Cyber AI Profile COI Working Sessions: Profile Elements

April 28, 2026

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Agenda

- Cyber AI Profile Project Overview
- Control Overlays for Securing AI Systems (COSAiS) Overview
- Today’s Plan
- Proposed Priorities Discussion
- Break
- “Standard Practices” Discussion
- Open Discussion
- Close-out

Engagement

Via Zoom:

We would love to hear from you!

- Submit questions during Q&A
- Raise your virtual hand to be unmuted to speak (remember to unmute on your end, too!)
- Enable captioning

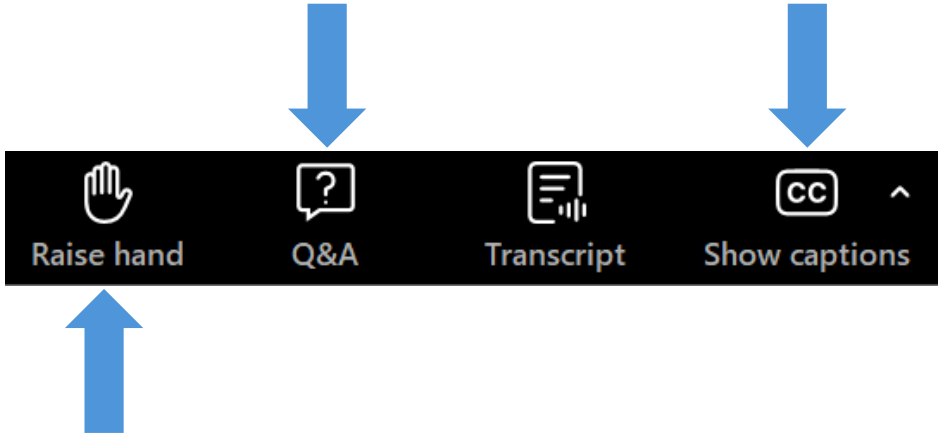

Via Slido:

Please use Slido to participate throughout the session. Scan the QR code or go to Slido.com and enter the access code to access the polls as they are opened.

Slido.com
#CyberAI_Spring2026-1

Cyber AI Profile Project Overview

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Cybersecurity, Privacy, and AI

The diverse use and rapid proliferation of Artificial Intelligence (AI) promises unique value for industry, consumers, and broader society, but like many technologies, to recognize these benefits to the greatest potential, new risks from these advancements in AI must be managed.

In NIST's Applied Cybersecurity Division (ACD), our key concern is how advancements in the broad adoption of AI may impact current cybersecurity and privacy risks and risk management approaches.

https://www.nist.gov/itl/applied-cybersecurity/cybersecurity-privacy-and-ai

The Case for a Cyber AI Profile

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Purpose:

Support cybersecurity programs as they manage the impacts of advancements in AI to their organization

Areas of focus:

- Cybersecurity risks that arise from the use of AI by organizations, including securing AI systems, components, and machine learning infrastructures, and minimizing data leakage.
- Determining how to defend against AI-enabled attacks.
- Assisting organizations in the use of AI with their cyber defense activities and using AI to improve privacy protections.

Outcomes:

- Establishes a shared understanding of AI-related cybersecurity terminology and considerations
- Fosters collaboration and communication across the AI and cybersecurity communities
- Enables organizations to measure their current practices, understand the available references, and identify gaps to update Organizational Profiles and roadmaps

Cyber AI Profile Roadmap

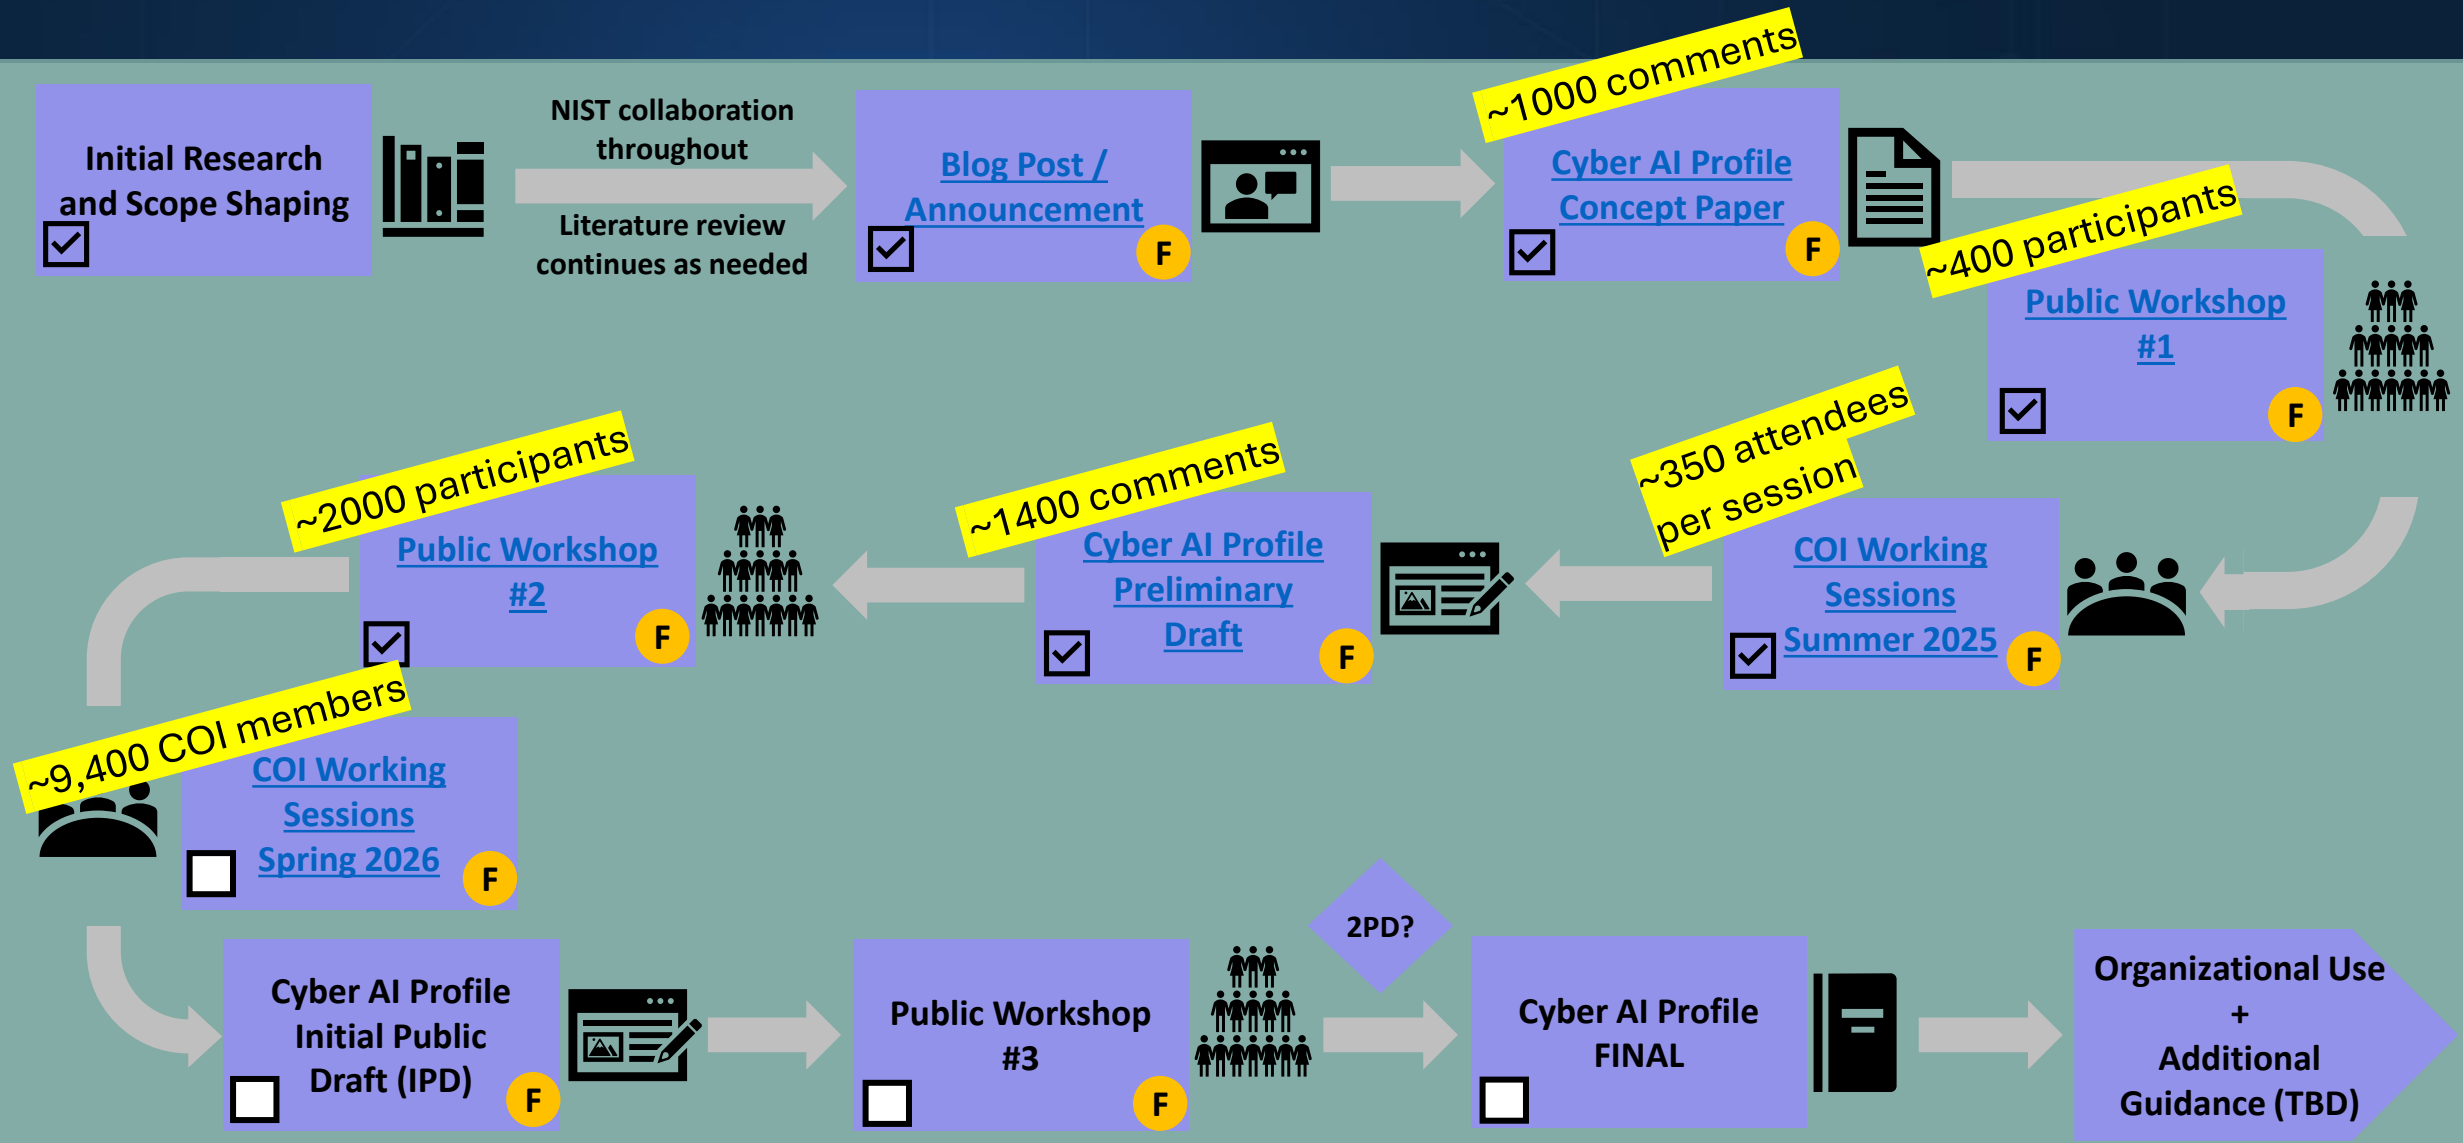

F Opportunities for COI/public stakeholder feedback (NOTE: Internal NIST collaboration occurs throughout)

Overall Outcomes of COI Engagement

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Concept Paper

- Confirmed the need for guidelines
- Agreed the NIST CSF was the right starting point
- Identified key challenges

Workshop #1

- Validated direction of the Profile
- Identified critical topics
- Proposed priorities and resources

Summer 2025 COI Working Sessions

- Refined cybersecurity priorities for each Focus Area

Workshop #2

- Continued refining Profile direction and content
- Shared insights regarding implementation

Preliminary Draft Comments

- Informed continued build out of Profile
- Raised areas for further engagement with the COI

Workshop #1 Blog Post

Workshop #2 Blog Post

Additional Resources

Cyber AI Profile

- NIST Cybersecurity, Privacy, and AI Program
- Blog post: Managing Cybersecurity and Privacy Risks in the Age of Artificial Intelligence: Launching a New Program at NIST | NIST
- NCCoE Project Page: Cyber AI Profile
- Cybersecurity and AI Workshop Concept Paper (posted in advance of the April 3, 2025, workshop)
- April 3rd 2025 Cyber AI Profile Workshop recording
- Blog post: Reflections from the First Cyber AI Profile Workshop
- Blog post: Reflections from the Second Cyber AI Profile Workshop
- Cyber AI Profile COI Working Sessions Introduction Video

NIST Cybersecurity Framework

- NIST CSF
- NIST CSF FAQs
- NIST CSF 2.0 Informative References
- NIST CSF Events

NIST Resources for Applying NIST Frameworks

- Resources for Applying NIST Frameworks

Community Profiles

- Examples of Community Profiles
- Creating Community Profiles FAQs

Cyber AI Profile Project Page

Control Overlays for Securing AI Systems (COSAiS) Update

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Control Overlays for Securing AI Systems

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY U.S. DEPARTMENT OF COMMERCE

The controls to manage cybersecurity risks to AI systems will ***largely be similar*** to those required for any type of software.
Many organizations are ***familiar with the SP 800-53 controls*** and may already be implementing them.
The SP 800-53 controls offer ***flexibility*** to meet the ***unique security considerations for AI systems***.

SP 800-53 CONTROLS TO MANAGE RISK FOR SPECIFIC **TYPES** AND **USES** OF AI SYSTEMS

COMMON **TECHNICAL FOUNDATION** FOR CYBERSECURITY OUTCOMES

IMPLEMENTATION-FOCUSED FOR DIFFERENT **AI USE CASES**

ORGANIZATIONS **USING** AI SYSTEMS
AI SYSTEM **DEVELOPERS**
CYBERSECURITY COMMUNITY

CAN LEVERAGE EXISTING **ASSESSMENT** GUIDELINES (SP 800-53A)

PROVIDES LINKS TO OTHER KEY CYBER & AI **NIST PUBS**

Development Methodology

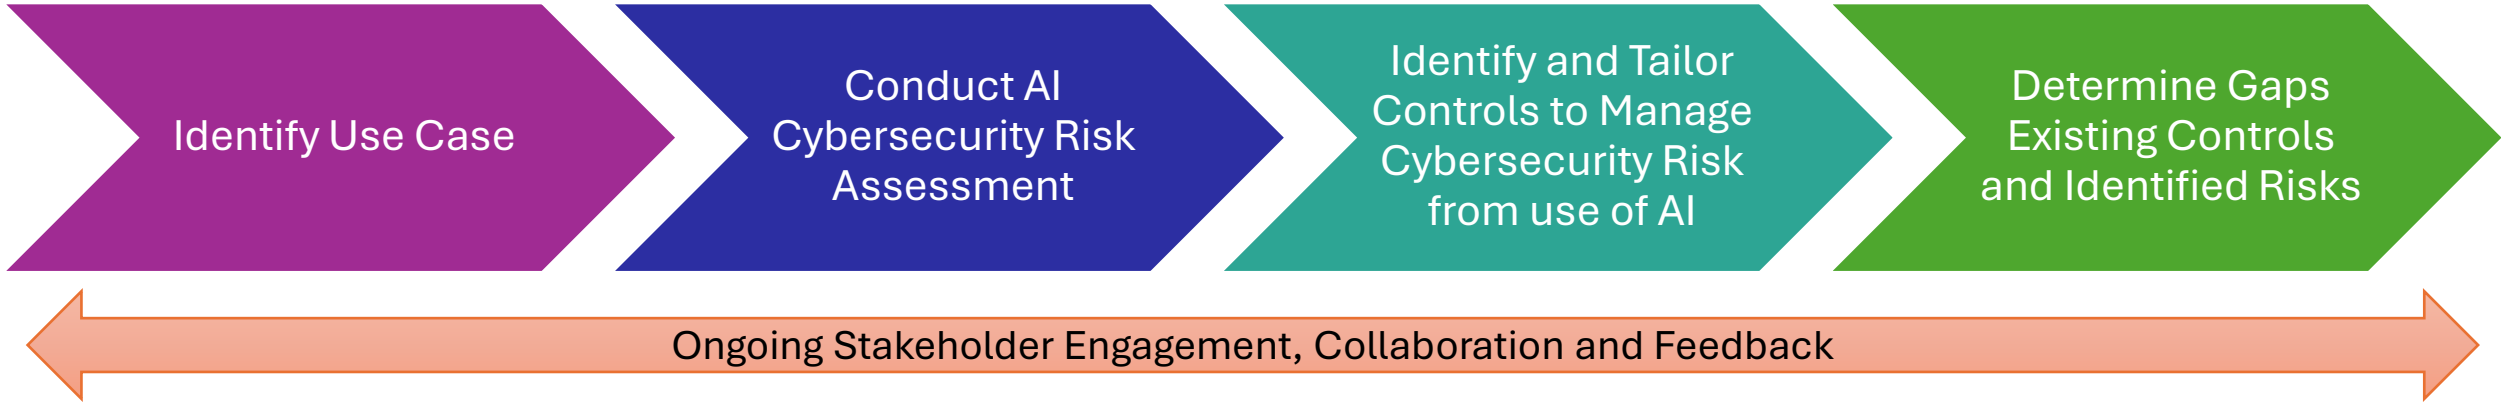

This methodology assumes that the organization has:

- an organization-wide information security program in place to manage cybersecurity risk and includes a risk management strategy with explicit risk tolerance
- used and/or is familiar with the NIST Risk Management Framework and SP 800-53 controls
- selected and implemented controls to manage IT and systems cybersecurity risk (including governance)
- implemented a process to assess and monitor controls

https://csrc.nist.gov/Projects/cosais

Current Progress: Overlay on Securing Predictive AI Systems

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY U.S. DEPARTMENT OF COMMERCE

Access the Annotated Outline:

Identify Use Case → Conduct AI Cybersecurity Risk Assessment → Identify and Tailor Controls to Manage Cybersecurity Risk from use of AI → Determine Gaps Existing Controls and Identified Risks

Predictive AI*

A) On-premises and third-party (cloud) model; proprietary data

B) On-premises and third-party (cloud) model; public-facing input data

Notional Risk Assessment

Leverage: ***Use Case*** and ***NIST AI 100-2e2025*** to identify ***attacks and mitigations***, and ***SP 800-53 controls*** to manage risk

Drafting – Methodology & Predictive AI Overlay

NISTIR 8605, Overview and Methodology

NISTIR 8605A, Controls for Securing Predictive AI Systems

**Updated from concept paper*

https://csrc.nist.gov/Projects/cosais

Annotated Outline: Overlay on Using & Fine-Tuning Predictive AI

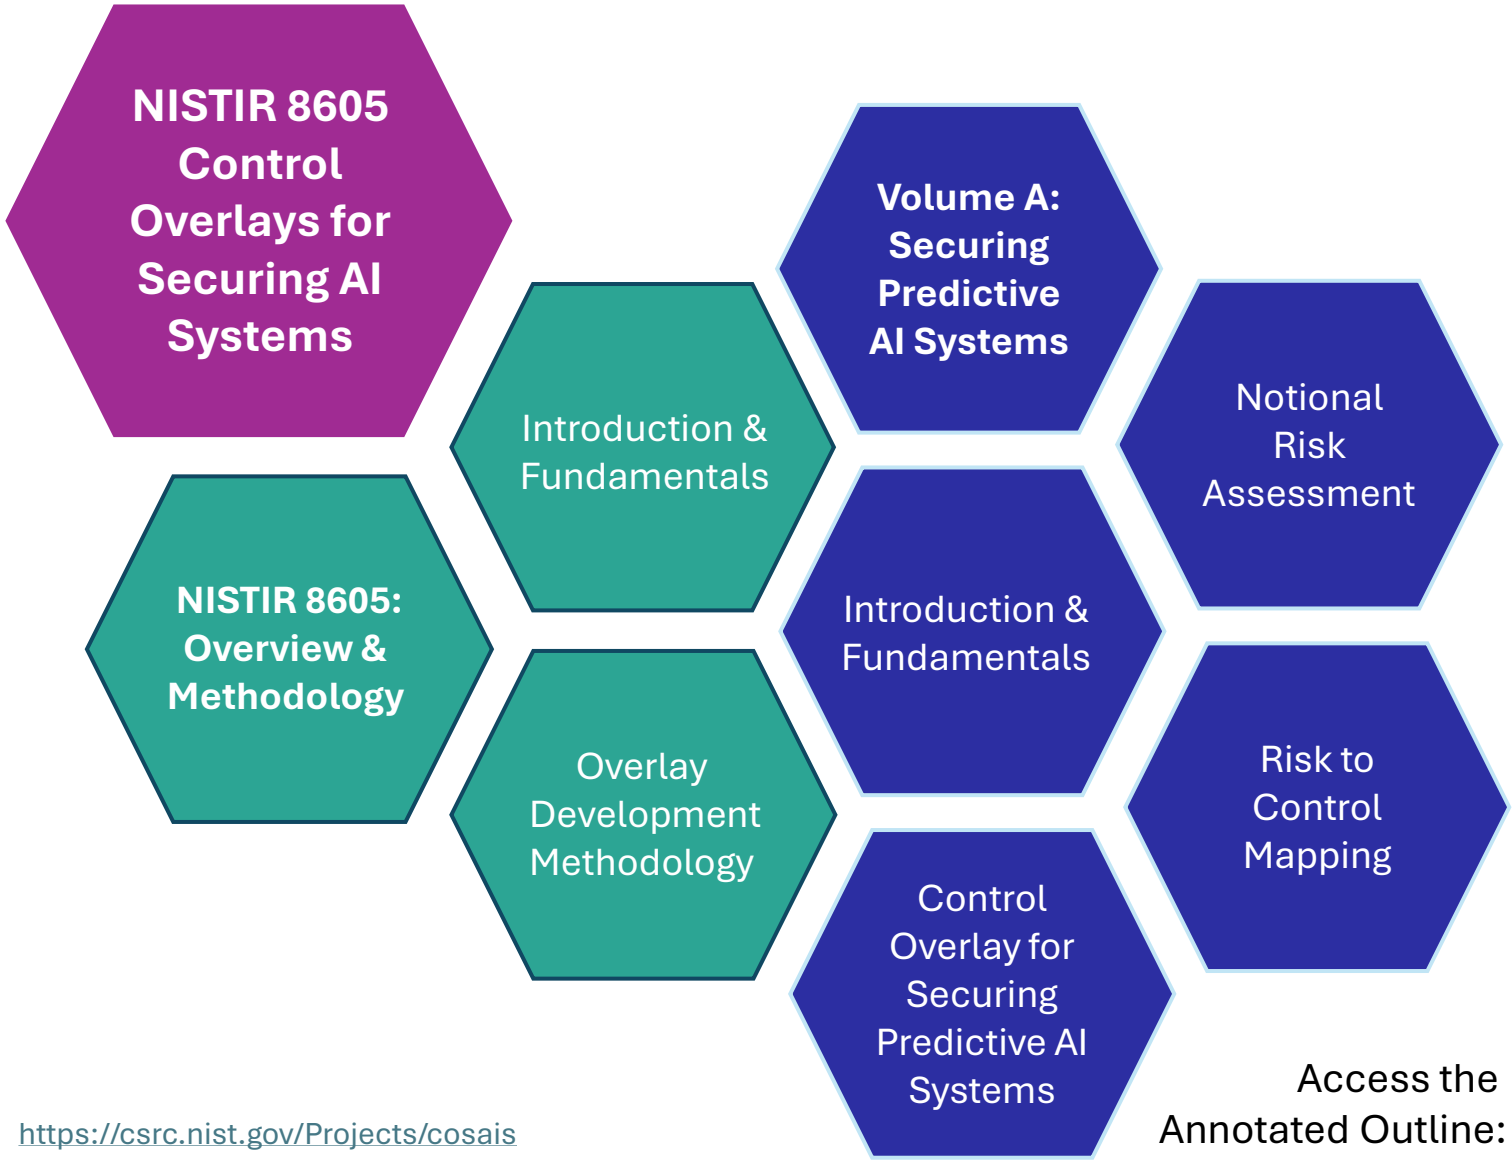

Access the Annotated Outline:

https://csrc.nist.gov/Projects/cosais

Planned for 2026 - early 2027

Volume B: Securing Generative AI Systems

Volume C: Controls for Developing Secure AI Systems

Volume D: Controls for Agentic AI Systems

Next Steps and Get Engaged

Proposed Deliverables: Control Overlays for Securing AI Systems

- NISTIR 8605, Overview and Methodology
- NISTIR 8605A, Controls for Securing Predictive AI Systems
- NISTIR 8605B, Controls for Securing Generative AI Systems
  NISTIR 8605C, Controls for Developing Secure AI Systems
- NISTIR 8605D, Controls for Securing Agentic AI Systems

https://csrc.nist.gov/Projects/cosais overlays-securing-AI@list.nist.gov

Today's Plan

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

How You Contribute Today

- **Please raise your virtual hand or type in the chat to contribute**
- Members of the press, please identify yourself and your organization
- Be respectful of others
- Please don’t be shy – we would love to hear from everyone!
- **Please remain on mute when not speaking**
- **We will use Slido to facilitate some of our discussions**

Using Slido

- We will be using Slido to facilitate some of our discussions
- Options to join via QR code or URL + event code
- Works on mobile phone and computer
- Responses are anonymous

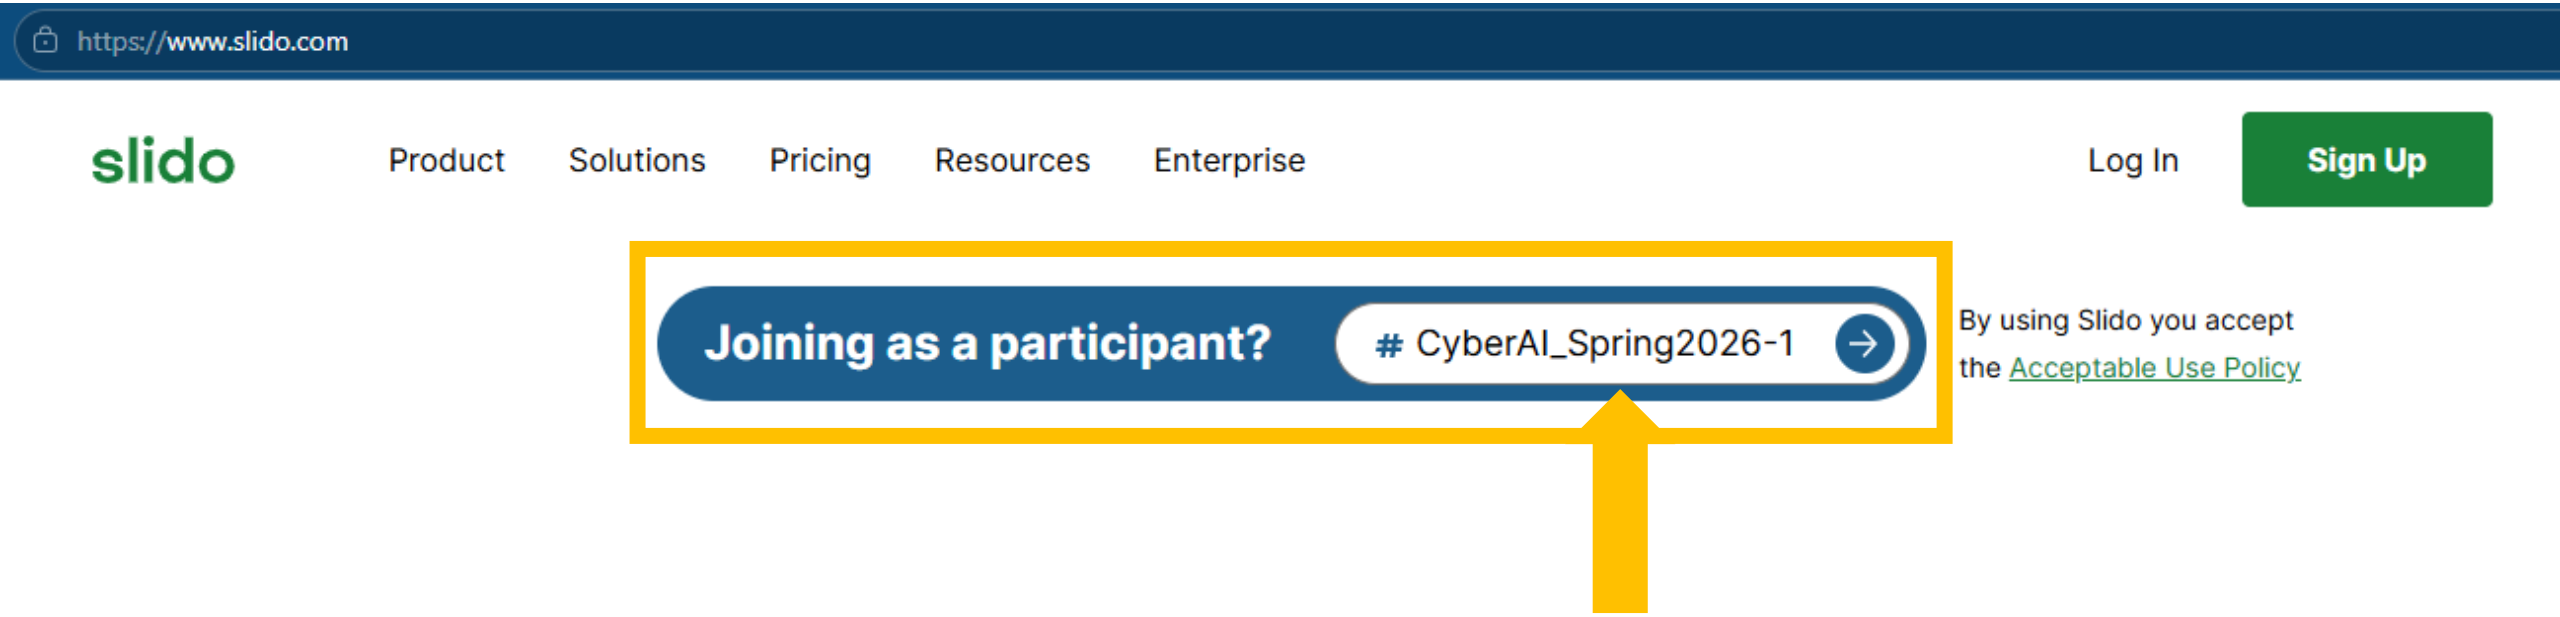

Slido Screens

No Active Polls

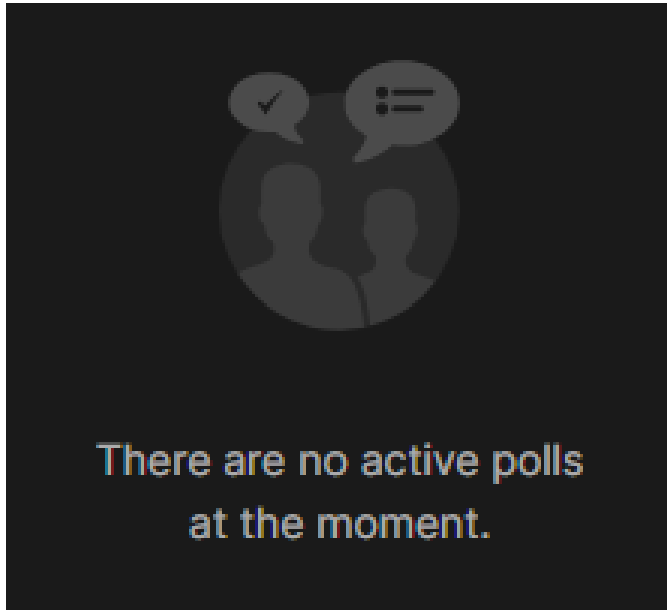

Active Polls

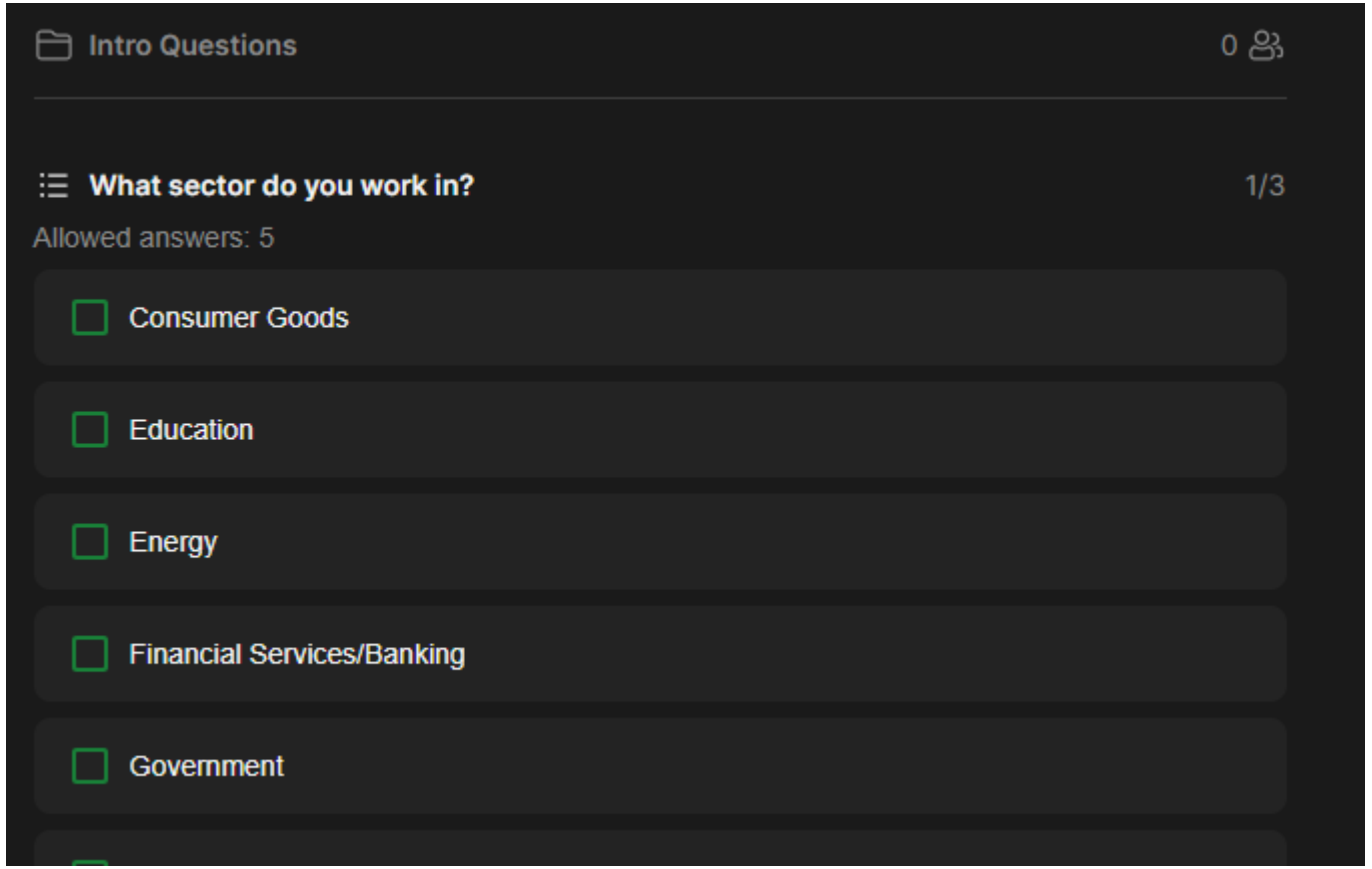

Slido: Getting to Know You

- What sector do you work in?
- Which NIST Frameworks does your organization use?

Slido.com
#CyberAI_Spring2026-1

Slido: Getting to Know You

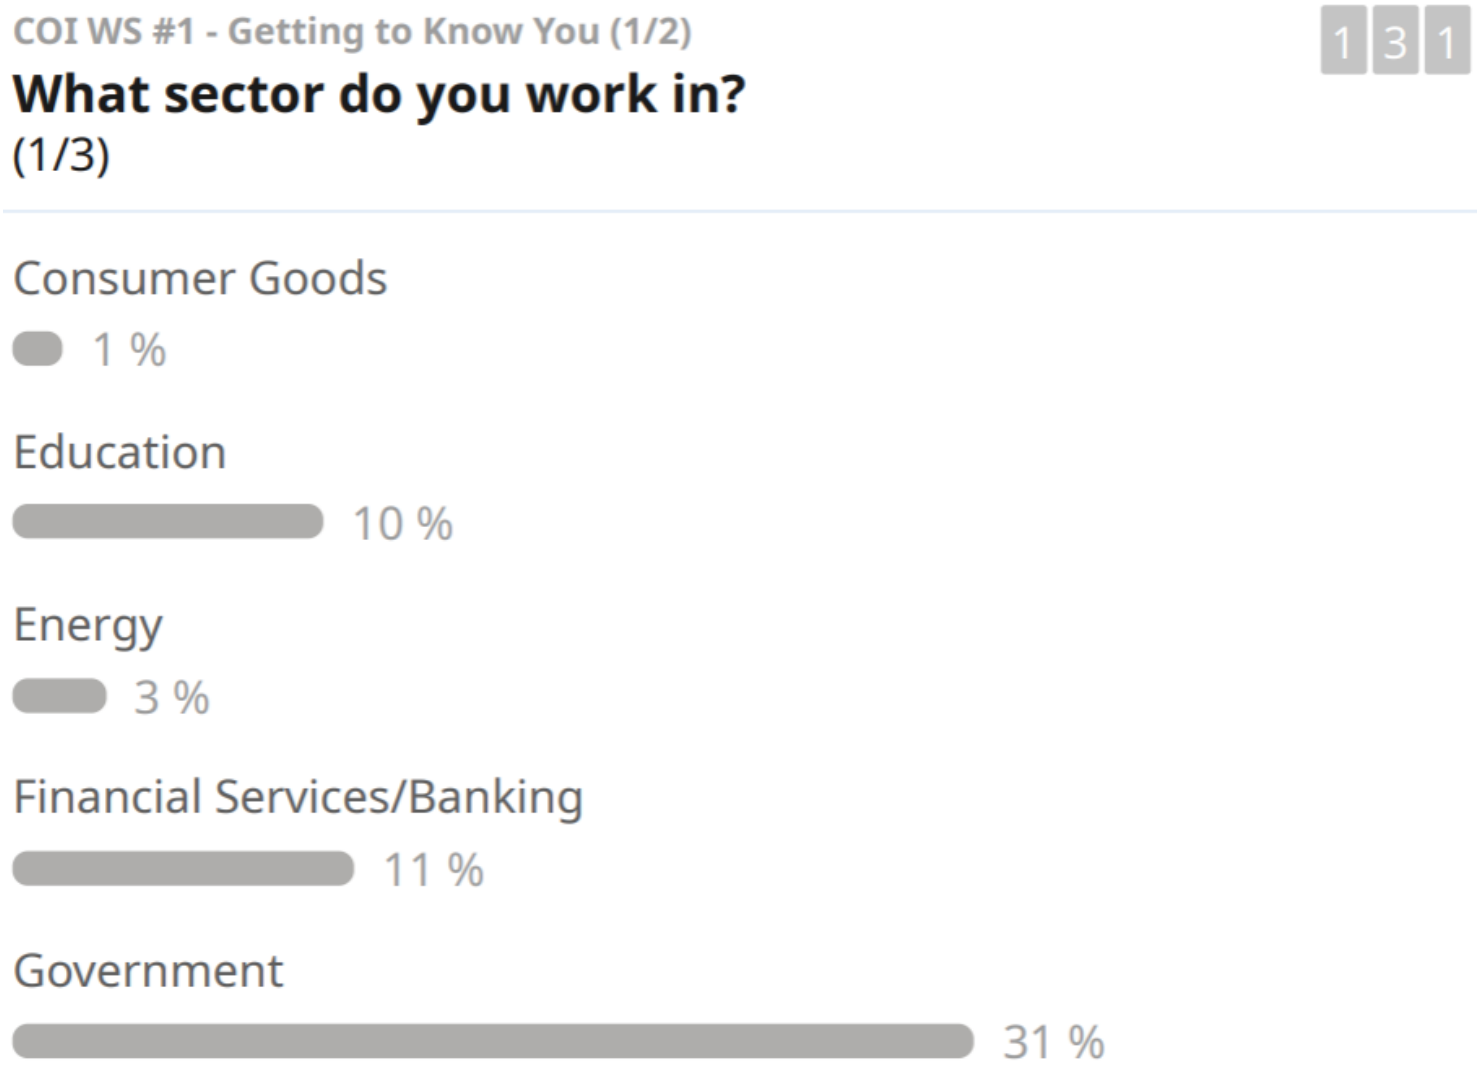

Slido: Getting to Know You

Slido: Getting to Know You

COI WS #1 - Getting to Know You (1/2)

1 3 1

What sector do you work in?
(3/3)

Telecommunications
4 %

Think Tank
1 %

Trade Association
1 %

Transportation
1 %

Other (please add in Zoom chat)
4 %

Slido: Getting to Know You

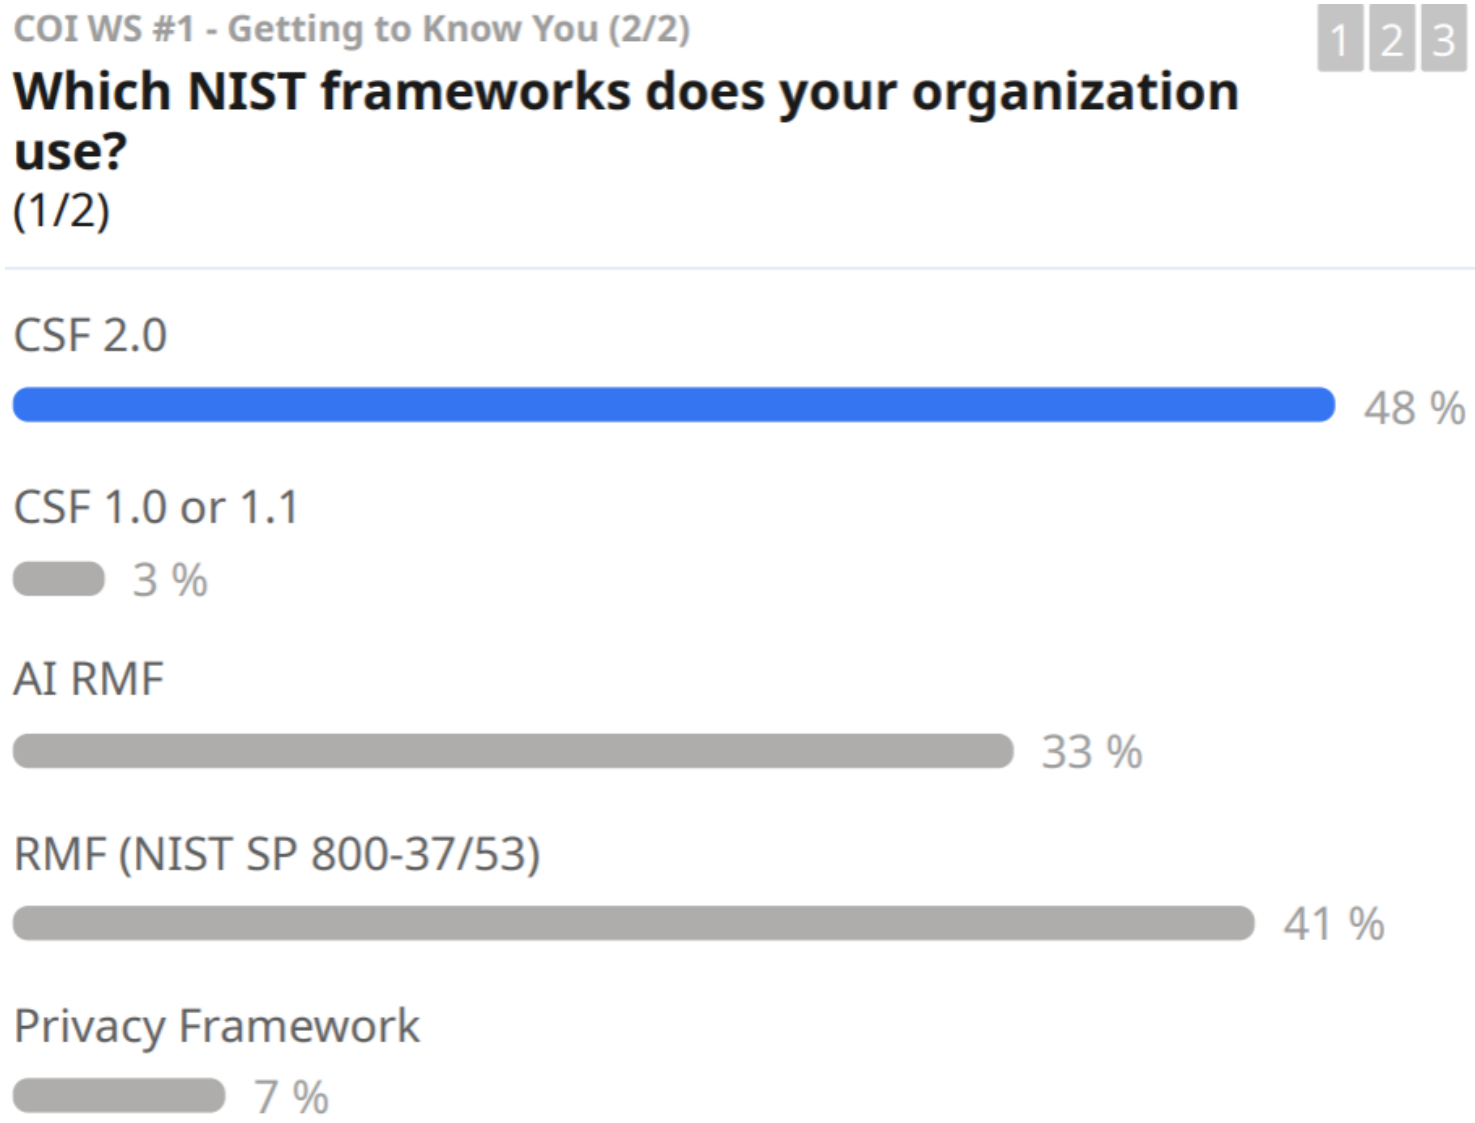

Slido: Getting to Know You

General Discussion Plan

- Summary of feedback on the topics (Proposed Priorities and “Standard Practices”)
- Review proposed options
- Facilitated discussion
- How we plan to use this feedback

Today's Focus: Profile Elements

- Options for the Subcategory proposed priorities
- Use and meaning of the phrase "standard cybersecurity practices apply" in the Subcategory considerations

Discussion Essay #1

Cyber AI Profile Draft

Proposed Priorities Discussion

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Proposed Priorities: Summary of Feedback

- Increased clarity around definitions and the methodology used to determine the prioritizations
- Change the prioritization terms
- Add a new prioritization level
- Refine the definitions of the priority level terms
- Clarify meaning of different priorities for the same Subcategory across the three Focus Areas

Proposed Priorities: Potential Options

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Option 1

No Change

Retain the existing 3-level prioritization schema and refine levels and definitions

Option 2

Add a new prioritization level

Replace the current schema with a 4-level prioritization

Option 3

Replace Schema

Reflect the approximate relative degree of differences between "traditional" cybersecurity considerations and AI-specific cybersecurity considerations

Option 4

Remove Priorities

Do not include a prioritization schema and therefore no proposed Subcategory priorities

NIST's preferred option

Proposed Priorities: Potential Option 1

Option 1

No Change

Retain the existing 3-level prioritization schema and refine levels and definitions

CSF 2.0 Core: PROTECT	General Considerations	Focus Area Proposed Priorities & Considerations		
		Secure	Defend	Thwart
PROTECT (PR)	Safeguards to manage the organization's cybersecurity risks are used			
Identify Management, Authentication, and Access Control (PR.AA)	Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access			
PR.AA-01: Identities and credentials for authorized users, services, and hardware are managed by the organization	**General Considerations:** Give AI systems unique and traceable identities and credentials to better track their activity. **Example Informative References:** NIST SP 800-53, Rev 5: A… 01; AC-02: A…	**Proposed Priority: 1** **Sample Focus Area Considerations:** AI systems may need their own identities and credentials (i.e., AI service level ac…unts) to interact with a …der system. …izations need …bility between AI …ems and their actions. **Example Informative References:** DASF 1, 3, 21-22, 32, 40, 48; ATLAS AML.M0005; …	**Proposed Priority: 2** **Sample Opportunities:** AI catches credential misuse that previous rules might miss by flagging unusual authentication activity. **Sample Focus Area Considerations:** Assign and manage unique and traceable identities and credentials to AI defense agents to support defensive response activities. **Example Informative References**: DASF 39; WASP AI Exchange: Model Access Control; …	**Proposed Priority: 1** **Sample Focus Area Considerations:** Standard cybersecurity practices apply. (Rationale) AI-enabled cyber attacks will lower the barrier of entry to gaining access to identities and credentials, services, and hardware. **Example Informative References:** NIST SP 800-172 – Configuration Management (3.4); NIST SP 800-172 Identification and Authentication (3.5); …

Proposed Priorities: Potential Option 2

Option 2

Add a new prioritization level

Replace the current schema with a 4-level prioritization (to include a new level for cybersecurity activities that should be addressed first)

CSF 2.0 Core: PROTECT	General Considerations	Focus Area Proposed Priorities & Considerations		
		Secure	Defend	Thwart
PROTECT (PR)	Safeguards to manage the organization's cybersecurity risks are used			
Identify Management, Authentication, and Access Control (PR.AA)	Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access			
PR.AA-01: Identities and credentials for authorized users, services, and hardware are managed by the organization	**General Considerations:** Give AI systems unique and traceable identities and credentials to better track their activity. **Example Informative References:** NIST SP [illegible]; IA-10; IA-	**Proposed Priority: 4** **Sample Focus Area Considerations:** AI systems may need their own identities and credentials [illegible] AI service level [illegible]nts) to interact with a [illegible]er system. [illegible]zations need traceability between AI systems and their actions. **Example Informative References:** DASF 1, 3, 21-22, 32, 40, 48; ATLAS AML.M0005; …	**Proposed Priority: 2** **Sample Opportunities:** AI catches credential misuse that previous rules might miss by flagging unusual authentication activity. **Sample Focus Area Considerations:** Assign and manage unique and traceable identities and credentials to AI defense agents to support defensive response activities. **Example Informative References**: DASF 39; WASP AI Exchange: Model Access Control; …	**Proposed Priority: 1** **Sample Focus Area Considerations:** Standard cybersecurity practices apply. (Rationale) AI-enabled cyber attacks will lower the barrier of entry to gaining access to identities and credentials, services, and hardware. **Example Informative References:** NIST SP 800-172 – Configuration Management (3.4); NIST SP 800-172 Identification and Authentication (3.5); …

Proposed Priorities: Potential Option 3

Option 3

Replace Schema

Reflect the approximate relative degree of differences between "traditional" cybersecurity considerations and AI-specific cybersecurity considerations

CSF 2.0 Core: PROTECT	General Considerations	Focus Area Proposed Priorities & Considerations		
		Secure	Defend	Thwart
PROTECT (PR)	Safeguards to manage the organization's cybersecurity risks are used			
Identify Management, Authentication, and Access Control (PR.AA)	Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access			
PR.AA-01: Identities and credentials for authorized users, services, and hardware are managed by the organization	**General Considerations:** Give AI systems unique and traceable identities and credentials to better track their activity. **Example Informative References:** NIST SP 800-53 [illegible] 01; AC [illegible]	**Proposed Priority:** Minimal **Sample Focus Area Considerations:** AI systems may need their own identities and credentials (i.e., AI service level [illegible]unts) to interact with a [illegible]ler system. [illegible]izations need [illegible]ility between AI systems and their actions. **Example Informative References:** DASF 1, 3, 21-22, 32, 40, 48; ATLAS AML.M0005; …	**Proposed Priority:** Significant **Sample Opportunities:** AI catches credential misuse that previous rules might miss by flagging unusual authentication activity. **Sample Focus Area Considerations:** Assign and manage unique and traceable identities and credentials to AI defense agents to support defensive response activities. **Example Informative References**: DASF 39; WASP AI Exchange: Model Access Control; …	**Proposed Priority:** Substantial **Sample Focus Area Considerations:** Standard cybersecurity practices apply. (Rationale) AI-enabled cyber attacks will lower the barrier of entry to gaining access to identities and credentials, services, and hardware. **Example Informative References:** NIST SP 800-172 – Configuration Management (3.4); NIST SP 800-172 Identification and Authentication (3.5); …

Proposed Priorities: Potential Option 4

Option 4

Remove Priorities

Do not include a prioritization schema and therefore no proposed Subcategory priorities

NIST's preferred option

CSF 2.0 Core: PROTECT	General Considerations	Focus Area Proposed Priorities & Considerations		
		Secure	Defend	Thwart
PROTECT (PR)	Safeguards to manage the organization's cybersecurity risks are used			
Identify Management, Authentication, and Access Control (PR.AA)	Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access			
PR.AA-01: Identities and credentials for authorized users, services, and hardware are managed by the organization	**General Considerations:** Give AI systems unique and traceable identities and credentials to better track their activity. **Example Inform… Refe…** …A-05; IA-06; …A-08; IA-09; IA-10; IA-11	**Sample Focus Area Considerations:** AI systems may need their own identities and credentials (i… AI service level …unts) to interact with a …er system. …zations need …ceability between AI systems and their actions. **Example Informative References:** DASF 1, 3, 21-22, 32, 40, 48; ATLAS AML.M0005; …	**Sample Opportunities:** AI catches credential misuse that previous rules might miss by flagging unusual authentication activity. **Sample Focus Area Considerations:** Assign and manage unique and traceable identities and credentials to AI defense agents to support defensive response activities. **Example Informative References**: DASF 39; WASP AI Exchange: Model Access Control; …	**Sample Focus Area Considerations:** Standard cybersecurity practices apply. (Rationale) AI-enabled cyber attacks will lower the barrier of entry to gaining access to identities and credentials, services, and hardware. **Example Informative References:** NIST SP 800-172 – Configuration Management (3.4); NIST SP 800-172 Identification and Authentication (3.5); …

The Profile would retain rationales, considerations, opportunities, and Informative References to support informed risk management discussions and decisions by Profile implementers.

Proposed Priorities: NIST's Preferred Option

Option 4

Remove Priorities

Do not include a prioritization schema and therefore no proposed Subcategory priorities

- **Enables broad and flexible application**
- **Supports use of the Profile with other CSF Profiles**
- **Supports long-term use as priorities may become dated as AI evolves rapidly**
- **Places emphasis on capturing AI-specific considerations**
- **Allows the rationales, considerations, opportunities (for Defend), and Informative References to inform priorities**
- **Eliminates potential ambiguity as to whether any prioritizations in Defend reflect considerations, opportunities, or both**

Proposed Priorities: Questions for Discussion

- What benefits or challenges does having a prioritization schema introduce?

- What benefits or challenges would be introduced by not including priorities for each Subcategory?

Slido.com
#CyberAI_Spring2026-1

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(1/17)

- - I like opt2 because it adds a OWASP top10-style level: you have to do at least this/you have to be this high to go in this ride. - Why not have number and word together? - It may be easier to normalize against members of the same industry.

  Like those (annoying) magic quadrants show, businesses like to know how they are doing with respect to others. - Another thing about priorities is they do not need to be set in stone: take OWASP Top 10. Every year they publish them and they change. That show trends in time
- Prioritization schema will help is selecting without objective thought, without that

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(2/17)

the judgement matters. The challenge in having a prioritization schema may be of generalized in nature which may have e difference from the sector/ area the same is applied.

- Repeatability.
- Benefits guide organizations on what to focus on first. However, the prioritization is at risk of getting stale if not evaluated and updated from time to time.
- Benefit - the "highest" priority tends to be thr spearhead where the focus is Challenge - those on the "lowest" priority level tends to recieve no funding, even if it's also important
- This enables organisations to focus resources on the highest risk threats first, reducing mean time to detect and respond while avoiding alert fatigue from low-severity noise. It aligns security operations with business impact, ensuring critical assets

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(3/17)

receive proportional protection. However challenges include the risk of deprioritising novel or low-confidence threats that evolve rapidly, maintaining schema accuracy as attack surfaces change, and introducing bias into AI that may systematically underweight emerging threat categories.

- A priority might be high in say one year, but as threats and technology etc change and evolve, a high priority could become low.
- Data controls don't seem to be focused on transparency risk as a priority with the organizational considerations (re: number 4).
- The prioritization is crucial especially for companies with limited resources.
- I think this was covered in the talk - benefits are that it manages expectations of range

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(4/17)

of options and puts them in order, offering guidance and structure. Challenges would be keeping this relevant in specifics, while not becoming instantly dated once technology changes influence the prioritization frameworks. Also keeping items universally applicable across vast sectors and disparate use cases and resources, could be problematic.

- The prioritization schema offers efficiency in an ever-evolving AI environment. Controls must scale with: data sensitivity, autonomy of the system, and downstream impact. Without prioritization, every AI system has the same weight as to controls, which in a live environment, is inefficient at best, unrealistic at worst--teams would be forced into checkbox compliance, instead of taking a risk mitigation

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(5/17)

approach. o Note: Contractual requirements are key to focus upon from an enterprise risk perspective, especially as to the preservation of privilege—the lack of which can result in professional sanctions, financial penalties, and reputational harm. The fair use of copyrighted material is already covered within IP already though litigation continues to evolve.

- Prioritization schema introduces what is given most to least importance in Cyber Ai.
- Prioritization lends a roadmap to maturity or as mature as you can get with AI. It is a starting point for orgs to begin work. However, I do understand that this overlay has a very broad audience that may not be able to keep up with shifting priorities.
- Provide a clear and objective structure for organizing and

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(6/17)

managing tasks. They help quickly identify what is most critical and what can be scheduled for a later time, avoiding delays and resource waste.

- apply priorities only where the priority is concrete. Where priorities are variable the considerations should be provided.
- Organizations will be very likely to plan around these, budget around these, etc. It will give them a starting point and take first steps. Other regulatory standards, auditors or Cybersecurity insurance for example may follow in their models and recommendations.
- Challenge: While there are a number of foundational prerequisites for closing the AI gap in this category, any priority ordering will not be the same for all

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(7/17)

communities and scenarios. So either a more contextualized schema (to address the variance above) or elimination of priorities (for simplification and to eliminate the misalignment) seem viable.

- B. The benefit to prioritization is that it focuses efforts for those that are struggling with what to do first. The NIST experts are best positioned based or their risk and threat research to help those with less expertise to determine what is most important. C. Engineers need to be told what to do. They don't like to spend a lot of time doing criticality analysis or risk management.
- 1) For smaller orgs who don't have a fully established cybersecurity or risk management programs, the prioritization schema will

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(8/17)

be helpful. 2)The prioritization schema has to align with the organization's established risk levels, especially if the expectation is that the AI cyber should fold into an existing cybersecurity and risk program. 3) Should the priority levels align with NIST RMF? Management will ask risk staff to align priority levels with what's already existing.

- The benefit is

straightforward, providing adopters with an authoritative basis for allocating resources and planning effort and timelines. The challenges are 1. maintaining the priorities in a rapidly evolving technological landscape, and 2. Ensuring adopters understand the importance of organizational and technological context in utilizing a priority schema

- Prioritization helps organizations figure out

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?

(9/17)

where to start, especially when they are resource-constrained and can't tackle everything at once. Upside is focus. Downside is that lower-priority items tend to get permanently deprioritized once labeled that way.
In critical infrastructure, that's a real problem because an AI system influencing physical operations doesn't

care what priority level a subcategory was assigned. Context matters more than a generic label!

- Benefits is flexibility in approach to schema due to organisational structure
- this helps an org design tiers to address the priorities, set the attention and sequence. to select most important first vs all at once.
- The absence of a prioritization schema will

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(10/17)

create inconsistencies across organizations when applying controls for reciprocity agreements and an extensive variable that may trigger chaos.

- Challenges are with connecting the purpose of the function with the function as its applied. Approach it from how the person accesses the service. ❤️ Everyones been working real hard thank yous!
- B: language that non-tech executives can use to measure AI activities in their C: For tech leaders it is additional work that does not add value, but confusion, as AI expansion of technology, not a stand alone difference from technology
- Challenges of potential rigidity if the schema is too static, difficulty in accurately scoring emerging

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(11/17)

AI risks, and the risk of over-engineering prioritization without sufficient data to support decisions.

- Challenges: implementing organizations have different governance, risk, and compliance postures. Therefore, my priorities do not equal your priorities.
- The benefit is huge. If we consider that we are acknowledging that fundamentally security requirements don't change, but we acknowledge that there might be some things to tackle first that offers actionable guidance.
- It gives agency a starting point to manage risk of AI. They would have a plan to spend resources on those items most critical and then work their way to less critical. Resources are not unlimited to address.

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(12/17)

- Prioritization might force the company focus on hot spots and miss the others areas in low prioritization.
- Challenge: we have no way of providing universally applicable rank ordering for the organizations using the prioritization Benefit: organizations need some help wrapping their heads around the controls Option 2 had a new prioritization level for prerequisites. This may be universally applicable? Priority required/optional?
- No having prioritizing hinders thay to address first. Categories can be useful but their relative value isn't as clear or possibly subjective
- Benefit are there is a (more) clear directive of what needs more (or less) urgency

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(13/17)

- 1. Fixed priorities may be dangerous for privacy. If a subcategory related to "Data Minimization" is labeled "Moderate Priority", but an organisation is processing highly sensitive health data, that "moderate" label might lead to under-resourcing a high-risk privacy area 2. Data protection legislation may mandate certain outcomes that a NIST "Moderate" priority might inadvertently downplay, creating a conflict in "common taxonomy"
- the costs associated of understanding / implementing priorities in the manufacturing supply chain is the biggest challenge - they still struggle with the cybersecurity CMMC level

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?

(14/17)

2 costs for implementation and audit. for small / medium coming up with an easy to use agnostic roadmap of acceptable use of AI applications

- C= challenges created by the prioritization: Moral hazard. Organizations will predictably focus on priority 1 subcategories and deprioritize or ignore those rated 2 or 3, even when the lower-rated subcategories may be critical in their specific threat model, regulatory environment, or deployment context. Accountability displacement. A priority 3 rating from NIST may be used to justify underinvestment in a subcategory that

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?

(15/17)

later proves to be the vector for a breach. Conflict with local/ sectoral regulatory requirements: NIST's priorities may conflict with other regulatory instruments B = Benefits: offer a quick and actionable starting point for organizations

- Applying to the continuous changes in AI tech.
- B: Having a proper way to categorize based on priority and usage C: Executive buy in/ understanding
- Good starting point . Good to say it is starting point as starting reference for different industries
- The biggest benefit of prioritization is that clearly identifies what is important and what is not
- I think the priorities would add benefits, it provides a baseline to a starting point.
- Benefits: starting point

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce?
(16/17)

for resource-constrained orgs. Challenges: dates fast, doesn't reflect sector/maturity, conflicts with other Profiles.

- It helps focus effort on the highest-impact issues and supports consistent, defensible decision-making, but it can oversimplify complex realities if treated as a substitute for judgment or left static over time.
- A challenge of adding one more layer of complication instead of simplifying.
- Priorities changes across space (industry, size, etc) & time (tech/threat evolution)
- N
- B is to align with business goals and risk appetite
- Both pros and cons. Like Change management there should be ongoing assessment and updates.
- Challenges is that since AI

Proposed Priorities: Questions for Discussion

Proposed Priorities (1/3)

0 5 4

What benefits or challenges does having a prioritization schema introduce? (17/17)

is new so it will not be as accurate and will change quickly.

- The benefits of using prioritizations allows us to provide extra protection levels for HVA assets.
- Easier for orgs new to AI or CSF
- I love the examples, but the prioritization is not needed
- Reference and guidance

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?

(1/11)

- - Businesses like simple things, in a checkbox style; opt4 feels a bit more vague, and harder to normalize. - Which business is the audience? If smaller ones, make it easier on them!
- No including priorities is one way good to have the specific area addressed rather than generalizing it. It will be challenging to have multiple areas to be observed/ analyzed when there is no prioritization is in place.
- Challenge would be to not know what should be completed first.
- C : increases difficulty in getting started, where to focus resources first, i.e., “how to eat the elephant”
- B: nothing is automatically not prioritized C: teams likely will

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(2/11)

be forced to prioritize where the funding where go anyways

- B: not including priorities for each subcategory allows organizations flexibility to adapt implementation based on their unique risk posture, industry context and resource constraints without being locked into a perspective hierarchy. It encourages holistic adoption rather than selective compliance with only high priority items. C: without subcategory priorities, organizations may struggle to allocate limited resources effectively leading to inconsistent implementation where critical controls are treated equally with lower impact ones. Teams may experience decisions paralysis delayed remediation of high risk gaps and difficulty justifying investment to leadership without clear prioritisation guidance
- Organisations are very diverse

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(3/11)

and what a priority is for one size or sector is not necessarily the same for others. Perhaps to highlight the concept of risk as risks vary among organisations and perhaps is already understood among more mature organisations. More education is needed about risk for less mature entities.

- What about managing primarily based on roles of these relationships rather than primarily on priorities?
- Beneficial
- B: avoids industry specific conflicts of interest and values , avoids the "one-size-fits-all" assumptions inherent in setting up how one's priorities should fit the framework.
- Without prioritization, every AI system has the same weight as to controls, which in a live environment, is inefficient

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(4/11)

at best, unrealistic at worst--teams would be forced into checkbox compliance, instead of taking a risk mitigation approach.

- Without priorities, I worry about speaking the same language with my sister agencies. My program is audited by another agency so we would be beholden to whatever they decide is a priority.
- B: Supports long-term use as priorities may become dated as AI evolves rapidly C: Lack of clarify that priorities allow to visualize
- the primary audience should be knowledgeable responsible people who are able to make the value-based risk management decisions, rather than slavishly complying. Maybe provide a separate guide for newbies who might need more guidance or to have decisions made for them.
- C - Organizations may prioritize 1 or 2 initially and have

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

048

What benefits or challenges would be introduced by not including priorities for each Subcategory?

(5/11)

significant gaps and risks arise by putting later prioritized items on the back burner. Organizations may budget for Priority 1 or 2 only. C- priority depends on many things - the environment, the rest of an organization's security stack and tools, the industry, the types of resources they are protecting C- this space and it's threats is so rapidly evolving there may be a need to frequently re-prioritize and update

- B: simplification B: eliminate misalignment C: may miss foundational prerequisite
- C. Engineers and developers entire world is about prioritization. Based on resources, threats etc. If NIST experts don't identify prioritization it is a proposition is based on risk management which typically is based upon cost and resources and

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(6/11)

if the threat is likely to occur and with so much unknown about AI then there are large opportunities for mistakes. NIST is best positioned know best. If things change as it develops then change the guidance.

- Smaller orgs or orgs with no established cybersecurity or risk management program, they may not be able to self-prioritize.
- Not including priorities allows for breadth in application and flexibility, but it might be limiting in helping organizations make decisions easily, especially for small-to-medium businesses that do not have the capacity to do rigorous risk assessment from time to time.
- B: Allows organizations to define priorities based on their own deployment context rather than inheriting a generic label that may not reflect their actual risk.

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(7/11)

C: Without any prioritization signal, less mature organizations won't know where to start and may implement subcategories in an order that doesn't match their risk exposure.

- C over what priority to prioritize over the already set guidelines. One would be left rudderless if every business sets it's own priorities to get things done. Especially on the size of the business
- where do we start? what is most relevant? Is risk binary, or are there priorities?
- Not including priorities will create inconsistencies across organizations when applying controls for reciprocity agreements and an extensive variable that may trigger chaos.
- It makes it more difficult to understand connections. I cant really say now if

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(8/11)

I would negate the subcategories its contingent on what the objective might be.

- B: maintain focus on cyber C: lack of priority to leadership, as cyber has been ignored and AI is an attention grabber
- Allows flexibility for organizations to tailor prioritization based on their own risk context, maturity, and use cases. And reduces complexity and avoids false precision where data or consensus is lacking.
- benefits: remove confusion on what to implement. challenges: no guidance on "what's important".
- Challenges - Lacking priorities gives the sense that everything is equally important which is not true.

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(9/11)

With the introduction of AI certain things are accelerated and need to be addressed first.

- Challenge - each agency will have to select their own prioritization so there is no standard or guidance. Each agency may derive priority from different areas.
- People need to watch over all areas and measure the risks they are facing first.
- priorities might be misleading

Option 2 and 3 are essentially the same. numbers versus names? For option 2, the new level, given its definition as "prerequisite" should be priority 0, shouldnt it?

- C- I don't know what to address first in the subcateogy
- Similar to above. Decision making or prioritization. However not including it allows agencies and organizations to define their own priorities
- B= alleviate the risks

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

0 4 8

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(10/11)

abovementionned C= need to include a guidance to help with prioritization

- Difficult to organize and communicate.
- B: More high level understanding C: Less broken down/ subtasks
- Not having priorities may muddle the boundaries
- B: provides guidance on where to start. Not having any reference, it enables sprawl of interpretation.
- Benefits: more durable as AI shifts; composes cleanly with other CSF Profiles; pushes attention to the rationale text, which is more actionable than a number. Challenges: harder for low-maturity orgs to know where to start; risks being read as "everything is equally important"; loses the explicit signal of where AI changes the calculus most.
- Simplifies the framework and avoids false precision, but it can make it

Proposed Priorities: Questions for Discussion

Proposed Priorities (2/3)

048

What benefits or challenges would be introduced by not including priorities for each Subcategory?
(11/11)

harder to consistently compare, sequence, or act on items across teams.

- C. Simpler approach. Each organisation has its own priorities.
- -
- B move faster C higher risks
- More detail is of benefit as it may unveil trends that would otherwise be missed.
- this would help focus on identifying what is Cyber AI
- It may prohibit us from prioritizing the most important assets and AI control requirements for our flagship apps and services
- Harder for orgs new to AI or CSF
- Organizations having to being able to gauge the risk to their organization so they can rank.

Proposed Priorities: Potential Options

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Option 1

No Change

Retain the existing 3-level prioritization schema and refine levels and definitions

Option 2

Add a new prioritization level

Replace the current schema with a 4-level prioritization

Option 3

Replace Schema

Reflect the approximate relative degree of differences between "traditional" cybersecurity considerations and AI-specific cybersecurity considerations

Option 4

Remove Priorities

Do not include a prioritization schema and therefore no proposed Subcategory priorities

NIST's preferred option

Slido.com
#CyberAI_Spring2026-1

Proposed Priorities: Potential Options

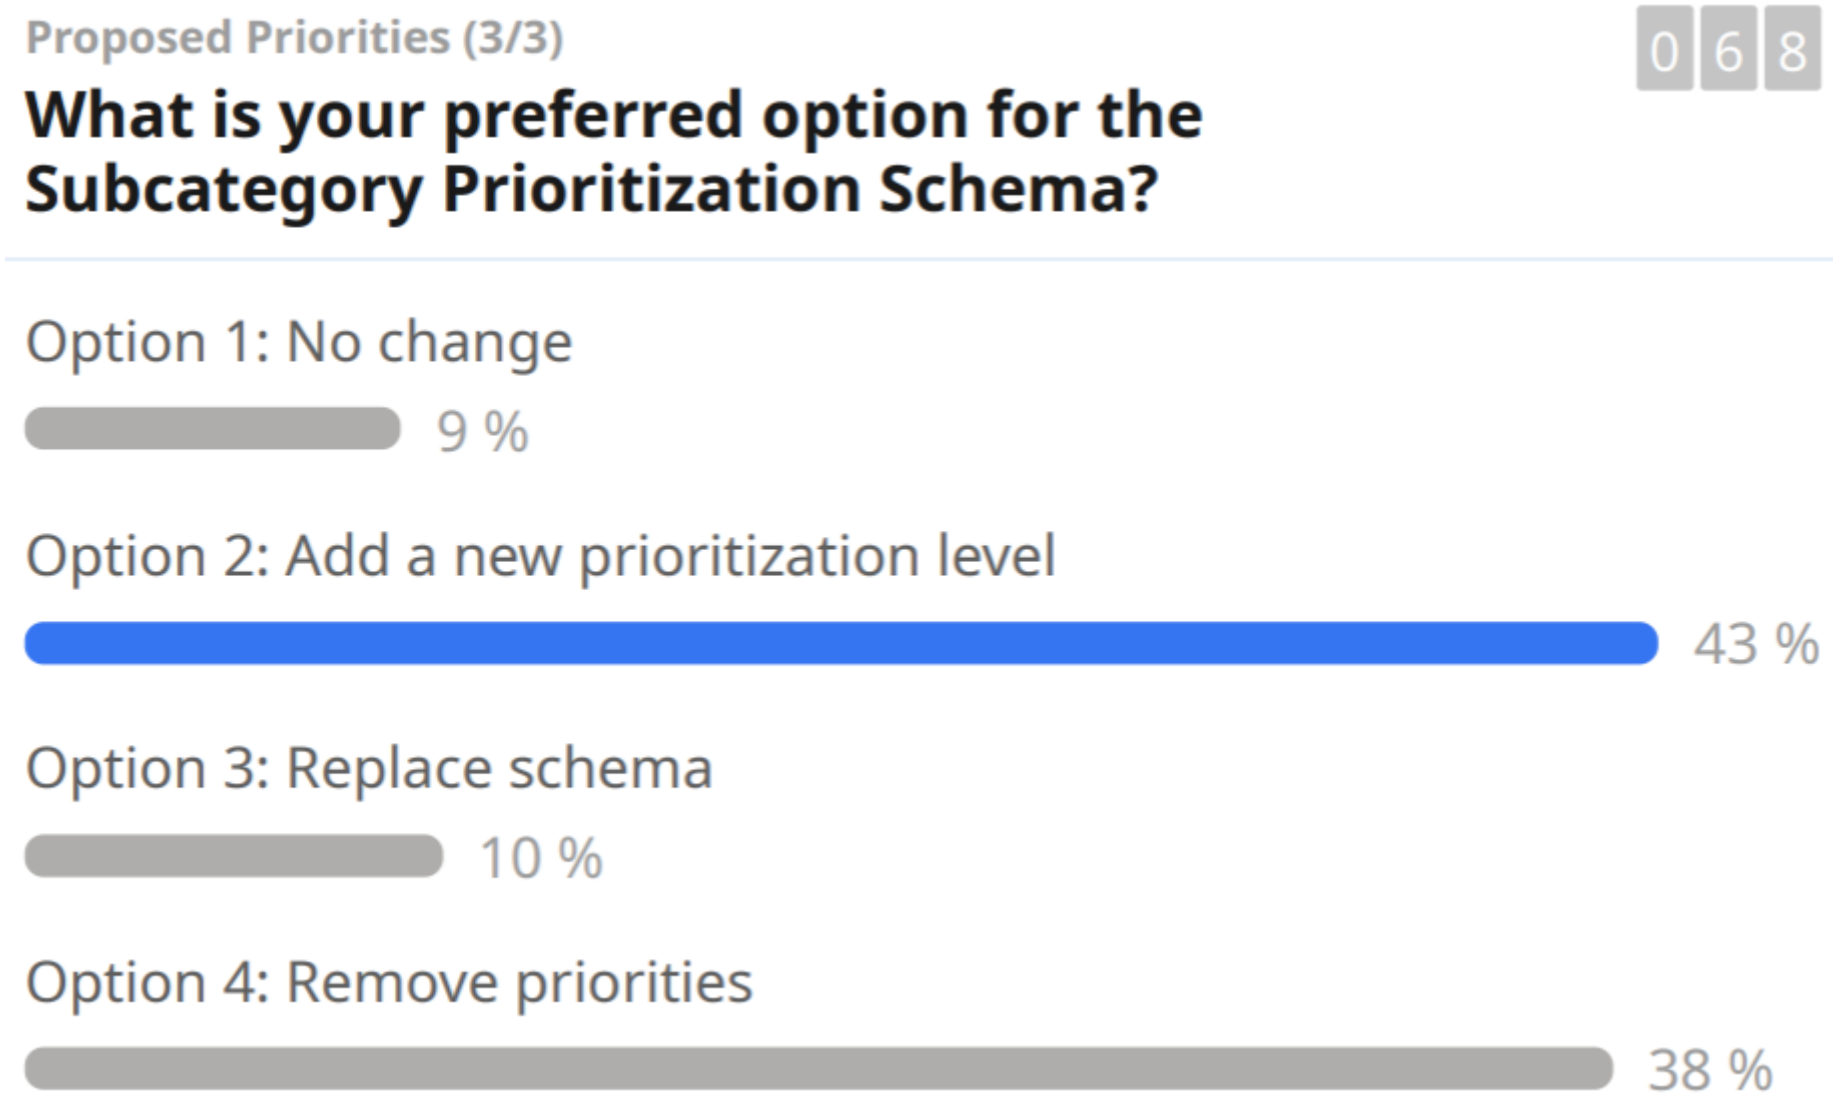

Proposed Priorities: Potential Options

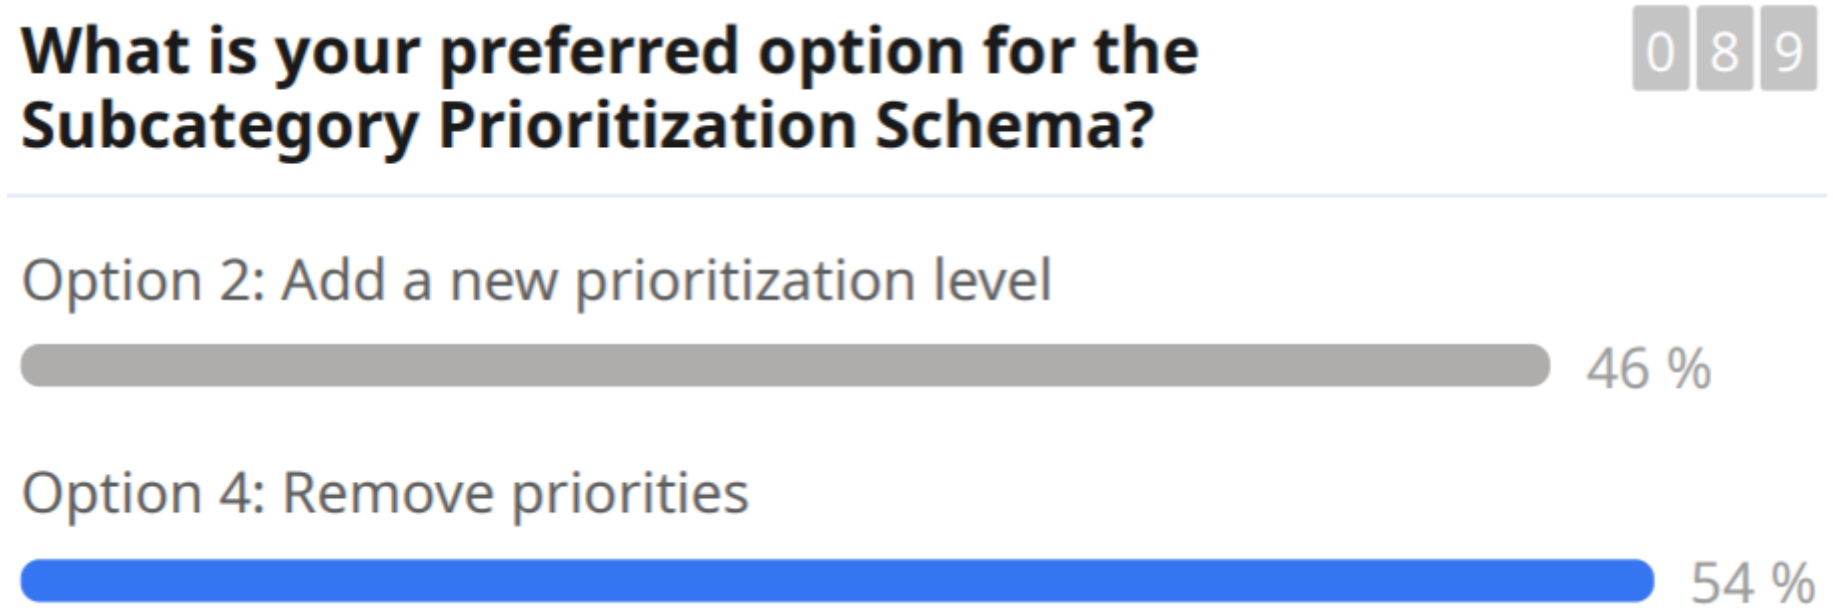

Proposed Priorities Discussion Part 2

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Proposed Prioritization Level Definitions: Potential Options

Level	Option 1: No change	Option #2: Revise Prioritization Levels	Terminology Options
Highest Priority	• Most critical to address the challenges for a Focus Area • Should be addressed immediately given available resources • *Current level terms: "1" / High Priority*	• Most critical Subcategories to enable a Focus Area • Should be addressed immediately given available resources	• Very High • Advanced • Elevated • Surpasses • High
Middle Priority	• Next priority after implementing High Priority • *Current level terms: "2" / Moderate Priority*	• Has a dependency on the Subcategory, but it is not as critical as High Priority • Next priority after implementing High Priority • May become a higher priority in certain contexts or environments	• Improve • Moderate • Medium
Lowest Priority	• Generally important to the Focus Area but may not require the same level of urgency as higher priorities • "Foundational" does not equate to low priority. All Subcategories should receive consideration • *Current level terms: "3" / Foundational*	• Generally important to a Focus Area but may not require the same level of urgency as higher priority Subcategories • Lowest priority does not equate to no priority. All CSF Subcategories should receive consideration by an organization	• Foundational • Supporting • Average

Proposed Prioritization Level Definitions: Potential Options

Level	Option 1: No change	Option 2: Revise Prioritization Levels
Highest Priority	• Most critical to address the challenges for a Focus Area • Should be addressed immediately given available resources • *Current level terms: "1" / High Priority*	• Most critical Subcategories to enable a Focus Area • Should be addressed immediately given available resources
Middle Priority	• Next priority after implementing High Priority • *Current level terms: "2" / Moderate Priority*	• Has a dependency on the Subcategory, but it is not as critical as High Priority • Next priority after implementing High Priority • May become a higher priority in certain contexts or environments
Lowest Priority	• Generally important to the Focus Area but may not require the same level of urgency as higher priorities • "Foundational" does not equate to low priority. All Subcategories should receive consideration • *Current level terms: "3" / Foundational*	• Generally important to a Focus Area but may not require the same level of urgency as higher priority Subcategories • Lowest priority does not equate to no priority. All CSF Subcategories should receive consideration by an organization.

Slido.com
#CyberAI_Spring2026-1

Proposed Prioritization Level Definitions: Potential Options

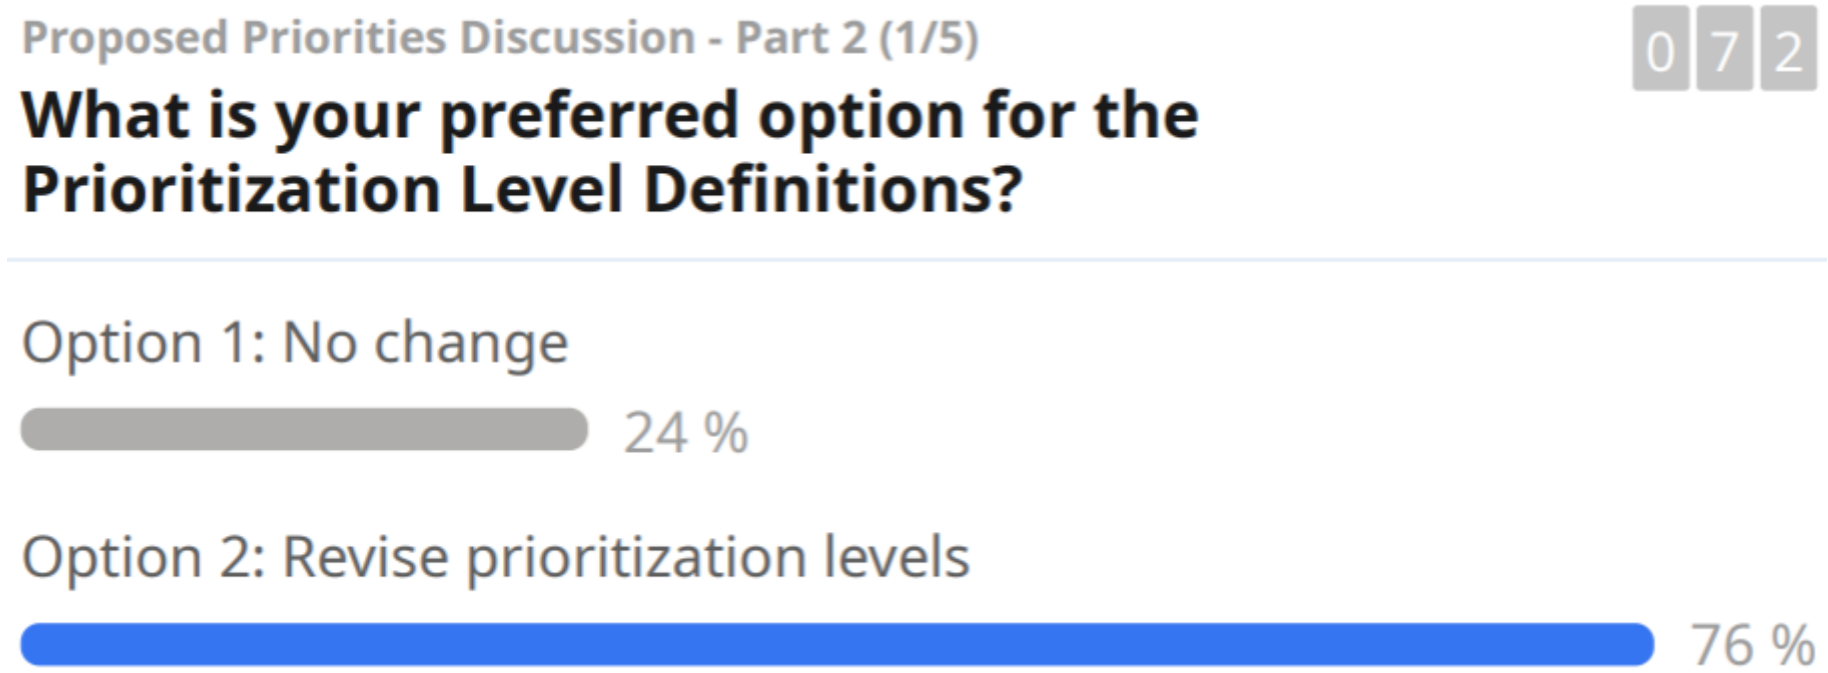

Proposed Prioritization Level Terms: Potential Options

Level	Option 2: Revise Prioritization Levels	Terminology Options
Highest Priority	• Most critical Subcategories to enable a Focus Area • Should be addressed immediately given available resources	• Very High • Advanced • Elevated • Surpasses • High
Middle Priority	• Has a dependency on the Subcategory, but it is not as critical as High Priority • Next priority after implementing High Priority • May become a higher priority in certain contexts or environments	• Improve • Moderate • Medium
Lowest Priority	• Generally important to a Focus Area but may not require the same level of urgency as higher priority Subcategories • Lowest priority does not equate to no priority. All CSF Subcategories should receive consideration by an organization.	• Foundational • Supporting • Average

Slido.com
#CyberAI_Spring2026-1

Proposed Prioritization Level Terms: Potential Options

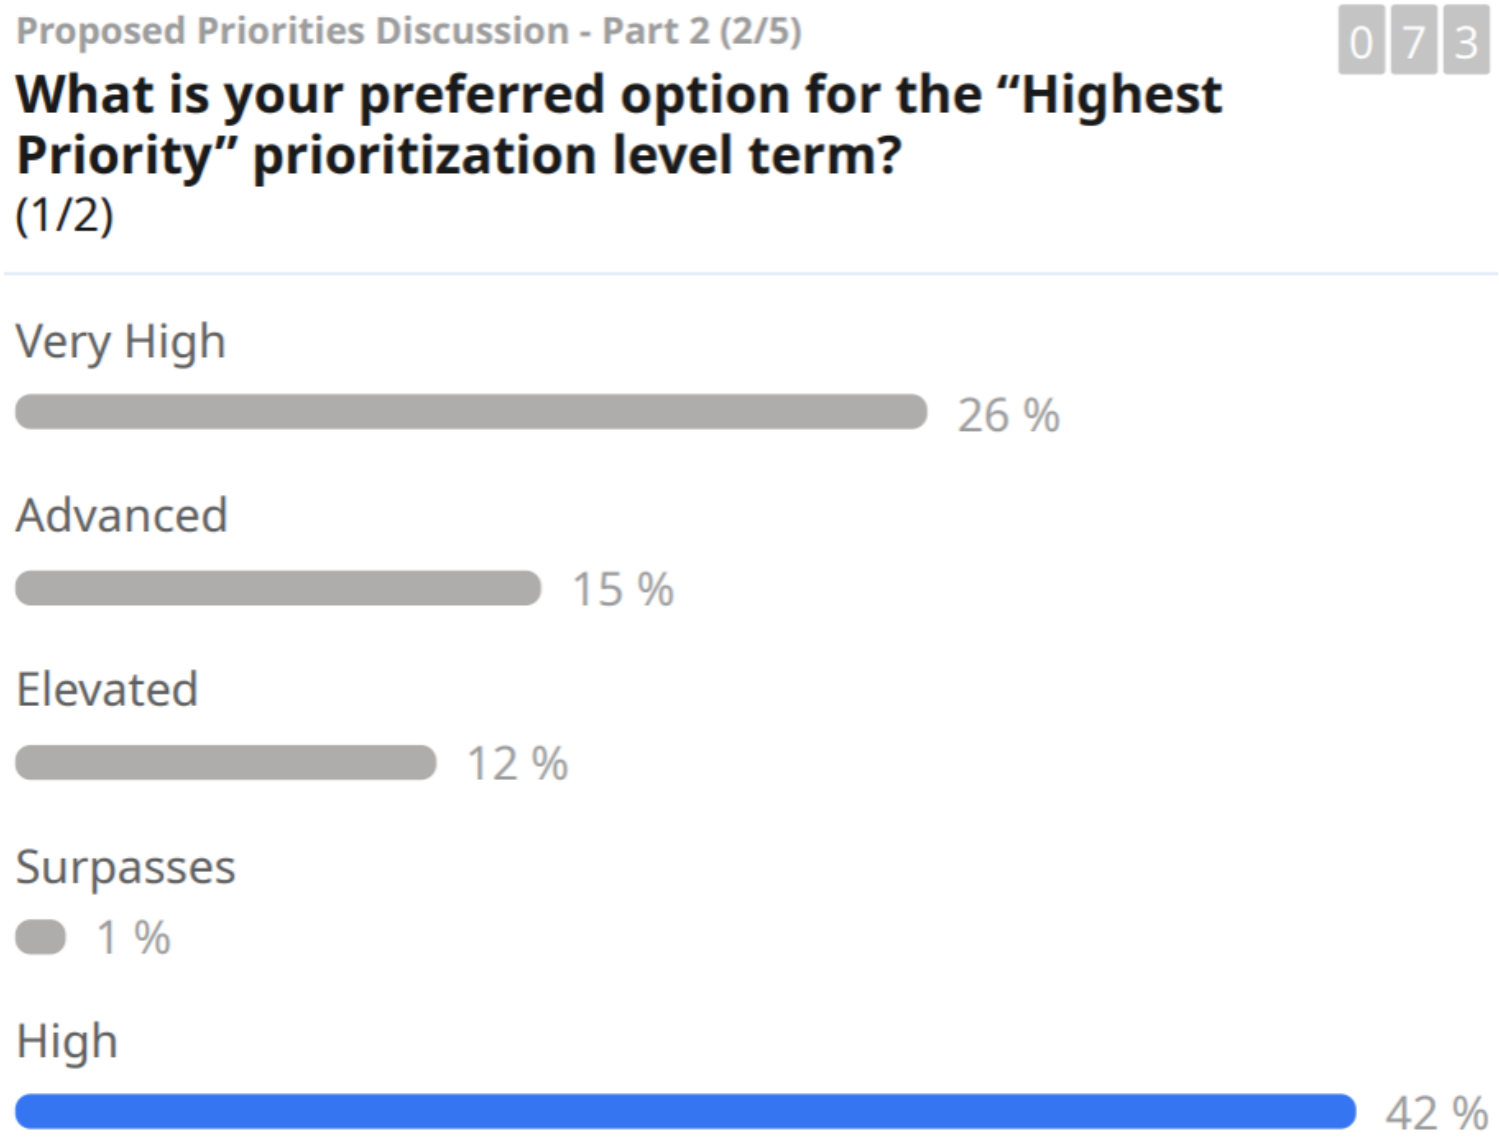

Proposed Prioritization Level Terms: Potential Options

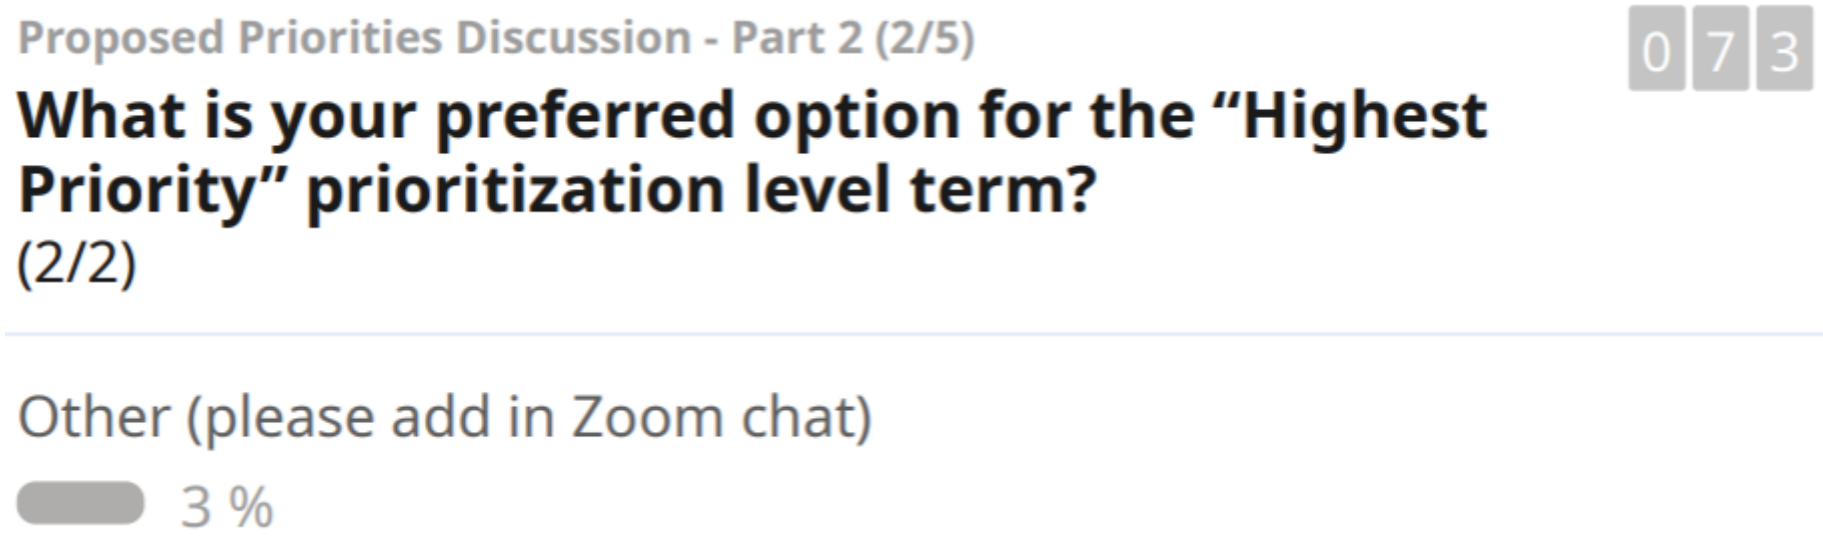

Proposed Prioritization Level Terms: Potential Options

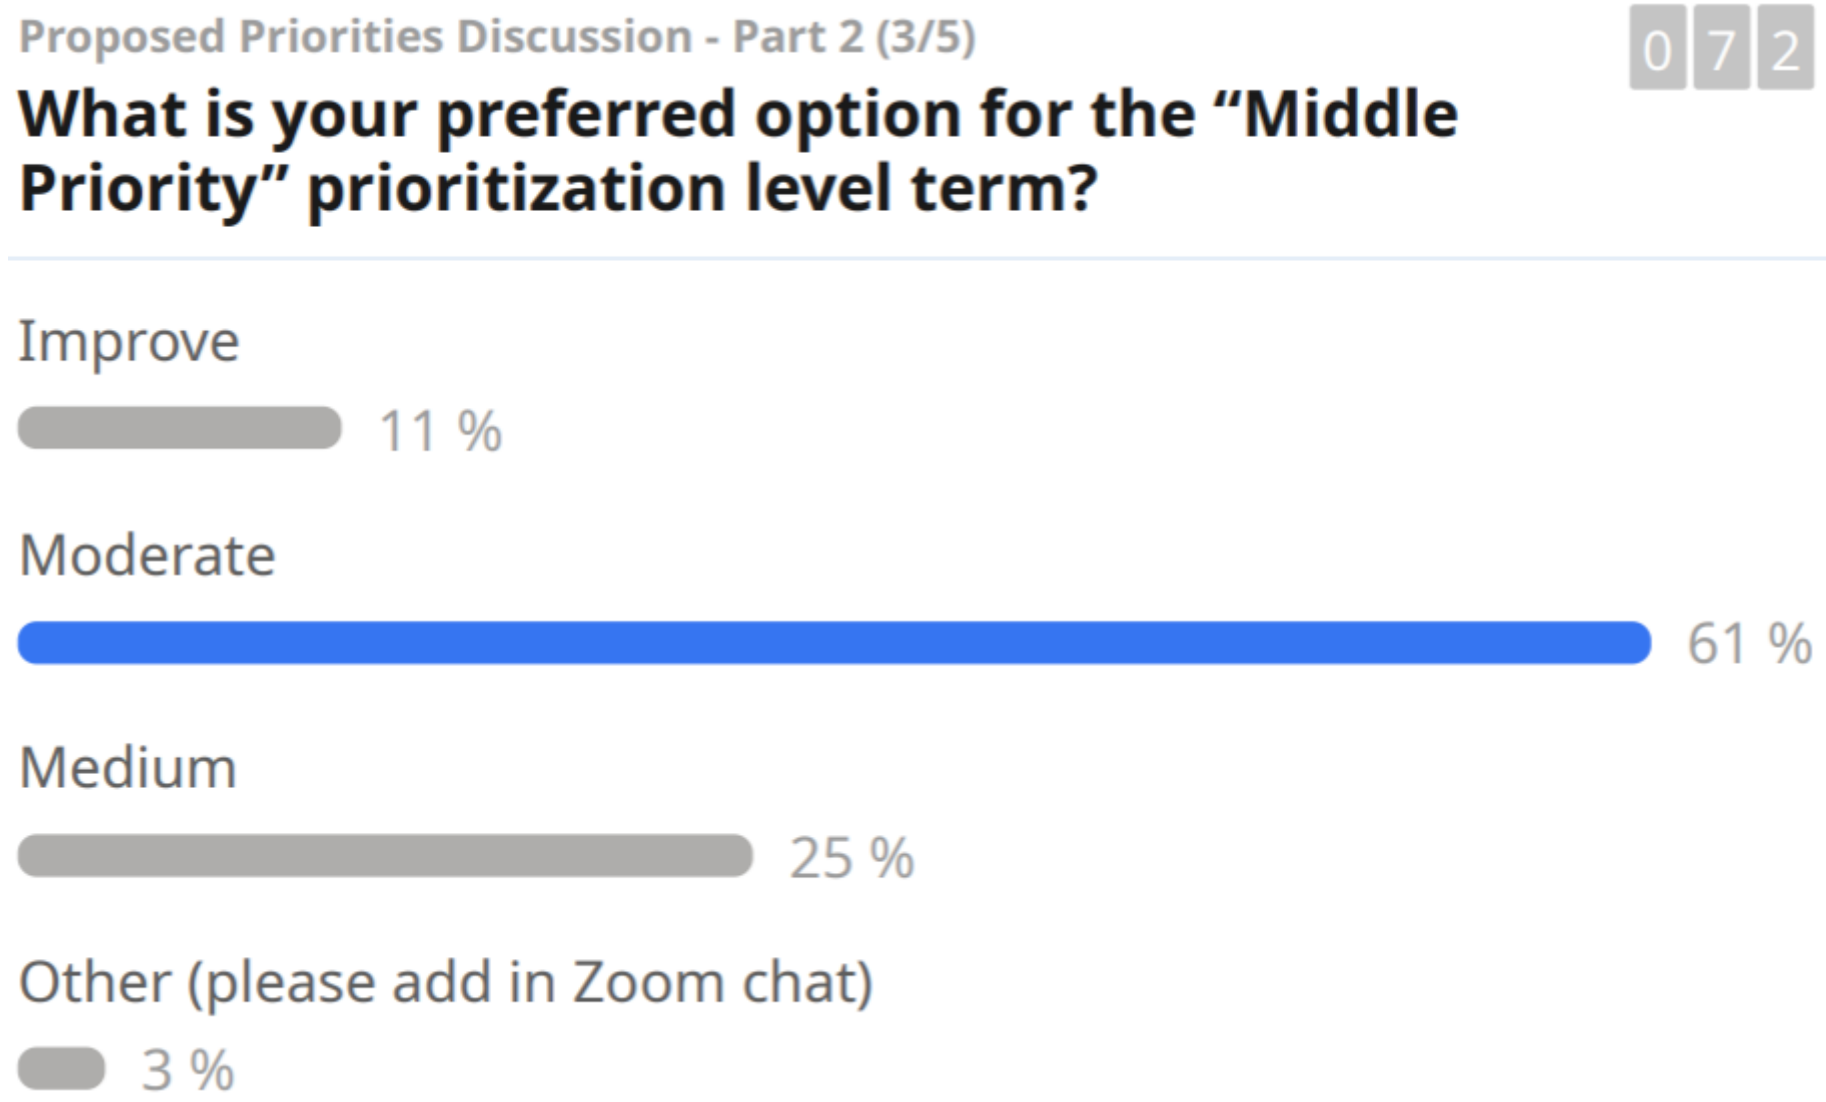

Proposed Prioritization Level Terms: Potential Options

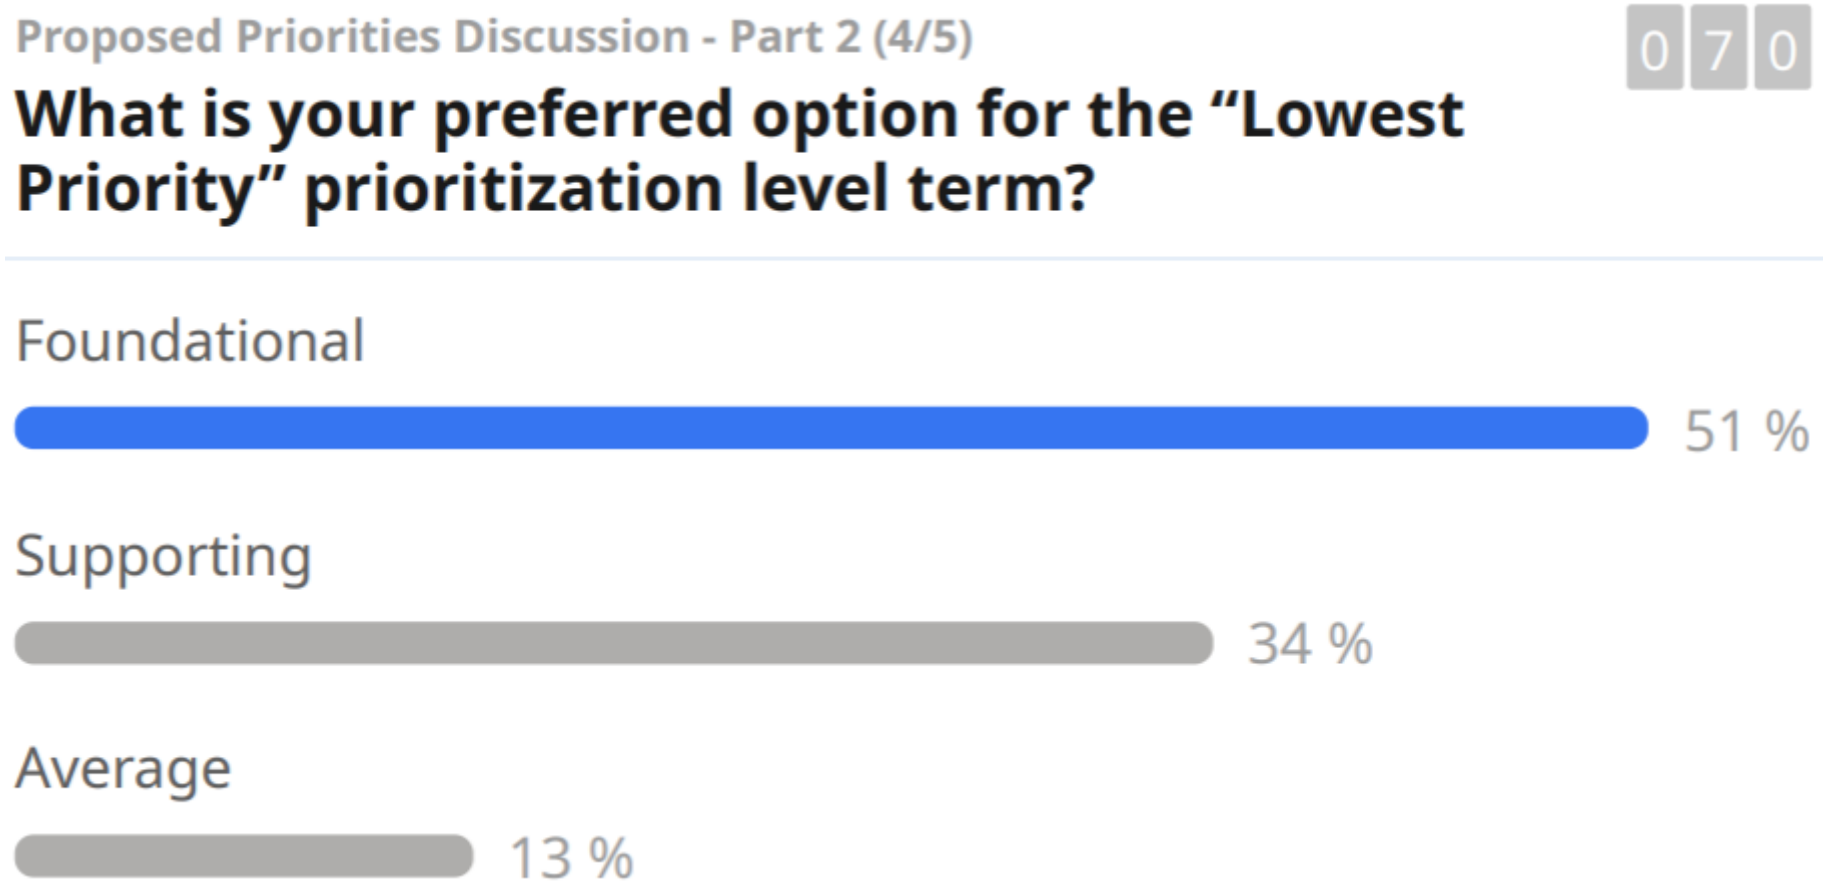

Proposed Prioritization Level Terms: Potential Options

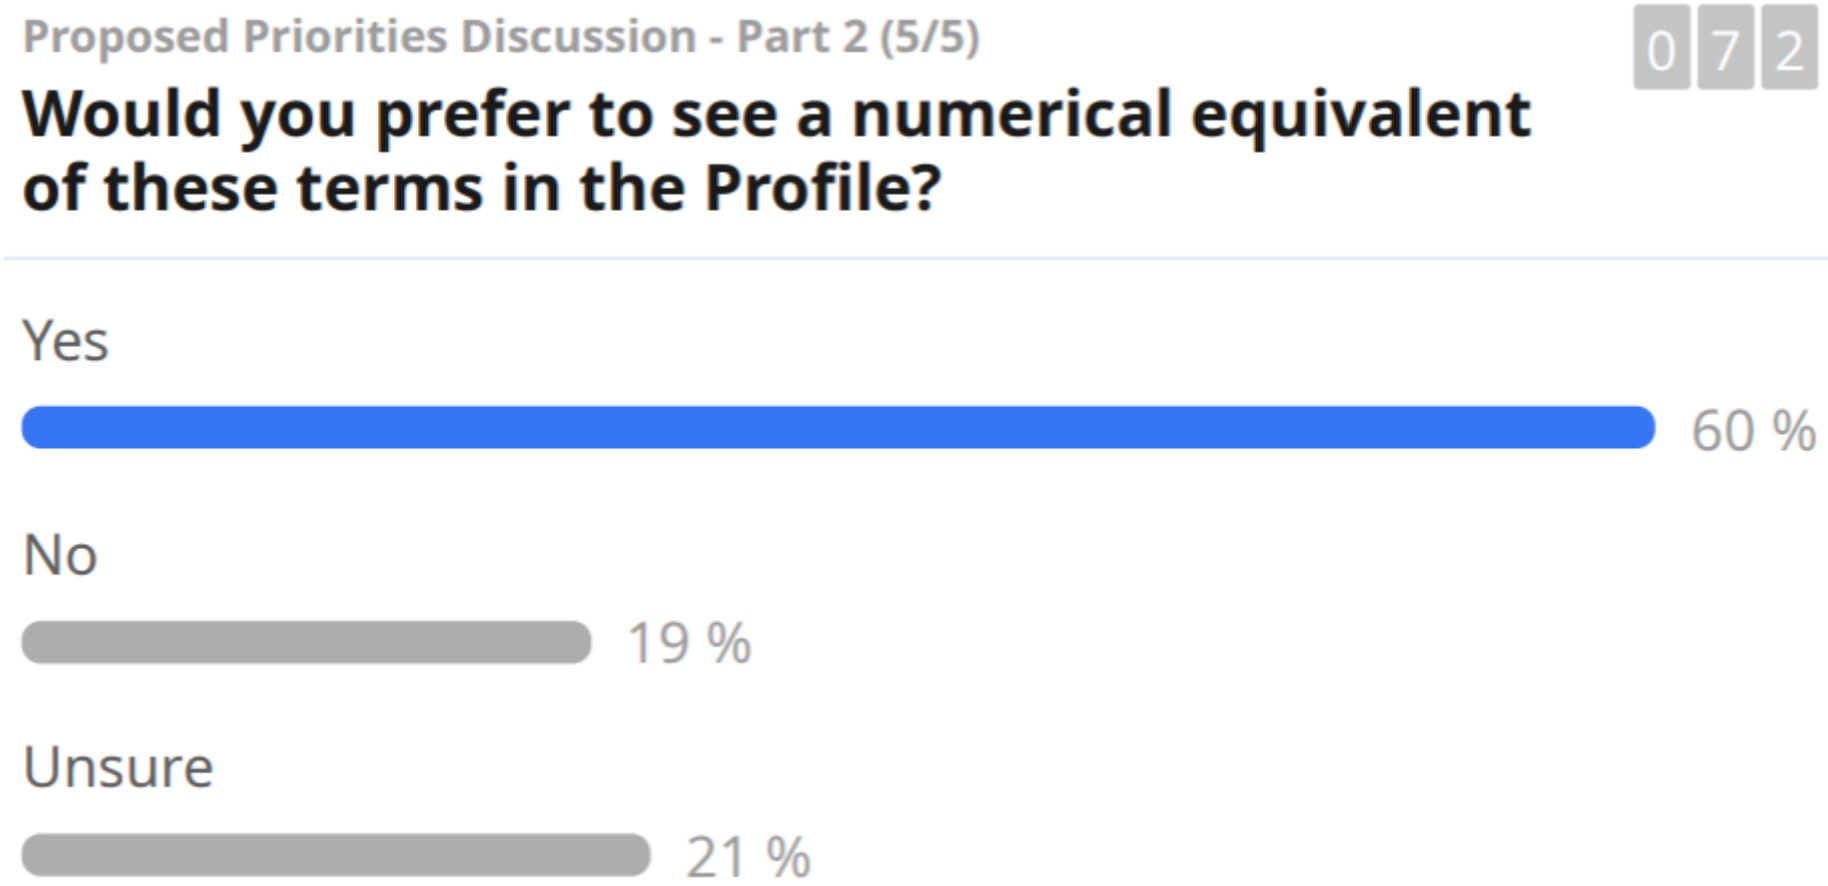

“Standard Practices” Discussion

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

“Standard Practices”: Summary of Feedback

- Ambiguous in meaning
- Should focus on AI-specific considerations
- May conflict with existing regulatory requirements or already be differently defined by individual organizations

- A new phrase will enable the Profile to allow for innovation
- New phrasing should be linked to cybersecurity best practices or acknowledge that organizations may already be using them
- New phrasing should address that organizations should still ensure that established cybersecurity practices are applied to AI

“Standard Practices”: Potential Definition Options

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Option 1: No change

Indicates that:

- There are no unique considerations identified for the Focus Area
- The activities of the cybersecurity program are sufficient for AI systems

(Refer to Section 2.2, page 13)

Option 2: New definition A

Indicates that:

- No AI-specific considerations 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.

Option 3: New definition B

Indicates that:

- No AI-specific considerations for the Focus Area were identified
- The activities of the cybersecurity program are sufficient for AI systems and within the context of acceptable risk.

NIST’s preferred option

*Potential statement options will be discussed shortly

“Standard Practices”: Potential Definition Options

Option 1: No change

- Leave definition as is, which reads, “When the phrase ‘[placeholder for phrase]’ is used, it indicates 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.”

“Standard Practices”: Potential Definition Options

Option 2: New definition A

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.”

“Standard Practices”: Potential Definition Options

Option 3: New definition B

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 option

“Standard Practices”: Potential Definition Options

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Option 1: No change

Indicates that:

- There are no unique considerations identified for the Focus Area
- The activities of the cybersecurity program are sufficient for AI systems

(Cyber AI Profile Section 2.2, page 13)

Option 2: New definition A

Indicates that:

- No AI-specific considerations 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.

Option 3: New definition B

Indicates that:

- No AI-specific considerations for the Focus Area were identified
- The activities of the cybersecurity program are sufficient for AI systems and within the context of acceptable risk.

NIST’s preferred option

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?

Slido.com
#CyberAI_Spring2026-1

“Standard Practices”: Potential Definition Options

"Standard Practices" Discussion (1/6)

027

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
(1/6)

- No
- Why would we not include AI-specific? It is just another set, like cybersecurity. And then you have the AND set
- Option 3 is widest ranging.
- Really good question will yous please revisit that question later?
- Do we assume 1 degree is separation for “AI Systems”? What about upstream or downstream AI systems (3 degrees of separation) ? Included or not?
- A
- New Definition A - for the current AI threats and attacks. 3 is closer and 1 may slightly differ in comparison to the risk
- No, if this is coupled to NIST 800-53 Overlay mapping
- The only difference I see is the context of acceptable

“Standard Practices”: Potential Definition Options

"Standard Practices" Discussion (1/6)

027

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
(2/6)

risk. This could cause differences between 1 and 2/3.

- I can imagine somebody puzzling over option 2, with its attempt to list out AI Integrations, etc... and deciding that there is a distinction being made when in fact, the simpler phrase conveys that standard safeguards should apply, period.
- I feel as option 3 could offer a false sense of security for executives. If x is done we are secure. No more need for investment in this area.
- Low medium high, colour code. Focus on the business requirement. What requires mitigation and medium no issues for now etc.,
- No
- I have to choose NIST recommendation

“Standard Practices”: Potential Definition Options

"Standard Practices" Discussion (1/6)

0 2 7

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
(3/6)

as it is most accurate

- Yes. An organization could apply standard cybersecurity safeguards to AI components as Option 2 requires, but still fail Option 3 if those safeguards don’t produce verifiable and actionable outputs within the operational context. In critical infrastructure, this gap is real. Controls get applied but AI outputs influencing physical operations are never validated against operational risk thresholds. Option 3 is stronger because it closes that gap.
- Not sure a situation will arise
- The level of IT experience with technology. Less experienced have more challenges applying existing terms and stds

“Standard Practices”: Potential Definition Options

"Standard Practices" Discussion (1/6)

0 2 7

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
(4/6)

- When considering the usage of AI as/with cybersecurity tooling, since this is new it’s difficult to see how “standard practices” could be applicable
- Yes, Option 2 sets a higher bar than Options 1 or 3.
- I like applicable cybersecurity safeguards language because it makes it clear you need to do something beyond just status quo which everyone may not be mature enough to be doing. Perhaps merge option 2 and option 3
- Option 3 - Relying on "standard" practices can obscure where AI introduces new harms, such as algorithmic bias or unauthorized re-identification. NIST's preferred replacement-"No AI-specific

“Standard Practices”: Potential Definition Options

"Standard Practices" Discussion (1/6)

0 2 7

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
(5/6)

cybersecurity program considerations were identified" - is more precise because it admits that while the security is standard the context of the data remains critical.

- Language that implies AI-exempt is concerning.
- It depends on what the definition of unique is. Yes something could be unique but not AI specific. Also I have concerns about “acceptable risk”
- None of the stated choices adds value onthe AI TOPIC. Note today's evidence based performance in cybersecurity.
- I'm leaning toward Option 2 (New definition A) with a caveat that NIST will provide guidance and also create an AI overlay

“Standard Practices”: Potential Definition Options

"Standard Practices" Discussion (1/6)

0 2 7

Looking at the definitions, is there a situation that would meet one of these definitions but not the other?
(6/6)

so that any system using AI components can be easily identified, assessed and secured based on its sensitivity and prioritization level.

- No
- Until the process/methodology matures, a specific category should be chosen,

“Standard Practices”: Questions for Discussion

- What contextual differences may influence how this phrase is applied for the same Subcategory across the three Focus Areas?

- Is the intent / value behind including such a phrase helpful or does its inclusion possibly create confusion related to distinct requirements posed by AI systems?

- What language does the phrase need to characterize the difference that AI systems introduce in comparison to non-AI systems?

Slido.com
#CyberAI_Spring2026-1

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (2/6)

0 2 1

What contextual differences may influence how this phrase is applied for the same Subcategory across the three Focus Areas (Secure, Defend, Thwart)?

(1/4)

- secure
- Secure seems like suitable in this context of cybersecurity.
- Really good question. Cross ref with list of industries.
- Secure
- Secure
- Each focus area has a different context. So there needs to be context added for each focus area
- The definition of acceptable risk.
- Secure is ambiguous it gives a sense of security and safety but I'm unclear on outcome as a result of an incident Defend is also ambiguous on outcome Thwart is very clear and gives a a sense of
- Thwart
- Kiss. Keep sweet and simple large audience. Data soverienity

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (2/6)

0 2 1

What contextual differences may influence how this phrase is applied for the same Subcategory across the three Focus Areas (Secure, Defend, Thwart)?
(2/4)

across countries and data. GDPR attracts enormous fines.

- Defend seems like suitable in cybersecurity context
- I believe that new definitions must add new concepts and outcomes
- The phrase means different things across the three focus areas. In Secure, verification is relatively straightforward around asset protection. In Defend, actionability depends on whether AI outputs can trigger response within operational timelines.3 In Thwart, verifiability requires outputs to remain trustworthy under adversarial conditions, which is a materially higher bar than the other two.
- Defend because that's the

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (2/6)

0 2 1

What contextual differences may influence how this phrase is applied for the same Subcategory across the three Focus Areas (Secure, Defend, Thwart)?
(3/4)

typical cybersecurity posture

- If internal attack vs extetnal, sone have challenges with linguisticts of what the person did vs technology.
- A Subcategory may be "standard" for Thwart but require AI-specific tailoring for Secure. Defend sits in between.

Sector regulation and the in-house-vs.- 3rd party AI sourcing model also shift what counts as "standard."

- Secure - focuses on the "supply chain" of data. So for privacy, it is about ensuring training data was legally

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (2/6)

0 2 1

What contextual differences may influence how this phrase is applied for the same Subcategory across the three Focus Areas (Secure, Defend, Thwart)?

(4/4)

obtained (PIPA compliance). Defend - focuses on protecting the system while in use. So in a privacy scenario, "Standard practices" might apply to the server, but not to the "prompt" risks that could leak citizen data. Thwart - focuses on stopping those using AI for harm. This should protect the public from AI-driven phishing or identity theft.

- secure
- There needs to be further detailed discussion.
- Agree with current implementation
- None

“Standard Practices”: Questions for Discussion

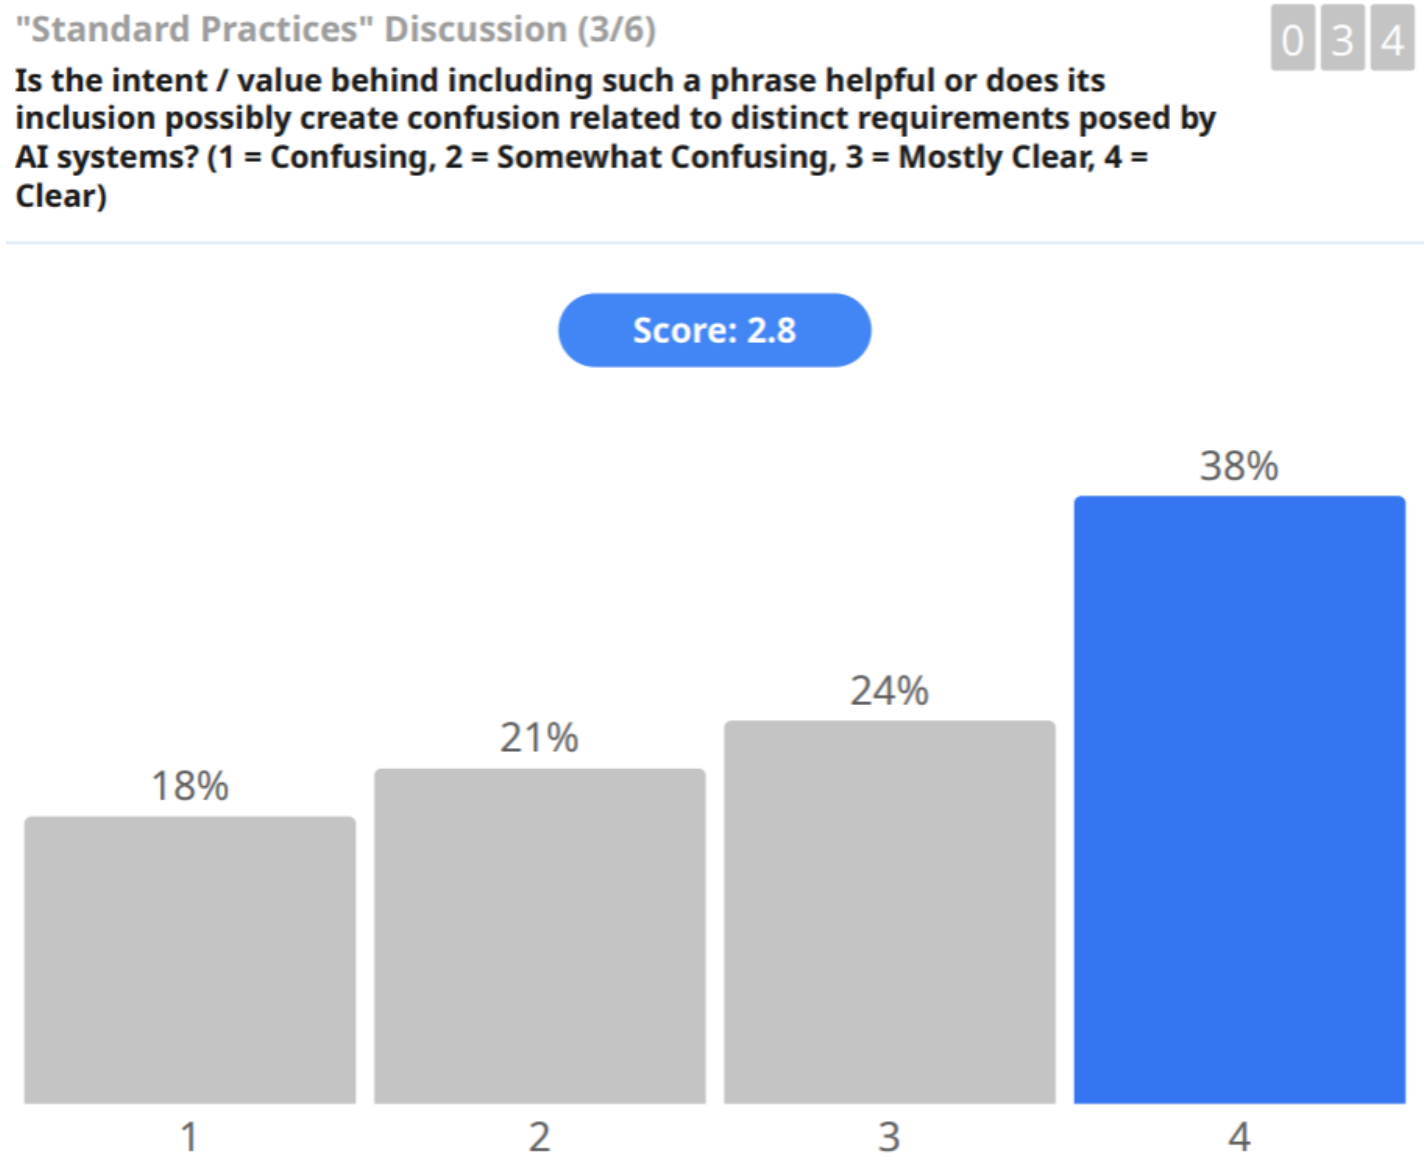

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (4/6)

0 2 3

What language can we use in this phrase to characterize the differences that AI systems introduce in comparison to non-AI systems? (1/4)

- “Focus Area” os a good expression to indicate exactly what you intend to communicate.
- As AI can hallucinate, and thus for audit purposes (the results of which are to be accurate and verifiable) there should be some wording about where AI was used in the auditing process so hallucinations are to be minimised etc.
- cyber - eventually ai cyber will be integrated within it
- Take time on this youre doing excellent its vast and intricate.
- Ai systems vs, non-Ai systems.
- AI needs traditional and current practices plus the AI specific security practices to ensure a comprehensive level of security practices are in place.
- I agree with the no specific AI Tailoring identified
- AI - Specific Industry Agnostic

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (4/6)

0 2 3

What language can we use in this phrase to characterize the differences that AI systems introduce in comparison to non-AI systems? (2/4)

- AI-System Risk
- Cx systems are complex, containers etc., third party. Immutable backups, zero trust networks, informed by policy bcm, board etc otherwise they can go out of business
- Concise and objective, making distinction clearly in text
- Secure4
- AI systems introduce behavioral variability that non-AI systems don’t. A non-AI system behaves deterministically within its programmed logic. An AI system can shift outputs over time through model drift, data distribution changes, or adversarial inputs without any change to code or configuration. Language should capture that: something like “dynamic behavioral risk” or “runtime variability.”

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (4/6)

0 2 3

What language can we use in this phrase to characterize the differences that AI systems introduce in comparison to non-AI systems?
(3/4)

- The language of non-human consideration to given AI systems in terms of output or ML training influenced system
- AI id data, math and code, which is wjy the definition is hard for folks to understand
- Replace "Standard Practices" with "No AI-specific tailoring required" or "AI-specific considerations: none identified."
- - the phrase should explicitly state that traditional practices must be applied to AI components - Explicit language ensures that the "verifiable and actionable" nature of security remains intact for audit purposes.
- Best practices irrespective of industry.
- What about using DAMA definitions? Many orgs are already using

“Standard Practices”: Questions for Discussion

"Standard Practices" Discussion (4/6)

0 2 3

What language can we use in this phrase to characterize the differences that AI systems introduce in comparison to non-AI systems? (4/4)

this and continuity for the definitions would help with consistency.

- future state - it will be part of cyber so doesnt have to be separate term
- TBD - None of the below options is acceptable.
- AI vs. Non-AI as the system category
- Ai -specific

“Standard Practices”: Potential Definition Options

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Option 1: No change

Indicates that:

- There are no unique considerations identified for the Focus Area
- The activities of the cybersecurity program are sufficient for AI systems

(Cyber AI Profile Section 2.2, page 13)

Option 2: New definition A

Indicates that:

- No AI-specific considerations 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.

Option 3: New definition B

Indicates that:

- No AI-specific considerations for the Focus Area were identified
- The activities of the cybersecurity program are sufficient for AI systems and within the context of acceptable risk.

NIST’s preferred option

Slido.com
#CyberAI_Spring2026-1

“Standard Practices”: Potential Definition Options

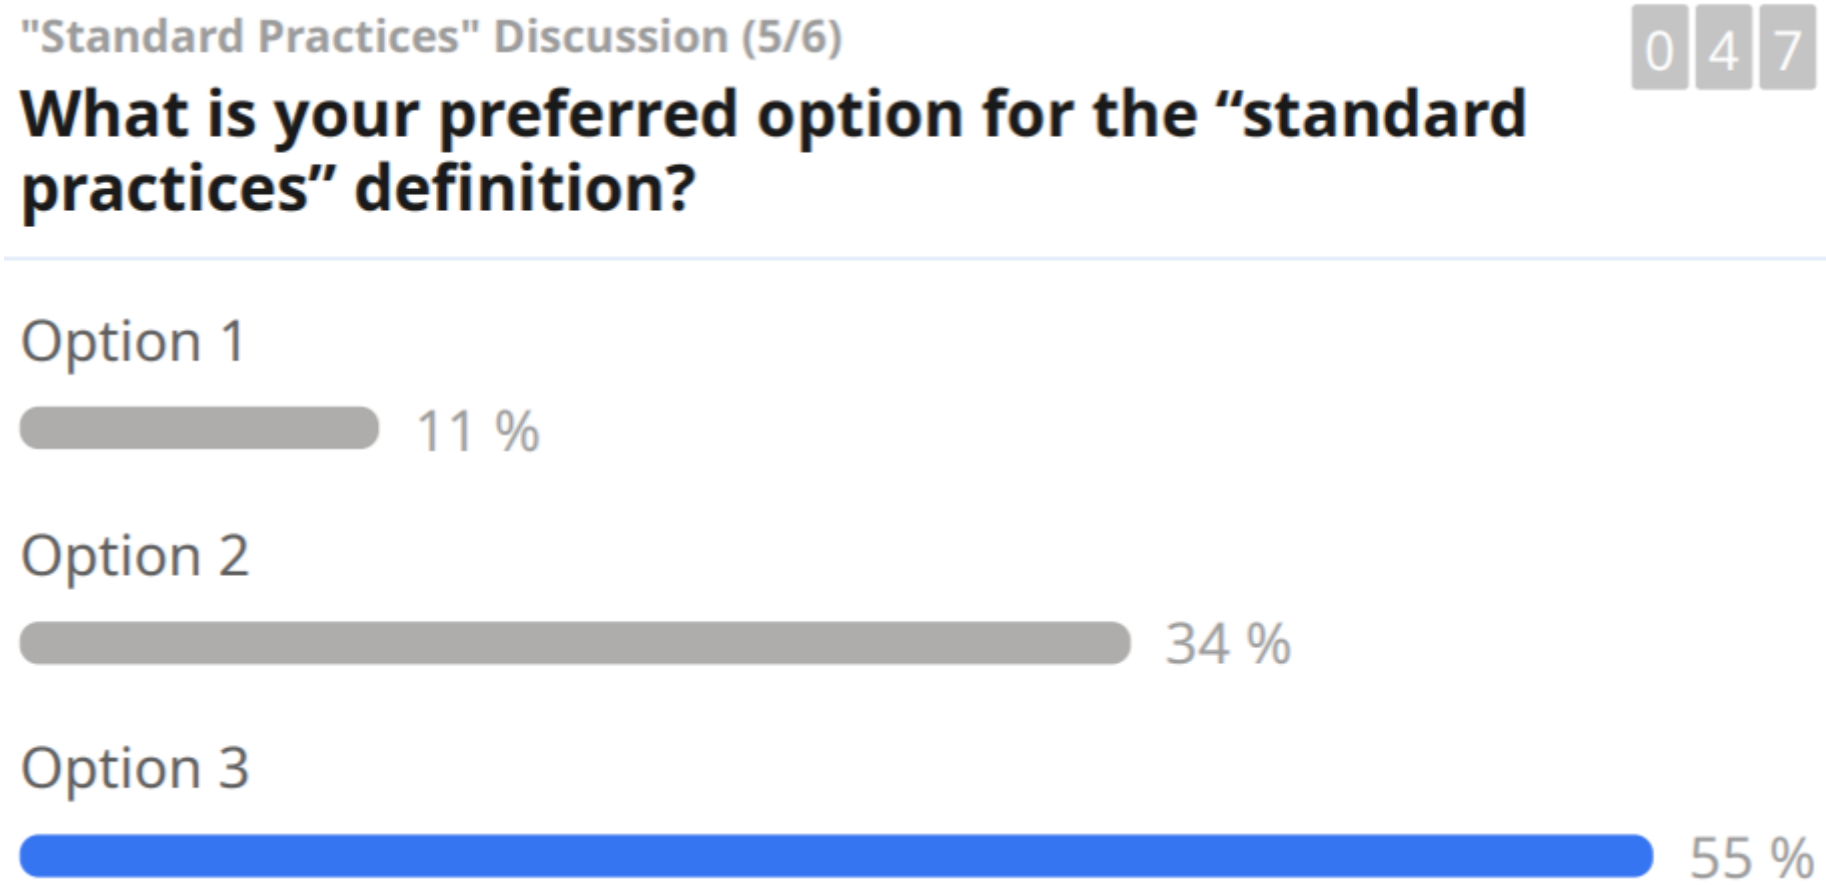

“Standard Practices”: Potential Phrase Options

#1	“Standard cybersecurity practices apply.” (current phrase)
#2	“No AI-specific related considerations were identified.”
#3	“No AI-specific cybersecurity program considerations were identified.”

NIST’s preferred option

“Standard Practices”: Potential Phrase Options

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

#1 “Standard cybersecurity practices apply.” (current phrase)

#2 “No AI-specific related considerations were identified.”

#3 “No AI-specific cybersecurity program considerations were identified.”

NIST’s preferred option

Slido.com
#CyberAI_Spring2026-1

“Standard Practices”: Potential Phrase Options

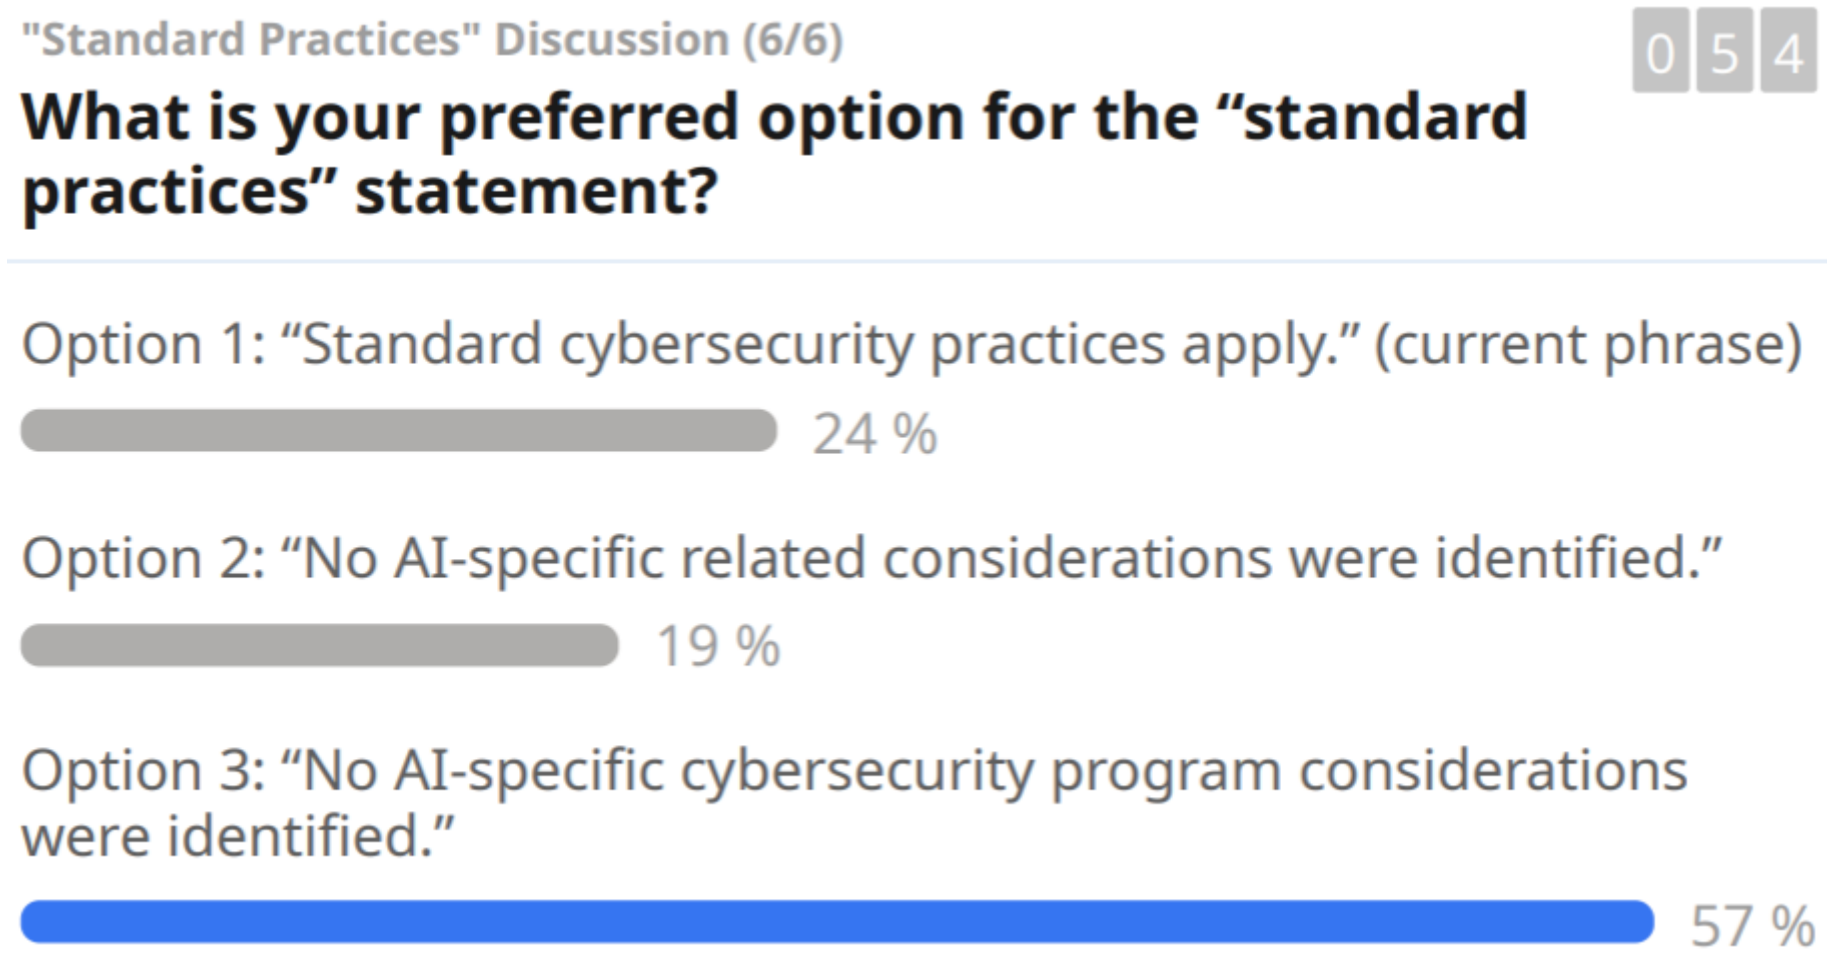

Open Discussion

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Questions for Open Discussion

- **Are there any other ideas or comments related to other Profile elements (i.e., rationales, considerations, opportunities, and Informative References) that you would like to share?**

Topics for future discussion:

- *Working Session #2 Focus: Extending the Technical Content (May 5)*
- *Working Session #3 Focus: Roles and Profile Delivery Formats (May 12)*

Questions for Open Discussion

Are there any other topics related to Profile elements that we should consider?

0 0 2

intellectual property

Close-out

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Next Steps After COI Working Sessions

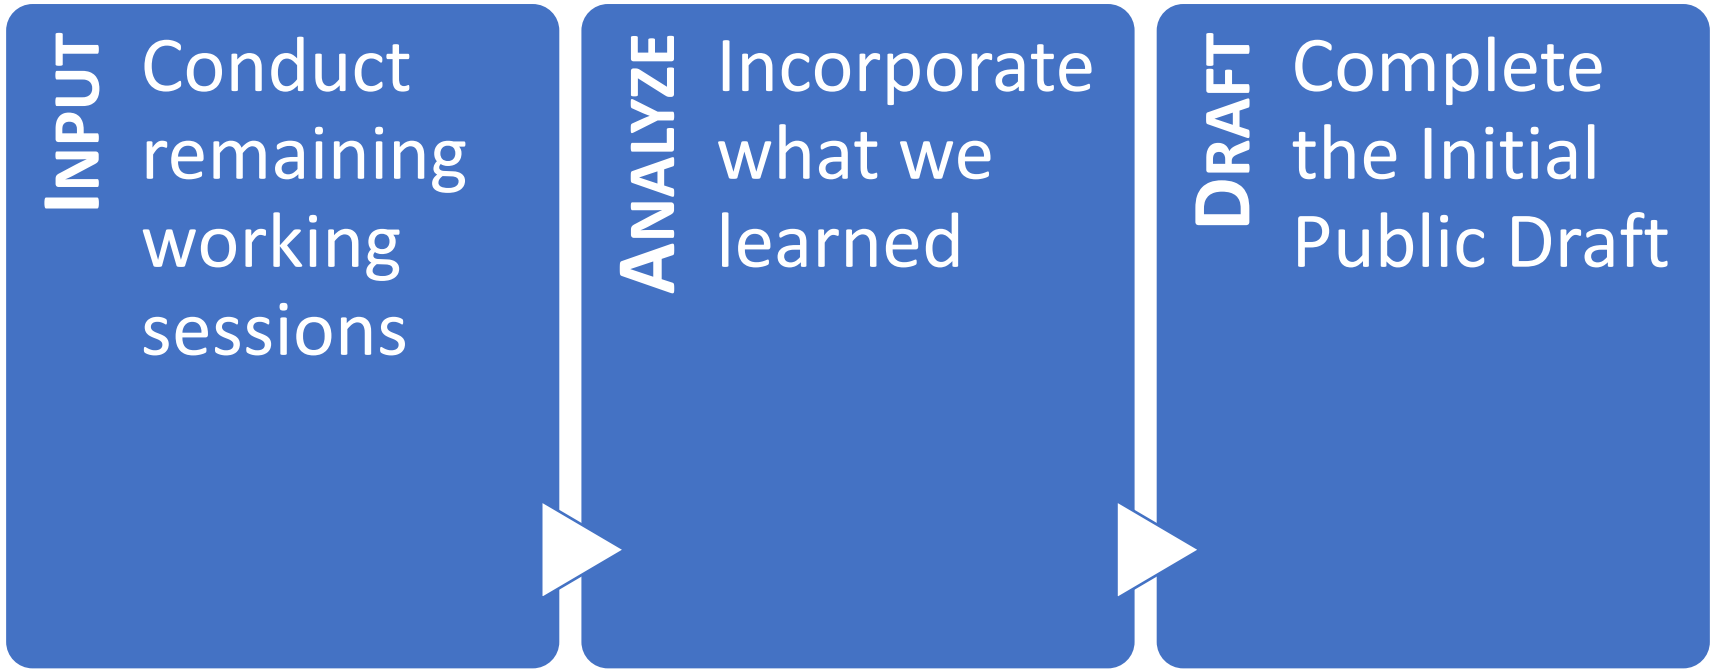

If we missed your input today, please feel free to email us: CyberAIProfile@nist.gov! Please send your inputs by May 15, 2026.

Working Session Schedule

April 28, 2026

Profile Elements

☑

May 5, 2026

Extensions of Technical Content

May 12, 2026

Roles and Profile Delivery Formats

NIST | NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

We Appreciate Your Input

THANK YOU

Your input is a critical part of this process! Thank you for contributing to the development of the Cyber AI Profile!

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

https://www.nccoe.nist.gov/projects/cyber-ai-profile

CyberAIProfile@nist.gov

nccoe.nist.gov

@NISTcyber

NIST AI and Cybersecurity Projects

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

Topic	Learn More!
AI Risk Management Framework (AI RMF) A framework to better manage risks to individuals, organizations, and society associated with artificial intelligence	
Center for AI Standards and Innovation (CAISI) Facilitates testing and collaborative research related to harnessing and securing the potential of commercial AI systems	
Adversarial Machine Learning Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations (NIST AI 100-2 E2025)	
Dioptra A software test platform for assessing the trustworthy characteristics of artificial intelligence systems	
Secure Software Development Framework (SSDF) AI Profile Secure Software Development Practices for Generative AI and Dual-Use Foundation Models: An SSDF Community Profile	

Topic	Learn More!
PETs Test Bed Evaluating Differential Privacy Guarantees	
DevSecOps Secure Software Development, Security, and Operations (DevSecOps) Practices	
Agent Identities Digital Identity Guidelines, Revision 4 (NIST SP 800-63)	
NCCoE Chatbot Secure, internal-use chatbot to assist with discovering and summarizing cybersecurity guidelines	
COSAiS NIST SP 800-53 Control Overlays for Securing AI Systems (COSAiS)	