Discussion Essay: Topics for Community Discussion at the NIST Cyber AI Profile Spring 2026 Community of Interest (COI) Working Session #1 – Updates to Profile Elements and Contents

Introduction

The NIST Cyber AI 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 1: Updates to Profile Elements and Contents, which will be held on April 28, 2026, from 1:00–4:00pm EDT. Input from the COI will help validate our proposed direction 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 overall outcomes of these discussions are to:

- Collect feedback from the COI regarding the proposals 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 using 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 April 28 working session will focus on the following:

- Options for Subcategory prioritization schema
- Use and meaning of the phrase “standard cybersecurity practices apply” in the Subcategory considerations.

Overview of Session 1 Discussion Topics and Potential Options

Updates to Profile Elements and Contents

Subcategory Prioritization

Context

Community Profiles tailor the CSF to focus on cybersecurity outcomes in a specific context. Prioritization can help organizations determine which cybersecurity outcomes to address first given limited resources. However, this presents significant challenges and

additional complexities as prioritization can differ depending on how the AI is used and the context and maturity of an organization.

The Cyber AI Profile provides structured and technology-neutral recommendations by CSF Subcategory and Profile Focus Area for a broadly available yet consistent and common approach for strategic adoption of AI as part of a cybersecurity risk management program. Within the context of the Cyber AI Profile and its Focus Areas (Secure, Defend, Thwart), the Profile intends to use a proposed priority for indicating how important it may be for organizations to achieve the Subcategory's outcome(s) to meet the objective of each Focus Area, and to help organizations determine which Subcategories they may wish to address sooner.

The Cyber AI Profile Preliminary Draft proposed a 3-level prioritization schema that is number based, with each number representing a qualitative priority (i.e., "1" for High Priority, "2" for Moderate Priority, and "3" for Foundational Priority). Organizations may decide that they need to adjust (e.g., increase, decrease) a Subcategory's proposed priority to reflect characteristics of the environment, needs, risk tolerance, or other factors, such as available resources, the degree of AI cybersecurity maturity, status of a cybersecurity program, sector regulations, or prevalent or evolving threats and vulnerabilities.

Overview of comments

Feedback indicates that the current proposed prioritization schema used in the draft is unclear and needs additional clarification to be used in practice. Feedback fell into the following primary themes: requests to change the prioritization terms, requests to add a new prioritization level, requests to refine the definitions of the terms used, clarify what it means when the same Subcategory is prioritized differently among the three Focus Areas, and more clarity around definitions and the methodology used to determine the prioritizations.

NIST's proposed options forward on the prioritization schema structure based on analysis of feedback received:

1. Retain the existing 3-level prioritization schema and refine levels and definitions for each level and consistently apply the rubric across all Focus Areas and Subcategories.
2. Replace the current prioritization schema with a 4-level prioritization schema and consistently apply the rubric across all Focus Areas and Subcategories.

- A 4-level prioritization schema would introduce a new level for cybersecurity activities that should be addressed first (e.g., similar in concept to pre-requisites)

3. Replace prioritization schema with an indicator that reflects the approximate relative degree of differences between "traditional" cybersecurity considerations and AI-specific cybersecurity considerations (note: this would be a minimal, significant, substantial scale).
4. Do not include a prioritization schema and therefore no proposed Subcategory priorities. The Profile would retain rationales, considerations, opportunities, and Informative References to support informed risk management discussions and decisions by Profile implementers. **(NIST's preferred option)**

Potential open questions for discussion during the working session:

- Are there specific benefits or challenges that having a prioritization schema introduces? If so, what are those?
- Are there specific benefits or challenges introduced by not including priorities? If so, what are those?

Based on community feedback thus far, NIST has identified removal of prioritization as the most viable option (Option 4). This approach reflects the following considerations:

- Enables broad application of the Cyber AI Profile to multiple sectors and organizations of varying maturity who may be using AI in different ways or guided by different regulations. The Cyber AI Profile will have a very broad user base, and it should be structured in a way that allows for and enables the most flexible use. This would acknowledge that priorities may vary by sector (healthcare vs. retail vs. critical infrastructure) while not trying to exhaustively address all sector-specific considerations.
- Supports use of the Cyber AI Profile with other CSF Profiles without concern of conflicting prioritization schemas or Subcategory priorities. If an organization wants to use the Cyber AI Profile in combination with a different CSF Profile and if each Profile has a different prioritization schema, this can create implementation challenges with organizations having to determine a prioritization schema resolution simply to use both Profiles. Additionally, organizations may resolve this dilemma in different ways, which could contribute to different organizations in the same sector or community losing the common taxonomy and risk perspective

benefits that CSF Profiles are intended to provide. While this challenge is not unique to AI, AI has a wide breadth of impact across and within organizations and communities. AI's speed of adoption across different domains and the rapidly growing landscape of guidelines, policies, and products make it critical to have a common taxonomy and risk perspective to assist organizations and individuals in how they use and communicate about AI.

- AI is evolving rapidly, which can signal rapid shifts in AI guidance or which cybersecurity activities are priorities. Today's AI priorities may be normal business practices in the near future. Including a prioritization schema could date the Profile and prohibit it from being flexible to AI shifts and changes.
- As a technology Profile, the emphasis should be on capturing AI-specific considerations for Subcategories in each Focus Area.
- The rationales, considerations, opportunities (for Defend), and Informative References in the Profile help organizations prioritize by informing risk management and resource-related discussions when implementing a specific Subcategory.
    - Prioritization is an optional component of Community Profiles. Prioritization can work well for sector-focused Community Profiles. However, technology-focused Community Profiles apply to a much broader notion of "community", and it can be difficult to identify priorities that are appropriate for this wider variety of applications of a Community Profile.
    - There are existing technology CSF Community Profiles (e.g., Position, Navigation, and Timing [PNT] Profile, Hybrid Satellite Networks [HSN] Profile) that do not include a prioritization schema. These Profiles provide guidance in tables for applicable Subcategories from Categories across all CSF Functions (absent priorities) that Profile implementers can use to make informed decisions for how to apply the Profile to specific missions, systems, or data.
- For consistency among all Focus Areas, priorities would need to be based on considerations and not include opportunities (Defend has Subcategories that may include narrative for both considerations and opportunities but only considerations impact prioritization). Removing proposed Subcategory priorities eliminates potential ambiguity or confusion as to whether any Subcategory prioritizations in Defend reflect considerations, opportunities, or both.

- Eliminates confusion around prioritizations labeled “foundational” or “supporting” and how those may be the most critical for some organizations, or conflating urgency with prioritization. Organizations may use the Profile to justify resource planning requests. If there are Subcategories prioritized with the lowest priority in the Profile but would be considered critical and therefore most important / highest priority to implement in an organization, the Profile may inhibit an organization’s ability to justify cybersecurity outcomes that do not align with the Profile’s prioritization designations.

If, as the community continues to provide feedback to the NCCoE on the best approach, prioritization is retained, the following list contains proposed ways forward with the prioritization schema. Terminology discussion may highlight there is a challenge or concerns about what prioritization means in the context of the Cyber AI Profile and its implementation by a broad user community.

Proposed ways forward on the prioritization level definitions (*if prioritization is retained)*:

1. Leave prioritization level definitions as currently defined in the Profile:
   - “1” / High Priority: These Subcategories are considered the most critical to address the challenges for a Focus Area. High priority Subcategories should typically be addressed immediately given available resources.
   - “2” / Moderate Priority: These Subcategories should be the next priority after implementing High Priority Subcategories.
   - “3” / Foundational Priority: These Subcategories are generally important to the Focus Area but may not require the same level of urgency as higher priorities. Note that “Foundational” does not equate to low priority. All Subcategories should receive consideration.
2. Revise prioritization level definitions to simplify:
   - Highest priority represents the most critical Subcategories to enable a Focus Area and should be addressed immediately given available resources. (Options for highest priority level terms: Very High, Advanced, Elevated, Surpasses, High, 1)
   - Middle priority represents that a Focus Area has a dependency on the Subcategory, but it is not as critical as High Priority Subcategories. They should be the next priority after implementing High Priority Subcategories and may

become a higher priority in certain contexts or environments. (Options for middle priority level terms: Improve, Moderate, Medium, 2)

- Lowest priority represents Subcategories that are generally important to a Focus Area but may not require the same level of urgency as higher priority Subcategories. Note that the lowest priority does not equate to no priority. All CSF Subcategories should receive consideration by an organization. (Options for lowest priority level terms: Foundational, Supporting, Average, 3)

Subcategory Considerations

The Profile currently includes the phrase "standard cybersecurity practices apply" to indicate that no additional AI-specific considerations were identified. However, feedback suggests this phrase is ambiguous and may obscure where AI changes cybersecurity risk.

Overall, the phrase is used to indicate that there are no unique Subcategory-related considerations identified for a Focus Area and the activities of an organization's existing cybersecurity program are sufficient for AI systems. These considerations also include the rationale for Subcategories that are designated as a High (1) or Moderate (2) proposed priority.

The phrase identifies where AI-specific needs do not extend beyond what a "typical" cybersecurity program does for "traditional" systems. In such situations, "traditional cybersecurity practices" would imply non-AI practices and that these cybersecurity practices are sufficient to apply to AI applications.

Throughout collaboration and outreach engagement for this Profile, stakeholders have consistently felt strongly the Profile should not create something new for AI cybersecurity, and that the Profile should build on what already exists.

While a Subcategory with this phrase may be applied to AI applications, the addition of AI does not change cybersecurity practices in a material way (e.g., encrypting data at rest - encrypting model weights is not different from encrypting data in cold storage).

Overview of comments

Analysis of the comments related to the "standard cybersecurity practices apply" statement indicate that clarification in the statement is desired, to potentially include revising the statement itself. This need for clarification is informed by the following factors:

- "Standard cybersecurity practices" is an ambiguous phrase and should be replaced with a phrase that focuses on AI-specific considerations (or explicit rationales where feasible).

- “Standard cybersecurity practices” may already be used by individual organizations. Without clarification, this phrase could be interpreted as introducing additional, parallel, or even conflicting expectations beyond those already addressed through regulatory review or organizational guidance.
- In regulated environments, “standard cybersecurity practices” may include existing regulatory cybersecurity, safety, and quality system requirements.
- The phrase “Standard cybersecurity practices apply” appears often in guidance documents, and sometimes this can cloud the distinct requirements posed by AI systems. To clarify these areas, it could be worthwhile to agree on minimum standards for things like telemetry, configuration manifests, provenance checkpoints, agent privilege policies, and recovery acceptance tests. These measures can support secure and well-monitored AI operations.
- LLMs (e.g., GenAI) involve unique elements like dynamic inference, model weights prone to extraction/theft, prompt attacks, and training data risks that exceed standard controls. Modifying the phrase enables the Profile to allow for innovation rather than reductive for AI complexities.
- There would be value in linking the phrase to cybersecurity best practices or acknowledge their existing guidance as organizations may already be using the phrase.
- There should be language (within the phrase or as part of the context that defines the phrase) that addresses that organizations should still ensure those established practices are applied to AI components, their integrations, and their output and action channels so results remain verifiable and actionable.

Potential open questions for discussion during the working session:

- Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
- Is there a difference in the way the phrase may be applied for the same Subcategory and Focus Area? Are there contextual differences that would influence how this phrase may be applied for the same Subcategory across the three Focus Areas?
- Does the intent behind including such a phrase provide value or does its inclusion possibly create confusion related to distinct requirements posed by AI systems?
- Does the phrase need language that addresses if traditional cybersecurity practices meet a safeguarding need?

NIST's proposed options forward on the definition based on analysis of feedback received:

1. Leave definition as is, which reads, "The phrase 'standard cybersecurity practices apply' is used to indicate that there are no unique considerations identified for the Focus Area and the activities of the cybersecurity program are sufficient for AI systems. These considerations also include the rationale for Subcategories that are designated as a High (1) or Moderate (2) proposed priority. For Defend, these considerations are only related to the opportunities listed."
2. Modify the definition to read, "When the phrase '[placeholder for phrase]' is used, it indicates that no AI-specific considerations (e.g., techniques, dependencies, needs, modifications) for the Focus Area were identified; organizations should still ensure applicable cybersecurity safeguards are applied to AI components, their integrations, and their output and action channels so results remain verifiable and actionable within the context of acceptable risk."
3. Modify the definition to read, "When the phrase '[placeholder for phrase]' is used, it indicates that no AI-specific considerations (e.g., techniques, dependencies, needs, modifications) for the Focus Area were identified, and the activities of the cybersecurity program are sufficient for AI systems and within the context of acceptable risk." **(NIST's preferred approach)**

Based on community feedback thus far, NIST identified Option 3 as the most viable option. The language in this revised definition offers an exact description of what the phrase is intended to convey.

NIST's proposed ways forward on the phrase based on analysis of feedback received:

1. Leave "standard cybersecurity practices apply" as is.
2. Revise phrase to "No AI-specific related considerations were identified."
3. Revise phrase to "No AI-specific cybersecurity program considerations were identified." **(NIST's preferred approach)**

NIST's preferred approach is Option 3 (Revise the phrase to "No AI-specific cybersecurity program considerations were identified."). Below are points for why NIST initially prefers this approach and would like to hear the community's feedback on:

- Eliminates confusion about what is considered a standard cybersecurity practice.

- Enables the reader to more easily understand where there are, and are not, any AI-specific considerations that should be documented relevant to the Focus Area and a Subcategory.
- Allows for flexible and broad use by various communities and organizations as what is considered a standard cybersecurity practice may be different among different implementers.
- Inclusion of a phrase minimizes potential confusion that could happen from absence of content in the table (i.e., blanks).
- Indicates that traditional cybersecurity practices remain important for AI and that the Profile is building on them, not creating something totally different or "reinventing the wheel" for AI. Inclusion of a phrase such as this is consistent with the approach taken in other NCCoE Community Profiles and has been validated by Cyber AI Profile COI.

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 will discuss 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	April 28, 2026 / 1:00-4:00 P.M. EDT
Session #2: Extending the Technical Content	This session will explore how the Profile is being strengthened in critical technical areas including Agentic AI and Zero Trust	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 Discussion Essays for Sessions #2 and #3. We look forward to hearing from you!