assessment are being used to inform CISA's approach to assisting agencies with insight and knowledge to make prioritized cybersecurity investment decisions to enhance cybersecurity and reduce risk.

DoDCAR introduced the concept of a threat-based, end-to-end analysis of a typical cybersecurity architecture. It was used to provide direction and justification for cybersecurity investments during the DoD financial planning process. DHS developed an organization, known as .govCAR, based on the DoDCAR model. DHS .govCAR produces results in increments or "spins," where each spin comprises a set of

cybersecurity capabilities for security architecture assessment. The benefit of adapting this methodology and applying it to risk profiling include:

- The use of a proven, standardized, and repeatable process to score capabilities against threats
- The use of a well-defined set of definitions and a scoring rubric

Threat-Based Risk Profiling Methodology

We developed a comprehensive methodology to attain an effective threat-based approach to risk profiling. This methodology consists of three phases:

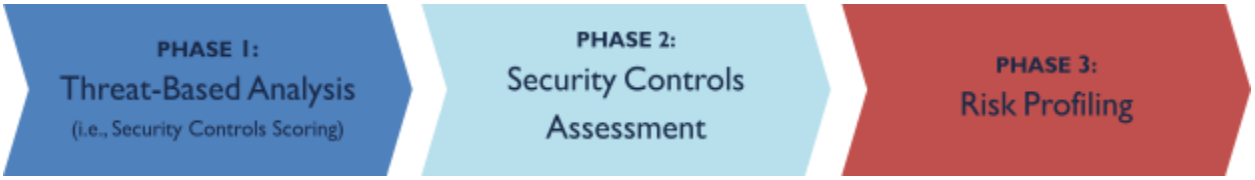

Phase 1: Threat-Based Analysis (i.e., Security Controls Scoring)

At the outset of this endeavor, the scoring teams recognized that a baseline of acceptable implementation parameters needed to be defined. With current processes, Agencies or organizations are required to define their own implementation parameters for a subset of the NIST security controls[1], which contain embedded assignment and selection statements. This approach can result in differing security implementations that need to be reviewed individually by each agency to determine acceptability. Normalizing these parameters creates the ability to avoid potential roadblocks in achieving maximum cloud adoption among the federal agencies as it may increase the reuse of security authorization packages from agency to agency and/or decrease the level of effort for each authorization.

After an extensive analysis of data provided by the CISO Council, a set of common values for these parameters was identified. These common values were compared against the FedRAMP defined parameters in the FedRAMP baselines and an overall recommended normalized value for each of the defined security control parameters was determined. These normalized parameters were further evaluated during control scoring sessions by representatives from the DHS CISA .govCAR program, the FedRAMP PMO, and the DHS Continuous Diagnostics and Mitigation (CDM) program. The parameters were adjusted during these sessions to establish the most reasonable level of security and to protect against the most prevalent threat actions.

[1] NIST Special Publication (SP) 800-53 Rev. 4 - Security and Privacy Controls for Federal Information Systems and Organizations

The NIST SP 800-53 security controls were scored using DHS CISA .govCAR methodology and rated for their ability to Protect, Detect, and Respond (P/D/R) against a series of threat actions enumerated using a cyber threat framework (i.e. NTCTF or MITRE ATT&CK). Each security control was assigned a value of Limited, Moderate, Significant, or Not Applicable for the functions of Protect, Detect, and Respond for each threat action.

The scoring effort for NIST 800-53, rev. 4 evaluated each of the NIST 800-53 security controls and its associated controls items against the 200+ techniques in the NTCTF. This process was tedious and the resulting scoring effort took almost a year to complete. The scoring effort for NIST 800-53, rev 5 included the evaluation of 686 unique MITRE ATT&CK tactic/technique pairs. To streamline scoring effort and accommodate the additional workload associated with the larger MITRE ATT&CK framework, the scoring team turned its attention toward using existing data to create the control scores.

In December 2020, Engenuity published a mapping of NIST 800-53 controls to ATT&CK, along with the assumptions used in creating the mapping. The scoring team leveraged this mapping to streamline the scoring effort of NIST 800-53, rev 5 controls. However, the Engenuity mapping did not cover all the control families in NIST 800-53, rev 5. To ensure the scoring effort for NIST 800-53 rev 5. was comprehensive, each control was scored using one of the following methods;

1. **Scored Using Engenuity** - these were primary controls covered by the Engenuity analysis and leveraged the mappings to the ATT&CK framework and their associated mitigations.
2. **Control Enhancements** – these control enhancements either inherited the score from the base control (i.e., same) or the enhancement received and updated score with justification.
3. **Correlated Controls** – these controls were not covered by the Engenuity analysis but inherited the mapping from a control that was covered as the control was identified as correlating to the covered control.
4. **Non Engenuity Controls** – these controls were scored using a method that did not rely on the Engenuity analysis (i.e., non-engenuity). **For these items the scoring team isolated specific attack techniques that the associated control provided mitigations against.**

The scoring was performed by Subject Matter Experts (SMEs) from several agencies and organizations. The scoring process utilized a form of the Delphi method (a method commonly used to reach consensus among multiple experts) who exchanged scoring opinions with supporting arguments. A scoring rubric helped normalize the scores based on the characteristics of the capability. The outcome was a collective group consensus for each capability scored with supporting rationale. The results of the control scoring were used to calculate an overall protection value for each control. The higher the value, the greater the level of protection provided by the control (i.e., more threat actions mitigated).

Appendix A: Security Controls Scoring, provides detailed, step-by-step descriptions of the control scoring process.

Phase 2: Security Controls Assessment

To enable automation and integration in the risk scoring process, the NIST SP 800-53 security controls were deconstructed into control items. Control items are the more granular parts of a security control that are

assessed by determination statements defined within NIST SP 800-53A, rev4[2]. Deconstruction of security controls into control items allows them to be grouped into security capabilities with defined, testable defect checks. This approach allows for data from automated and manual assessments to be integrated into the overall risk profiling process. An assessment of the system implementation of a control item results in a status of "Satisfied" or "Other than satisfied." The assigned value for each control item was then used as an input for risk profiling.

Appendix B: Security Controls Assessment provides example details and a step-by-step approach for utilizing the threat-based control scores from the security assessment.

Phase 3: Risk Profiling

The NIST SP 800-53, rev 4, security controls were mapped to the sixteen (16) NIST Interagency Report (NISTIR) 8011 capabilities[3]. The purpose of the NISTIR 8011 is to provide an approach for automating assessments of security controls in systems. This can be utilized for initial assessments, continuous monitoring, and ongoing security authorizations.

To derive a risk profile, the protection values that were assigned to each security control during the threat analysis/scoring phase and the assessment results (i.e., satisfied/other than satisfied) for each control item were leveraged to compute an overall risk maturity level for each security capability, as shown in Figure 1.

Appendix C: Risk Profiling provides details and a step-by-step approach for determining the risk maturity profile of a system.

Security Capability	Maturity Level
Manage and Assess Risk (RISK) (1)	100%
Perform Resilient Systems Engineering (SE) (2)	100%
Software Asset Management (SWAM) (3)	100%
Configuration Settings Management (CSM) (23)	95.65%
Manage Trust for Persons Granted Access (TRUST) (3)	96.86%
Manage Behavioral Expectations (BEHAVE) (4)	75%
Manage Credentials and Authentication (CRED) (20)	94.53%
Manage Privileges and Accounts (PRIV) (22)	95.03%
Manage Network Boundaries (BOUND-N) (3)	100%
Manage Preparation for Events (Incidents and Contingencies) (PREP) (4)	100%

Figure 1: Example Security Capability Maturity Level (Proof of Concept)

[2] NIST Special Publication 800-53r4, Assessing Security and Privacy Controls in Federal Information Systems and Organizations. Building Effective Assessment Plans December 2014

[3] NIST Interagency Report (NISTIR) 8011 Automation Support for Security Control Assessments. Volume 1 June 2017

Applications of Threat-Based Risk Profiling

The threat-based approach to risk profiling enables a wide range of opportunities for the government and industry to utilize the resulting data to make informed risk-based decisions. Additional opportunities for application of this approach are under exploration. Table 1, Potential Additional Opportunities for Application of this Approach, details a high-level set of initial opportunities, along with the potential impact of applying the threat-based model.

Opportunity	Impact
Inform updates to Future Control Baselines	Reduce the burden on CSPs and improve security by focusing on the most prevalent controls.
Incorporate into Annual Assessments and Continuous Monitoring	Enables annual assessments and continuous monitoring activities to focus on prioritized risks that address prevalent threats
Produce Risk Profiles using the Open Security Controls Assessment Language (OSCAL) (e.g., Security Assessment Report (SAR), Plan of Action & Milestones (POA&M), Continuous Diagnostics and Mitigation (CDM) sensors)	Enables adoption of automated, near real-time and current updates of the threat environment to generate a true risk profile for an information system
Assist Authorization decision-making	Provides threat-based data that better informs risk management decisions and authorizations
Prioritize Remediation Efforts	Enables information resource spending and allocation by allowing the government and industry to address the most significant problems first
Identify Desired Future State	Enables strategic planning to assist with roadmapping and cost benefit analyses

Table 1. Potential Additional Opportunities for Application of this Approach

Conclusion

FedRAMP is committed to evaluating ways to continuously drive efficiency and cost-savings by adapting and improving its processes to better service federal cloud cybersecurity needs. Today's authorization approach identifies residual risk based on high, moderate, and low security impact. The FedRAMP PMO anticipates that with the right tools and processes, vendors could enter the federal marketplace faster, using fewer of their own and federal resources, and with more secure systems that protect against the most current threat landscape.

With a threat-based risk profile, agencies, CSPs, and other industry partners can strategically manage and develop the protection of their systems. This threat-based methodology provides an innovative approach to inform risk management decisions across the government. Additionally, this approach provides an opportunity to expedite the authorization process by prioritizing controls that mitigate threats and vulnerabilities posing the most risk to our federal systems and data.

Industry and federal government entities are encouraged to review this approach, methodology, and intended impact and provide feedback to the FedRAMP PMO. All questions and comments regarding the details outlined in this paper should be directed to the FedRAMP PMO via info@fedramp.gov, with the subject line "Feedback: FedRAMP PMO Threat-Based Risk Profiling Approach."

> **Threat-based scoring will allow authorizing officials to leverage qualitative data based on a defensible methodology to inform risk-based authorization decisions.**

Appendix A: Security Controls Scoring

Step 1. Control Item Scoring

To support scoring, each NIST SP 800-53 control was deconstructed into its associated control items. Threat scoring was performed at the control item level so that each control item could be classified as Defined, Document (i.e., passive), or Implement (i.e., active). See Figure 2: Control Item Classification below for more.

Defined	Control item defines the implementation standards for the control
Document	Control item requires the documentation of related security relevant information
Implement	Control item implements a Protect (P), Detect (D), or Respond (R) capability

Figure 2: Control Item Classification

Control items classified as "Implement" can take on one of the following values:	• Significant (S) • Moderate (M) • Limited (L) • Applicable (A) • Not Applicable (NA) • None (N)
Control items classified as "Defined" or "Document" can take on one of the following values:	• Applicable (A) • Not Applicable (NA)
Each control item was scored for its ability to P/D/R to each threat action[4]. To calculate the overall protection value for each control item P/D/R functions were weighted as follows:	• P = .4 • D = .3 • R = .3

Figure 3: Control Item Scores and Weights

The weighting was applied for each control item that received a P/D/R score and the sum of the weighted P/D/R scores was multiplied by the threat action heatmap value to produce a protection value for each control. This process was repeated for every threat action that received a score for the associated control

[4] Representative examples of adversarial threat events expressed as tactics, techniques, and procedures (TTPs)

item and the summation of those individual protection values produced an overall Protection Value (PV) for the control. The formula below depicts the computation of the control PV:

Protection Value (PV) Formula:

$$\textbf{\textit{ProtectionValue (PV)}} = \sum_{TA} TA_{HeatMapValue} * \{.4 * (.9P_{S(0,1)} + .6P_{M(0,1)} + .3P_{L(0,1)} + .1P_{A(0,1)}) + .3 * (.9D_{S(0,1)} + .6D_{M(0,1)} + .3D_{L(0,1)} + .1D_{A(0,1)}) + .3 * (.9R_{S(0,1)} + .6R_{M(0,1)} + .3R_{L(0,1)} + .1R_{A(0,1)})\}$$

Notes: PS (0,1) takes the value 0 if there was no Significant Protect score and the value of 1 if the Protect score is Significant. Similarly, RL (0,1) takes the value 0 if there was no Limited Respond score and the value of 1 if the Respond score is Limited. Threat Action Heat Map Values were provided by DHS CISA .govCAR.

Step 2. Security Control Prioritization

To prioritize the security controls, the security control items were regrouped back into their associated security controls. As the control items were grouped for scoring by those items that were classified as "Implement," or "Defined/Document," it was possible to get two different protection values for each control. The overall protection value for each security control was calculated by summing the distinct protection values for all of the control items associated with that control.

AC-2		
AC-2(a)	63.25	
AC-2(b)	20.63	
AC-2(c)	20.63	
AC-2(d)	63.25	
AC-2(e)	63.25	
AC-2(f)	63.25	83.88
AC-2(g)	63.25	
AC-2(h)	63.25	
AC-2(i)	63.25	
AC-2(j)	63.25	
AC-2(k)	20.63	

Figure 4: Security Controls Protection Value

With the overall protection values for each security control now defined, it is possible to rank the security controls in priority order.

Appendix B: Security Controls Assessment

Ranking of security controls by their protection values enables the organization to establish an assessment threshold shifting the focus of the security controls assessment to evaluating only those security controls which fall above the established threshold. This prioritization has the potential to provide the foundation for streamlining the security authorization process.

The results of the security controls assessment for each control (satisfied or other than satisfied) and the protection value of the control can be leveraged to produce an implementation value for each control which is weighted based upon the protection values of the related control items. The formula below represents the calculation for control implementation value:

$$Control\ Implementation\ Value = [\sum_{PV} PV_{ControlItems} * (\%control\ items\ implemented)]/PV$$

For example, Figure 5 demonstrates the Security Control Implementation Value (AC-2 Example) illustrating how to utilize the assessment results to compute an overall implementation value for a security control.

Account Management (AC-2)		PV	Imp. Status	CI Score	Sub-totals	Imp. Value	%Imp
AC-2(a)	63.25	83.88	1	63.25	47.44	61.19	73%
AC-2(d)	63.25		1	63.25	13.75		
AC-2(e)	63.25		1	63.25			
AC-2(f)	63.25		1	63.25			
AC-2(g)	63.25		1	63.25			
AC-2(h)	63.25		0	0			
AC-2(i)	63.25		0	0			
AC-2(j)	63.25		1	63.25			
AC-2(b)	20.63		1	20.63			
AC-2(c)	20.63		1	20.63			
AC-2(k)	20.63		0	0			

Figure 5: Security Control Implementation Value (AC-2 Example)

Appendix C: Risk Profiling (i.e., Capability Maturity Levels)

To create an overall system threat-based risk profile (capability maturity scores), each of the NIST 800-53 security controls was mapped to the security capabilities listed in NISTIR 8011. It is important to note that a single control can support multiple capabilities. The implementation values for each of the security controls were then used to calculate an overall maturity level for each capability (see Table 2, Capability Maturity Levels.) Figure 6 below provides an example of this process for the security capability: Manage Trust for Persons Granted Access (TRUST).

Manage Trust for Person Granted Access (TRUST)

Control No.	Control Name	% Implemented
AC-2	Account Management	73%
AC-5	Separation of Duties	80%
AC-6	Least Privilege	100%
	Capability Maturity Level:	**84%**

Figure 6: Capability Maturity Level (TRUST)

Once each of the sixteen security capabilities has been assigned a maturity level it now becomes possible to produce an overall Prioritized Risk Profile (maturity level) for the information system.

Security Capability	Maturity Level
Manage and Assess Risk (1)	100%
Perform Resilient Systems Engineering (2)	100%
Software Asset Management (3)	100%
Configuration Settings Management (23)	95.65%
Manage Trust for Persons Granted Access (3)	96.86%
Manage Behavioral Expectations (4)	75%
Manage Credentials and Authentication (20)	94.53%
Manage Privileges and Accounts (22)	95.03%
Manage Network Boundaries (3)	100%
Manage Preparation for Events Incidents and Contingencies (4)	100%
Manage Anomalous Event Detection (17)	99.45%
Manage Anomalous Event Response and Recovery (6)	100%

Table 2: Capability Maturity Levels

Appendix D: Maintenance

DHS CISA .govCAR continues to evaluate the threat landscape and updates the threat framework and heat map accordingly. In addition, as the NIST SP 800-53 security controls catalog is updated, additional scoring initiatives will be required. These changes will be evaluated, and additional scoring sessions will be conducted.

Rev 5: To streamline this process the Open Security Controls Assessment Language (OSCAL) was implemented for the Rev 5 analysis to programmatically compare versions of NIST SP 800-53 (i.e., rev 4 and rev 5) and to automatically identify changes/gaps. Our model supports dynamic updates as new threat data becomes available.