Discussion Essay: Topics for Community Discussion at the NIST Cyber AI Profile Spring 2026 Community of Interest (COI) Working Session #3 – Usability of the Cyber AI Profile

Introduction

The NIST Cyber AI Profile ("Profile") is intended to help organizations strategically adopt AI while addressing and prioritizing cybersecurity risks stemming from its advancements. In December 2025, NIST NCCoE released the Preliminary Draft of the NIST Cyber AI Profile (NIST IR 8596). In January 2026, the NIST NCCoE Cyber AI Profile team hosted a workshop to obtain feedback on the Preliminary Draft and to identify cybersecurity priorities as AI adoption continues to grow. Feedback from the January 2026 workshop and ongoing analysis of the 1,400+ comments received on the draft confirmed support for the development of the Cyber AI Profile and validated that organizations need practical guidance that builds on existing cybersecurity practices while addressing the cybersecurity risks related to AI development and use. Participants emphasized the need for both enterprise risk management and implementation level resources, particularly for organizations managing AI adoption, integration, and use alongside existing cybersecurity operations. Key themes heard from participants during the workshop have been instrumental in both creating the next draft of the Profile and identifying the community's areas of priority as AI continues to evolve.

As a result of this analysis, the team identified key topics for discussion with the COI. The NIST NCCoE is holding a series of virtual COI working sessions to carry out these discussions. This read-ahead document focuses on two topics that reflect feedback received on the Profile and presents proposed ways forward that will be discussed with the community during COI Session 3: Usability of the Profile, which will be held on May 12, 2026, from 1:00–4:00pm EDT. Input from the COI will help validate our proposed directions, and identify any implementation considerations that should inform the next draft of the Profile.

About the COI Working Sessions

The Spring 2026 COI working session series is designed to support targeted, facilitated discussions to further refine the Cyber AI Profile. The working sessions will cover community feedback on adapting cybersecurity practices to AI, strengthening the Cyber AI Profile in key technical areas, and exploring revised Profile delivery formats to enhance usability for different roles in the AI ecosystem. The working sessions are designed to:

- Collect feedback from the COI regarding the proposed approaches in the discussion essay for each session
- Gather insights regarding current practices and challenges
- Identify key characteristics for successful approaches to addressing the discussion topics

During the working sessions, which are held on Zoom, participants will have the opportunity to share feedback by actively engaging in discussion and responding to Slido polls. The Preliminary Draft and this discussion essay should be reviewed prior to the working sessions as they provide critical context, an overview of the topics, questions for discussion, and proposed paths forward.

The May 12 working session will focus on the following:

- Options for if and how to incorporate roles into the Cyber AI Profile
- How to represent AI-specific needs in Supply Chain documentation and Bills of Materials (BOMs)
- Profile delivery formats

Overview of Discussion Topics and Potential Options

The NCCoE team received a significant number of comments, along with substantial suggested changes, related to the roles that the Cyber AI Profile addresses (e.g., developer, deployer, user), including whether and how roles relate to supply chain and AI bills of materials (AIBOMs) considerations, and the format(s) in which to share the Cyber AI Profile content. To better understand the community's feedback and thoughtfully plan potential revisions, the NCCoE team will continue focused discussions on these important topics. These discussions will help clarify key concerns, identify areas of consensus and divergence, and ensure that any updates are practical and aligned with real-world implementation needs. More detailed information on the context, feedback, and specific comments for each topic is provided in the subsections below.

Roles of Profile Audiences

Context

The cybersecurity expertise required to managing AI artifacts/assets differs from the knowledge base and expertise of traditional cybersecurity professionals. AI assets and artifacts have a distinct development and provenance lifecycle that conventional methods do not capture; this distinction matters within the NIST CSF because the same CSF Functions must be applied across all cybersecurity activities. Differentiating roles between developers, deployers, and end users ensures that responsibility for AI model-level risk is explicitly owned, that AI-relevant controls are implemented, and that incident response plans address AI-centric failures. Differentiating these roles therefore enables the CSF to achieve its intended outcomes - risk-aware governance and effective mitigation - while retaining clear lines of ownership and accountability and seamless collaboration with established cybersecurity teams.

Overview of comments

Analysis of comments indicated that the community is interested in better understanding how the Cyber AI Profile applies to the roles of developers, deployers, and users. These comments were also aligned with discussions regarding Supply Chain and AIBOMs. Feedback generally fell into the following primary themes:

1) Handling of roles: Developer vs. Deployer vs. User
   - Requests to include AI specific roles and responsibilities fields within the CSF Considerations Section 2.4.
   - Requests to explicitly recognize that developers, deployers, and other relevant actors across the value chain generally have differentiated responsibilities based on their roles in the context of AI.

2) Supply Chain
   - Feedback stressed that the Cyber AI Profile Considerations must go beyond generic mentions of “AI supply chain risk management” and provide concrete, actionable references; comments asked for richer citations to existing standards, tooling, and best-practice frameworks that specifically address the end-to-end AI supply chain (e.g., SBOM/AIBOM specifications and emerging model-card / data-card initiatives).
   - Feedback also highlighted that the considerations must also make clear that data is as critical an element of the AI supply chain as software: training-data provenance, licensing, quality metrics, and lineage must be captured, signed, and monitored with the same rigor applied to code, containers, and model binaries.

- The Cyber AI Profile draft should emphasize the criticality of the supply chain as a foundational component of any AI-enabled system, strengthening content in Subcategories throughout Section 2.4, as appropriate, where the concept of supply chain is included.
- Feedback indicated interest in illustrative, annotated case studies that walk a reader through a real-world adversarial attack on an AI/ML supply chain. While out of scope for the Cyber AI Profile, this input is useful for future planning.

3) AIBOMs versus AISBOMs

- Feedback indicated opposing desires regarding inclusion of "AIBOM" in the Profile, with some comments indicating it is necessary to aid organizations in auditability and provenance attestations for the supply chain of AI inventories while other comments indicated "AIBOM" should *not* be included in the Profile due to the immaturity of the standards and tooling. Feedback also indicated that removal of "AIBOM" mentions would allow the Profile to remain method agnostic when Profile users discuss implementation approaches.

- Feedback signaled interest in a formal definition of "AIBOM" and distinctions from traditional SBOMs, and emerging AI model cards and data cards.

Proposed options forward on handling roles in the Profile:

Option 1: Differentiate unique aspects of AI roles within an organization in Section 1.3 to highlight the shared responsibilities across the 3 Focus Areas: Provide a shared responsibility matrix that separates ownership of the distinct aspects of AI with a roles-based focus. Use Figure 11 in the Cyber AI Profile as the basis for this matrix to highlight that these new AI roles and responsibilities are shared across the Focus Areas. Use the AI tasks and actors as described in the NIST AI RMF[1] to show roles where AI development, deployment, and end users are jointly responsible for secure, safe, reliable, and trustworthy outcomes.

Option 2: Include specific roles, based on the AI actors and actor tasks discussed in the Cyber AI Profile's Sample Focus Area Considerations as referenceable elements to augment new AI-specific personnel (NIST's preferred option): Provide a representative list of AI roles, based on the AI actions and actors described in the AI RMF (NIST AI 100-1). While specific roles, job titles, and how they are integrated may differ within organizations, the AI RMF provides a representative list of AI tasks and actors across the AI system lifecycle. Based on comments from the Cyber AI Profile COI, the most relevant for roles and

responsibilities emphasizing AI cybersecurity-focused activities include[1]:

- **AI Developers:** provide the initial infrastructure of AI systems, responsible for model building and interpretation tasks, including the creation, selection, calibration, training, and/or testing of models or algorithm (e.g., machine learning experts, data scientists, developers)
- **AI Deployers:** responsible for contextual decisions relating to how the AI system is used to assure deployment of the system into production, including piloting the system, checking compatibility with legacy systems, ensuring regulatory compliance, managing organizational change, and evaluating user experience (e.g., system integrators, software developers, end users, operators and practitioners, evaluators)
- **AI Governance:** retains management, fiduciary, and legal authority and responsibility for the organization in which an AI system is designed, developed, and/or deployed (e.g., organizational management, senior leadership, and the Board of Directors).
- **End Users:** use the system for specific purposes within organizational restrictions, and can range in competency from AI experts to first-time technology end users (e.g., individuals, groups).

Introducing new, unique roles and responsibilities that AI integration necessitates and mapping those roles to each CSF Subcategory, as shown in Figure 1, can provide organizations a way to demonstrate to evaluators, customers, or auditors that *"someone* is responsible for every cybersecurity activity." Organizations could then, for example, create a *CSF-aligned responsibility matrix* (e.g., Responsible, Accountable, Consulted, and Informed [RACI] matrix) that satisfies both internal governance and external audit expectations when adopting AI. However, assigning roles to each Subcategory may introduce unnecessary complexity when implementing the Cyber AI Profile given the high degree of variability between notional roles in the Profile and organization-specific roles. This may also remove some of the inherent flexibility of the Cyber AI Profile by inadvertently implying rigidity in how roles are aligned to outcomes.

[1] For discussion purposes, some AI Tasks in NIST AI 100-1 v1.0 have been extrapolated into roles (e.g., AI Development tasks are represented in role form as AI Developers).

CSF 2.0 Core: GOVERN	General Considerations	Focus Area Proposed Priorities & Considerations Secure	Defend	Thwart
GOVERN (GV)	The organization's cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored			
Organizational Context (GV.OC)	The circumstances—mission, stakeholder expectations, dependencies, and legal, regulatory, and contractual requirements—surrounding the organization's cybersecurity risk management decisions are understood			
GV.OC-01: The organizational mission is understood and informs cybersecurity risk management	**General Considerations:** No general considerations identified—see Focus Area Considerations. **Example Informative References:** NIST SP 800-53, Rev 5: PM-11 **Roles:** AI Governance, End Users, ... SAMPLE TEXT PROVIDED FOR DISCUSSION PURPOSES ONLY	**Proposed Priority: 3** **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain **Roles:** AI Developers, AI Deployers	**Proposed Priority: 3** **Sample Opportunities:** Standard cybersecurity practices apply. **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** ENISA Threat Landscape 2025; DASF 50; ATLAS AML.M0020; OWASP AI Exchange: AI Security Overview; https://arxiv.org/pdf/2311.05232; NIST AI 100-2e2025 **Roles:** AI Developers, AI Deployers	**Proposed Priority: 3** **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** AI-specific Example Informative References pending additional inputs. **Roles:** End-Users

Figure 1: The Profile as it currently appears in the draft with the inclusion of "Roles" as defined in Option 1.

Proposed options forward on handling Supply Chain in the Profile:

Option 1: Create a separate section dedicated to AI Supply Chain to address relevant use cases (NIST's preferred option): This option would add a distinct section (for example, a new section in Section 1.2 or as an Appendix item) to emphasize the criticality of the supply chain of AI assets (AI models, AI model providers, training data, deployment and hosting hardware, supporting software, etc.). This new section could describe the cybersecurity considerations associated with the supply chain of AI assets and inventories, highlighting the unique role that model and data provenance plays in this context. This approach would allow the Profile to acknowledge the importance of the AI supply chain security and assurance as organizations rapidly adopt it.

Option 2: Elevate the visibility of supply chain in CSF Considerations: Emphasize the criticality of the supply chain as a foundational component of any AI-enabled system, strengthening content in Subcategories, as appropriate, where the concept of supply chain is included.

Proposed options forward on handling AI Bill of Materials (AIBOM) in the Profile:

Option 1: Do not reference it ("AIBOM") within the Cyber AI Profile and instead leverage known Software Bill of Materials, AI Model Cards, and Data Cards within the considerations: To avoid fragmentation, the Cyber AI Profile should not require the explicit use of an "AIBOM" label. Instead, the guidance can acknowledge the "AIBOM" as one possible method for documenting AI provenance

alongside other well-established instruments such as model cards, data cards, and the conventional software- and hardware-BOMs. NIST could clarify that the "AIBOM" is simply an additional option for organizations seeking greater accountability, while also recognizing that many practitioners already rely on AI model card and/or data card formats to convey the same information. Moreover, NCCoE should note that the implementation of SBOMs for AI is inherently more complex—AI artifacts involve large binary model files, evolving training data, and nuanced licensing constraints—so the standards and tooling ecosystem for "AIBOMs" are still maturing relative to the long-standing SBOM practices for conventional software.

Option 2: Standardize terminology within the context of existing Software Bill of Materials (BOM), Model Cards, and Data Cards guidance (NIST's preferred option): Regardless of the terminology, e.g., AISBOM (AI Software BOM) or AIBOM (AIBOM), the Cyber AI Profile should explicitly state that an "AIBOM" is not intended to replace the existing Software BOM framework, but to augment it when a product contains AI-specific artifacts such as trained models, curated datasets, or provenance metadata. By positioning the newly defined BOM as a supplemental layer, organizations can continue to use the mature SBOM, model card and data card tooling and processes that already support software and hardware components while gaining the additional transparency required for AI elements.

(Continued) Option 2: Accept/integrate "AIBOM" and further develop a concise, minimum standardized set of elements that can be adopted as best-practice guidance and ensure they are referenced throughout the Cyber AI Profile as an Appendix item (NIST's preferred option):

With a common terminology in place, NIST and NCCoE can collaborate with the broader AI-security community to publish a minimal checklist of fields that any "AIBOM" should contain. The core elements might include: (i) model version and provider information, (ii) model architecture, training-data sources, hyper-parameters, and known limitations, (iii) cryptographic signing and verification of model artifacts, (iv) model fingerprinting or watermarking for provenance tracking, (vi) runtime integrity verification for deployed models, (vii) fine-tuning datasets together with provenance and licensing data, (viii) embedding-model and vector-database configurations, (ix) system prompts and prompt-template specifications, (x) retrieval corpora for Retrieval-Augmented Generation (RAG) pipelines, (xi) third-party AI services and APIs, (xii) evaluation datasets and benchmark results, (xiii) runtime

telemetry or attestation data that enable auditors to confirm the deployed model's identity (e.g., a hash of the weights), (xiv) container or inference-runtime image hash, and (xv) a description of the execution environment. By publishing these items as a "minimum AIBOM schema," NIST would give developers a clear, interoperable target while still allowing flexibility for organizations that wish to adopt richer, domain-specific extensions. This approach balances the need for transparency with the reality that "AIBOM" standards are still evolving, and it provides a concrete stepping-stone toward broader, industry-wide adoption. However, this effort may somewhat distract from the role/objective of the Cyber AI Profile.

Open questions for discussion during the working sessions:

- Additional questions about roles within the Cyber AI Profile:
  - Are these newly defined, AI-specific, roles or just existing roles (developers, deployers, users) that require additional expertise, such as AI knowledge, etc.?
  - Is the adoption of NIST's AI RMF (NIST AI 100-1) definitions for AI Developers, AI Deployers, AI Governance, and End-Users sufficient?
- How should accountability be spread across the supply chain?
  - How are organizations today establishing the lines of responsibility among developers, deployers, and users in the context of AI systems?
  - Are there other parties which share responsibility?
- What should be included in an AIBOM that would not be in an SBOM?

Profile Delivery Formats

Context

The current iteration of the Cyber AI Profile delivers all of its content via a portable document format (PDF) document, which is a static format. However, the bulk of the content of the Profile exists in the form of tables. In these tables, the Subcategories of the CSF 2.0 are related to three Focus Areas, as well as General Considerations that apply across Focus Areas. Each one of these cells can contain multiple types of information that relate to a CSF Subcategory in the context of the Focus Areas (including General Considerations). The Preliminary Draft contains a numerical priority level, prose which describes how AI impacts a specific Subcategory and Focus Area, and a list of Informative References which helped inform the given considerations and priority level of a cell. The Defend Focus Area includes an additional type of information that describes opportunities for AI to be deployed to defend organizations.

Overview of comments

Feedback has demonstrated interest other preferred delivery formats for the Profile beyond the initial static document format in which the Preliminary Draft was provided. Alternative formats other than PDF were raised in order to support customization of the Profile, as well as integration into tooling. Preferences emerged for the following characteristics of the delivery format(s):

- Machine-readable (various specific formats were proposed ranging from an Excel workbook to CSV, JSON, or XML,
- Provides capabilities for readers to customize their view of the information (e.g., sorting, filtering, hiding),
- Takes advantage of NIST mapping tools, including the Online Informative References (OLIR) Program and the Cybersecurity and Privacy Reference Tool (CPRT),
- Ingestible by LLMs, and
- Provides the ability to be versioned.

Proposed way forward on Profile delivery formats

NIST is proposing to address this feedback by releasing an Excel Workbook alongside the PDF publication. These two formats will be released together **when the final version of the Profile is published**. NIST is proposing this path forward to address community feedback because:

- Excel is exportable to CSV. From a CSV file, users can convert to the machine-readable format they prefer.
    - Nearly all of these formats are then easily ingested by LLMs.
- Excel has many functions for sorting and filtering, as well as more advanced operations.
    - Using filtering and sorting, readers can choose which Focus Areas or CSF 2.0 Core elements (i.e., Functions, Categories, Subcategories) they want to view, if that better fits their use case.
    - Alternatively, for those who want a complete picture of AI's impact on Cybersecurity, the default view can be comprehensive.
- Producing an Excel format for the Profile will enable integration into OLIR and CPRT.
- The workbook is versionable.

However, there remains a question of how to format this workbook. NIST proposes five options:

1. **Option 1: Keep the exact same layout as the tables in the PDF version of the Profile.**

CSF 2.0 Core: GOVERN	General Considerations	Focus Area Proposed Priorities & Considerations		
		Secure	Defend	Thwart
GOVERN (GV)	The organization's cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored			
Organizational Context (GV.OC)	The circumstances—mission, stakeholder expectations, dependencies, and legal, regulatory, and contractual requirements—surrounding the organization's cybersecurity risk management decisions are understood			
GV.OC-01: The organizational mission is understood and informs cybersecurity risk management	**General Considerations**: No general considerations identified—see Focus Area Considerations. **Example Informative References:** NIST SP 800-53, Rev 5: PM-11	**Proposed Priority:** 3 **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain	**Proposed Priority:** 3 **Sample Opportunities:** Standard cybersecurity practices apply. **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** ENISA Threat Landscape 2025; DASF 50; ATLAS AML.M0020; OWASP AI Exchange: AI Security Overview https://arxiv.org/pdf/2311.05232; NIST AI 100-2e2025	**Proposed Priority:** 3 **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** AI-specific Example Informative References pending additional inputs.
GV.OC-02: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	**General Considerations**: AI use introduces considerations from multiple aspects of organizational operations, including legal, technical, procurement/ acquisitions, and governance teams. Collaboration across these areas is essential for addressing AI-related cybersecurity risks. Multidisciplinary approaches	**Proposed Priority:** 3 **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** DASF 13, 40, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025	**Proposed Priority:** 2 **Sample Focus Area Considerations:** Understanding the strengths and limitations of AI capabilities for cyber defense is important for meeting stakeholder expectations and to ensure the balance between the required human oversight and automation.	**Proposed Priority:** 2 **Sample Focus Area Considerations:** Standard cybersecurity practices apply. **Example Informative References:** AI-specific Example Informative References pending additional inputs.

Figure 1: The Profile as it currently appears in the draft

	A	B	C	D	E
1	CSF 2.0 Subcategory	General	Secure	Defend	Thwart
2	**GV.OC-01**: The organizational mission is understood and informs cybersecurity risk management	**Considerations**: No general considerations identified—see Focus Area Considerations. **Informative References**: NIST SP 800-53, Rev 5: PM-11	**Priority**: 3 **Focus Area Considerations**: Standard cybersecurity practices apply. **Informative References**: OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain	**Priority**: 3 **Opportunities**: Standard cybersecurity practices apply. **Focus Area Considerations**: Standard cybersecurity practices apply. **Informative References**: ENISA Threat Landscape 2025; DASF 50; ATLAS AML.M0020; OWASP AI Exchange: AI Security Overview; https://arxiv.org/pdf/2311.05232; NIST AI 100-2e2025	**Priority**: 3 **Focus Area Considerations**: Standard cybersecurity practices apply. **Informative References**: AI-specific Example Informative References pending additional inputs.
3	**GV.OC-02**: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	**Considerations**: AI use introduces considerations from multiple aspects of organizational operations, including legal, technical, procurement/acquisitions, and governance teams. Collaboration across these areas is essential for addressing AI-related cybersecurity risks. Multidisciplinary approaches facilitate a comprehensive enterprise view.	**Priority**: 3 **Considerations**: Standard cybersecurity practices apply. **Informative References**: DASF 13, 40, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025	**Priority**: 2 **Opportunities**: Standard cybersecurity practices apply. **Considerations**: Understanding the strengths and limitations of AI capabilities for cyber defense is important for meeting stakeholder expectations and to ensure the balance between the required human oversight and automation.	**Priority**: 2 **Focus Area Considerations**: Standard cybersecurity practices apply. **Informative References**: AI-specific Example Informative References pending additional inputs.

Sheet1

Figure 2: Option 1 rendered in Excel

2. **Option 2: Split the Focus Areas across different sheets within Excel and split the priority, considerations, and Informative References across columns.** The Subcategories are split across rows.

	A	B	C	D
1	CSF 2.0 Subcategory	Secure Priority	Secure Considerations	Secure Informative References
2	**GV.OC-01**: The organizational mission is understood and informs cybersecurity risk management	3	Standard cybersecurity practices apply.	OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain
3	**GV.OC-02**: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	3	Standard cybersecurity practices apply.	DASF 13, 40, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025
	GV.OC-03: Legal, regulatory, and contractual requirements regarding cybersecurity—including privacy	3	The legal framework regarding AI usage is evolving, particularly in areas like cybersecurity, privacy, fair use of	DASF 32, 40; OWASP AI Exchange: Controls to Limit the Effects of Unwanted Behavior; OWASP Conventional runtime controls: OWASP AI

General Secure Defend Thwart +

Figure 3: Option 2 rendered in Excel. The image depicts just the sheet for Secure, but similar sheets will exist for General, Defend, and Thwart.

	A	B	C	D	E
1	CSF 2.0 Subcategory	Defend Priority	Defend Opportunities	Defend Considerations	Defend Informative References
2	**GV.OC-03**: Legal, regulatory, and contractual requirements regarding cybersecurity—including privacy and civil liberties obligations—are understood and managed	1	AI capabilities can support compliance with legal, regulatory, and contractual requirements by analyzing and summarizing requirements, accelerating policy development, speeding up review processes, identifying and mapping similar concepts between documents, and even	Defensive AI tools handle logs and sensitive data in line with privacy obligations including consent, usage and aggregation controls and new AI specific laws. AI audits are designed to demonstrate	DASF 44,50; ENISA Threat Landscape 2025 ; ATLAS AML.M0005
3	**GV.OC-02**: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	2	Standard cybersecurity practices apply.	Understanding the strengths and limitations of AI capabilities for cyber defense is important for meeting stakeholder expectations and to ensure the balance between the required human oversight and automation.	DASF 38,50;OWASP AI Exchange: AI Transparency; ENISA Threat Landscape 2025; ATLAS AML.M0003
4	**GV.OC-01**: The organizational mission is understood and informs cybersecurity risk management	3	Standard cybersecurity practices apply.	Standard cybersecurity practices apply.	ENISA Threat Landscape 2025; DASF 50; ATLAS AML.M0020; OWASP AI Exchange: AI Security Overview; https://arxiv.org/pdf/2311.05232; NIST AI 100-2e2025
5					

General Secure Defend Thwart +

Figure4: Option 2 Defend sheet displaying the priority sorting capability.

3. **Option 3: Keep the Focus Areas in the same sheet within Excel and split the priority, considerations, and Informative References across columns.** The Subcategories are split across rows.

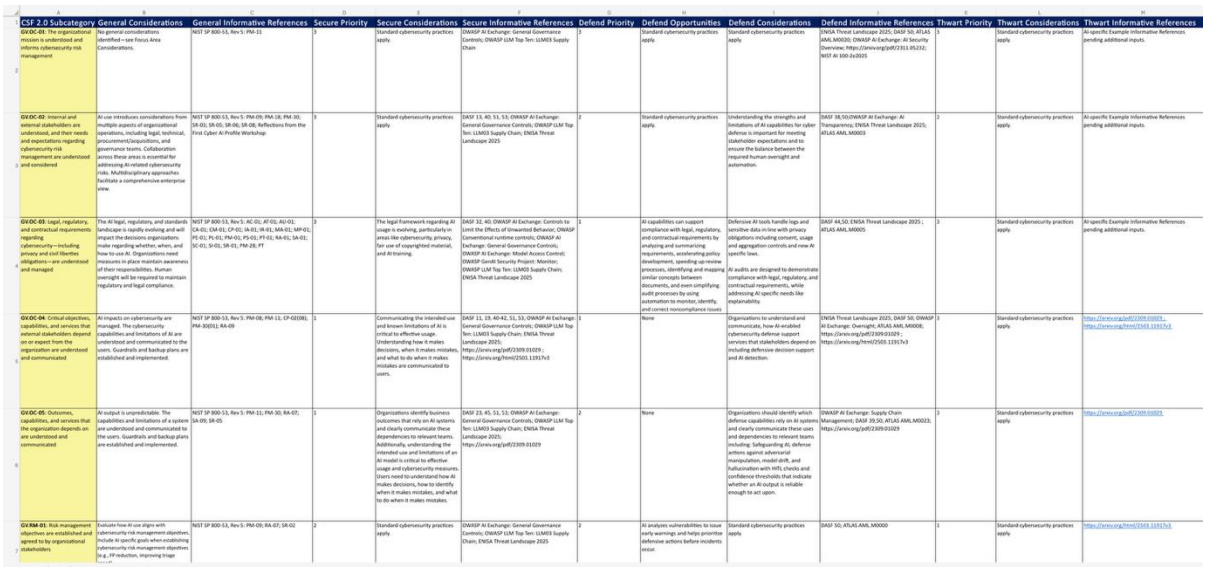

CSF 2.0 Subcategory	General Considerations	General Informative References	Secure Priority	Secure Considerations	Secure Informative References	Defend Priority	Defend Opportunities	Defend Considerations	Defend Informative References	Thwart Priority	Thwart Considerations	Thwart Informative References
GV.OC-01: The organizational mission is understood and informs cybersecurity risk management	No general considerations identified—see Focus Area Considerations.	NIST SP 800-53, Rev 5: PM-11	3	Standard cybersecurity practices apply.	OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain	3	Standard cybersecurity practices apply.	Standard cybersecurity practices apply.	ENISA Threat Landscape 2025; DASF 50; ATLAS AML.M0020; OWASP AI Exchange: AI Security Overview: https://arxiv.org/pdf/2311.05232; NIST AI 100-2e2025	3	Standard cybersecurity practices apply.	AI-specific Example Informative References pending additional inputs.
GV.OC-02: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	AI use introduces considerations from multiple aspects of organizational operations, including legal, technical, procurement/acquisitions, and governance teams. Collaboration across these areas is essential for addressing AI-related cybersecurity risks. Multidisciplinary approaches facilitate a comprehensive enterprise view.	NIST SP 800-53, Rev 5: PM-09; PM-18; PM-30; SR-03; SR-05; SR-06; SR-08; Reflections from the First Cyber AI Profile Workshop	3	Standard cybersecurity practices apply.	DASF 13, 40, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025	2	Standard cybersecurity practices apply.	Understanding the strengths and limitations of AI capabilities for cyber defense is important for meeting stakeholder expectations and to ensure the balance between the required human oversight and automation.	DASF 38,50;OWASP AI Exchange: AI Transparency; ENISA Threat Landscape 2025; ATLAS AML.M0003	2	Standard cybersecurity practices apply.	AI-specific Example Informative References pending additional inputs.
GV.OC-03: Legal, regulatory, and contractual requirements regarding cybersecurity—including privacy and civil liberties obligations—are understood and managed	The AI legal, regulatory, and standards landscape is rapidly evolving and will impact the decisions organizations make regarding whether, when, and how to use AI. Organizations need measures in place maintain awareness of their responsibilities. Human oversight will be required to maintain regulatory and legal compliance.	NIST SP 800-53, Rev 5: AC-01; AT-01; AU-01; CA-01; CM-01; CP-01; IA-01; IR-01; MA-01; MP-01; PE-01; PL-01; PM-01; PS-01; PT-01; RA-01; SA-01; SC-01; SI-01; SR-01; PM-28; PT	3	The legal framework regarding AI usage is evolving, particularly in areas like cybersecurity, privacy, fair use of copyrighted material, and AI training.	DASF 32, 40; OWASP AI Exchange: Controls to Limit the Effects of Unwanted Behavior; OWASP Conventional runtime controls; OWASP AI Exchange: General Governance Controls; OWASP AI Exchange: Model Access Control; OWASP Gen AI Security Project: Monitor; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025	1	AI capabilities can support compliance with legal, regulatory, and contractual requirements by analyzing and summarizing requirements, accelerating policy development, speeding up review processes, identifying and mapping similar concepts between documents, and even simplifying audit processes by using automation to monitor, identify, and correct noncompliance issues	Defensive AI tools handle logs and sensitive data in line with privacy obligations including consent, usage and aggregation controls and new AI specific laws. AI audits are designed to demonstrate compliance with legal, regulatory, and contractual requirements, while addressing AI specific needs like explainability.	DASF 44,50; ENISA Threat Landscape 2025 ; ATLAS AML.M0005	3	Standard cybersecurity practices apply.	AI-specific Example Informative References pending additional inputs.
GV.OC-04: Critical objectives, capabilities, and services that external stakeholders depend on or expect from the organization are understood and communicated	AI impacts on cybersecurity are managed. The cybersecurity capabilities and limitations of AI are understood and communicated to the users. Guardrails and backup plans are established and implemented.	NIST SP 800-53, Rev 5: PM-08; PM-11; CP-02(08); PM-30(01); RA-09	1	Communicating the intended use and known limitations of AI is critical to effective usage. Understanding how it makes decisions, when it makes mistakes, and what to do when it makes mistakes are communicated to users.	DASF 11, 19, 40-42, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025; https://arxiv.org/pdf/2309.01029 ; https://arxiv.org/html/2503.11917v3	3	None	Organizations to understand and communicate, how AI-enabled cybersecurity defense support services that stakeholders depend on including defensive decision support and AI detection.	ENISA Threat Landscape 2025; DASF 50; OWASP AI Exchange: Oversight; ATLAS AML.M0008; https://arxiv.org/pdf/2309.01029 ; https://arxiv.org/html/2503.11917v3	3	Standard cybersecurity practices apply.	https://arxiv.org/pdf/2309.01029 , https://arxiv.org/html/2503.11917v3
GV.OC-05: Outcomes, capabilities, and services that the organization depends on are understood and communicated	AI output is unpredictable. The capabilities and limitations of a system are understood and communicated to the users. Guardrails and backup plans are established and implemented.	NIST SP 800-53, Rev 5: PM-11; PM-30; RA-07; SA-09; SR-05	1	Organizations identify business outcomes that rely on AI systems and clearly communicate these dependencies to relevant teams. Additionally, understanding the intended use and limitations of an AI model is critical to effective usage and cybersecurity measures. Users need to understand how AI makes decisions, how to identify when it makes mistakes, and what to do when it makes mistakes.	DASF 23, 45, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025; https://arxiv.org/pdf/2309.01029	2	None	Organizations should identify which defense capabilities rely on AI systems and clearly communicate these uses and dependencies to relevant teams including: Safeguarding AI, defense actions against adversarial manipulation, model drift, and hallucination with HITL checks and confidence thresholds that indicate whether an AI output is reliable enough to act upon.	OWASP AI Exchange: Supply Chain Management; DASF 39,50; ATLAS AML.M0023; https://arxiv.org/pdf/2309.01029	3	Standard cybersecurity practices apply.	https://arxiv.org/pdf/2309.01029
GV.RM-01: Risk management objectives are established and agreed to by organizational stakeholders	Evaluate how AI use aligns with cybersecurity risk management objectives. Include AI specific goals when establishing cybersecurity risk management objectives (e.g., FP reduction, improving triage	NIST SP 800-53, Rev 5: PM-09; RA-07; SA-02	2	Standard cybersecurity practices apply.	OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025	2	AI analyzes vulnerabilities to issue early warnings and helps prioritize defensive actions before incidents occur.	Standard cybersecurity practices apply.	DASF 50; ATLAS AML.M0002	1	Standard cybersecurity practices apply.	https://arxiv.org/html/2503.11917v3

Figure 5: Option 3 rendered in Excel

4. **Option 4: A combination of option 2 and 3.** Create one sheet which contains a consolidated view of all Focus Areas as well as a sheet for each individual Focus Area. **(NIST's preferred option)**

CSF 2.0 Subcategory	General Considerations	General Informative References	Secure Priority	Secure Considerations	Secure Informative References
GV.OC-01: The organizational mission is understood and informs cybersecurity risk management	No general considerations identified—see Focus Area Considerations.	NIST SP 800-53, Rev 5: PM-11	3	Standard cybersecurity practices apply.	OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain
GV.OC-02: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	AI use introduces considerations from multiple aspects of organizational operations, including legal, technical, procurement/acquisitions, and governance teams. Collaboration across these areas is essential for addressing AI-related cybersecurity risks. Multidisciplinary approaches facilitate a comprehensive enterprise view.	NIST SP 800-53, Rev 5: PM-09; PM-18; PM-30; SR-03; SR-05; SR-06; SR-08; Reflections from the First Cyber AI Profile Workshop	3	Standard cybersecurity practices apply.	DASF 13, 40, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025
GV.OC-03: Legal, regulatory, and contractual requirements regarding cybersecurity—including privacy and civil liberties	The AI legal, regulatory, and standards landscape is rapidly evolving and will impact the decisions organizations make regarding whether, when, and how to use AI. Organizations need	NIST SP 800-53, Rev 5: AC-01; AT-01; AU-01; CA-01; CM-01; CP-01; IA-01; IR-01; MA-01; MP-01; PE-01; PL-01; PM-01; PS-01; PT-01; RA-01; SA-01; SC-01; SI-01; SR-01; PM-28; PT	3	The legal framework regarding AI usage is evolving, particularly in areas like cybersecurity, privacy, fair use of copyrighted material, and AI training.	DASF 32, 40; OWASP AI Exchange: Controls to Limit the Effects of Unwanted Behavior; OWASP Conventional runtime controls; OWASP AI Exchange: General Governance Controls; OWASP AI Exchange: Model Access Control; OWASP

Consolidated | General | Secure | Defend | Thwart | +

Figure 6: Option 4 consolidated view rendered in Excel

	A	B	C	D
1	CSF 2.0 Subcategory	Secure Priority	Secure Considerations	Secure Informative References
2	**GV.OC-01**: The organizational mission is understood and informs cybersecurity risk management	3	Standard cybersecurity practices apply.	OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain
3	**GV.OC-02**: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	3	Standard cybersecurity practices apply.	DASF 13, 40, 51, 53; OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain; ENISA Threat Landscape 2025
4	**GV.OC-03**: Legal, regulatory, and contractual requirements regarding cybersecurity—including privacy and civil liberties obligations—are understood and managed	3	The legal framework regarding AI usage is evolving, particularly in areas like cybersecurity, privacy, fair use of copyrighted material, and AI training.	DASF 32, 40; OWASP AI Exchange: Controls to Limit the Effects of Unwanted Behavior; OWASP Conventional runtime controls; OWASP AI Exchange: General Governance Controls; OWASP AI Exchange: Model Access Control; OWASP GenAI Security Project:

Consolidated | General | Secure | Defend | Thwart | +

Figure 7: Option 4 Focus Area view rendered in Excel

5. **Option 5: Keep the Focus Areas in the same sheet within Excel and split the priority, considerations, and Informative References for each Subcategory across rows.**

	A	B	C	D	E
1	CSF 2.0 Subcategory	General	Secure	Defend	Thwart
2	**(Priority) GV.OC-01**: The organizational mission is understood and informs cybersecurity risk management	None	3	3	3
3	**(Opportunities & Considerations) GV.OC-01**	No general considerations identified—see Focus Area Considerations.	Standard cybersecurity practices apply.	**Opportunities**: Standard cybersecurity practices apply. **Focus Area Considerations**: Standard cybersecurity practices apply.	Standard cybersecurity practices apply.
4	**(Informative References) GV.OC-01**	NIST SP 800-53, Rev 5: PM-11	OWASP AI Exchange: General Governance Controls; OWASP LLM Top Ten: LLM03 Supply Chain	ENISA Threat Landscape 2025; DASF 50; ATLAS AML.M0020; OWASP AI Exchange: AI Security Overview; https://arxiv.org/pdf/2311.05232; NIST AI 100-2e2025	AI-specific Example Informative References pending additional inputs.
5	**(Priority) GV.OC-02**: Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered	None	3	2	2
6	**(Opportunities & Considerations) GV.OC-02**	AI use introduces considerations from multiple aspects of organizational operations, including legal, technical, procurement/acquisitions, and governance teams. Collaboration across these areas is	Standard cybersecurity practices apply.	**Opportunities**: Standard cybersecurity practices apply. **Considerations**: Understanding the strengths and limitations of AI capabilities	Standard cybersecurity practices apply.

Figure 8: Option 5 rendered in Excel

Potential open questions for discussion during working session

- Which format is preferred?

- What else would you want to see in a Profile delivery format? Any additional functionalities?
- Is there another format which would work better? Other Excel Workbook formats could be emailed for reference.

Proposed ways forward on Profile delivery formats

1. **Option 1: Keep the exact same layout as the tables in the PDF version of the Profile**
    a. Pros: Users are familiar with the format from the Preliminary Draft. All information is provided together.
    b. Cons: Keeping multiple types of data (priority, considerations, opportunities, and Informative References) in the same cell decreases the flexibility of the format and is difficult to compare the same types of information across Focus Areas. Filtering / sorting on any one of these fields would require advanced processing.
2. **Option 2: Split the Focus Areas across different sheets within Excel and split the priority, considerations, and Informative References across columns.**
    a. Pros: Having the data split across multiple tabs increases the filtering and sorting capabilities of the format. This format increases the separation between Focus Areas for community members interested in just one Focus Area.
    b. Cons: Splitting the Focus Areas across different tabs within Excel decreases the exportability to machine readable formats (i.e., csv, json, etc.) though it still can be done with post-processing. For people who see a benefit in seeing the Focus Areas side by side, there is no easy way to combine them.
3. **Option 3: Keep the Focus Areas in the same sheet within Excel and split the priority, considerations, and Informative References across columns.**
    a. Pros: All data is sortable and filterable. If community members are only focused on one Focus Area, or even just one section of one Focus Area, the workbook can be filtered down to this (e.g., hide rows or columns). Similarly, if community members want to view multiple Focus Areas at once, they can set filters up so that this data is side by side. Additionally, exporting this format to CSV and other machine-readable formats is simple.
    b. Cons: The number of columns will likely require most community members to scroll to see all of the data for a given Subcategory.
4. **Option 4: A combination of option 2 and 3. (NIST's preferred option)**

    a. Pros: Format 4 retains all of the pros of Format 2 and 3. Additionally, this format allows the community to have both a consolidated view across all Focus Areas as well as single Focus Area views.
    b. Cons: For the consolidated view, users will need to scroll across the screen to see all of the data, but it can be narrowed by switching to a Focus Area sheet or utilizing filtering / hiding columns.
5. **Option 5: Keep the Focus Areas in the same sheet within Excel and split the priority, considerations, and Informative References for each Subcategory across rows.**
    a. Pros: All data is sortable and filterable, though the approaches to doing so will differ from other options and may be slightly more complex (e.g., keyword filtering for "Priority: 1" instead of a dropdown filter to select "1"). Community members can see all of the data pertaining to a given Subcategory from a single view without scrolling, switching to a different sheet, or zooming out.
    b. Cons: This design is more difficult to filter and export to CSV and other machine-readable formats when compared to Format 3.

Next Steps

We welcome feedback on these topics during the upcoming COI working sessions and encourage you to register using the links in the table below. In addition to this discussion essay, we also encourage participants to review the Cyber AI Profile Preliminary Draft prior to the working sessions to best prepare for the informed, targeted discussions. If you are unable to attend or want to provide written feedback, you are welcome to share it via email with cyberaiprofile@nist.gov by May 15, 2026.

Session	Description	Date/Time
Session #1: Updates to Profile Elements and Contents	This session discussed approaches to addressing feedback received regarding Profile content, including clarifying approaches and phrasing for Profile elements such as the priorities and considerations in Section 2 of the Preliminary Draft. Please also see the Session #1 discussion essay.	April 28, 2026 / 1:00-4:00 P.M. EDT
Session #2: Extending the Technical Content	This session explored how the Profile is being strengthened in critical technical areas including Agentic AI and Zero Trust. Please also see the Session #2 discussion essay.	May 5, 2026 / 1:00-4:00 P.M. EDT
Session #3: Usability of the Profile	This session will explore different delivery formats to ensure the Profile meets the needs of different stakeholders	May 12, 2026 / 1:00-4:00 P.M. EDT

These discussions are intended to inform NIST's efforts to craft the next Cyber AI Profile draft. Please stay tuned for the release of the Initial Public Draft. We look forward to hearing from you!