Check for updates

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

NIST Cybersecurity Framework 2.0:

Quick-Start Guide for Using Artificial Intelligence (AI) for CSF Analysis and Reporting

U.S. Department of Commerce
Howard Lutnick, Secretary of Commerce

National Institute of Standards and Technology
Arvind Raman, Under Secretary of Commerce for Standards and Technology and NIST Director

NIST Special Publication
Initial Public Draft - NIST SP 1353 ipd
https://doi.org/10.6028/NIST.SP.1353.ipd
Please send comments to csf@nist.gov
August 2026

INTRODUCTION

Why Use Artificial Intelligence (AI) with the NIST Cybersecurity Framework (CSF) 2.0?

The CSF 2.0 provides guidelines and a common language to help organizations of all sizes and sectors manage cybersecurity risks. In practice, organizations are increasingly leveraging Artificial Intelligence (AI) systems to better understand, assess, prioritize, and communicate their cybersecurity efforts in alignment with CSF 2.0. For example, AI can help with analyzing, planning, implementing, and monitoring the organization's progress toward achieving the outcomes of the CSF 2.0.

Purpose of this Guide

The following pages provide three notional use cases and examples of prompts for structuring natural language inputs to produce specified CSF 2.0 outputs from a generative AI model. **The examples are notional and illustrate some ways that an organization *might* use AI for planning and implementing CSF 2.0**. Pay attention to the orange exclamation triangles ⚠ throughout the document. These indicate important notes regarding AI usage with the NIST CSF 2.0.

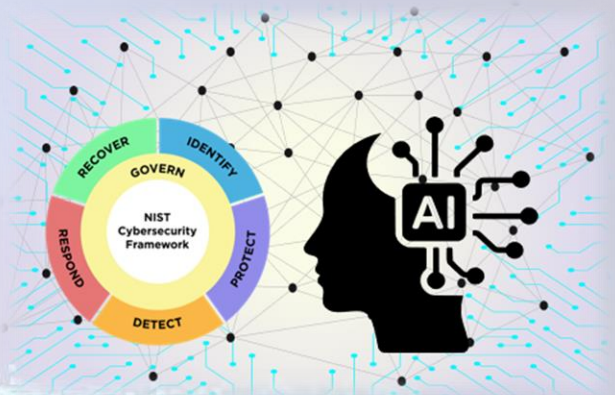

Use Cases Included in this Guide:	What the Use Cases Illustrate:
USE CASE 1: Review Organizational Policies and Strategy to Align with CSF Governance	**USE CASE 1** illustrates the use of AI-assisted review to evaluate organization cybersecurity policy, strategy, and **risk governance in alignment with the CSF 2.0 outcomes**.
USE CASE 2: CSF Current State Profile Development Based on Organizational Artifacts	**USE CASE 2** illustrates how to **produce a draft CSF Organization Current State Profile–** mapping artifacts and personnel interview notes to CSF 2.0 outcomes, documenting any assumptions, and recording observed gaps in the interviews and evidence.
USE CASE 3: Target State Profile Development Based on Organizational Artifacts	**USE CASE 3** illustrates how to draw upon internal and industry references to **create a draft CSF Target State Profile** describing desired outcomes to meet mission objectives, stakeholder expectations, address the risk landscape, and fulfill requirements.

Have an interesting AI use case for the CSF? Submit additional use case ideas to: csf@nist.gov.

INTRODUCTION

Effective Prompt Engineering Helps Support Consistent and Repeatable CSF 2.0 Outcomes

The CO-STAR framework was used to create prompts for the illustrative CSF use cases provided. There are many useful prompt frameworks (e.g., CRAFT, RISEN, RTF, APE, CREATE, etc.). Users should choose a prompt framework that best serves their particular use case.

C — Context
Describe background details that provide information, constraints, and the specific scenario that will affect the CSF outcomes desired.

O — Objective
Describing the purpose of the CSF activity being accomplished helps the AI model to focus its response on meeting that specific goal.

S — Style
Specifying the response style (e.g., technical, managerial, audit-based) helps frame the output for the CSF purpose.

T — Tone
Specifying the tone (i.e., objective, authoritative) helps ensure output is aligned with target audience.

A — Audience
Identifying the intended audience sets the level of abstraction, precision of terms, and response structure.

R — Response
Specifies format & structure of the CSF-related output. Clear, explicit instructions help ensure that the result fulfills the objectives.

Important Notes Before Reading Further:

- The following pages provide descriptions of three use cases and sample executable notional prompts using the CO-STAR framework to illustrate the use of AI for planning and implementing CSF activities. The use case examples illustrate a possible approach and are not prescriptive assessment or assurance methodologies.
- Although AI tools were not used to author this QSG, AI tools, prompts, profiles and fictitious data were used in the prompt research to produce this guide and the sample organizational documents. This usage does not imply that the models, software, profiles or services are the best available for this purpose, nor does it imply recommendation or endorsement by NIST.
- Always review an AI tool's privacy and security settings (e.g., data retention, training, access privileges, and confidentiality terms).
- Follow company policies (e.g., AI security, data management, privacy and security, etc.) before inputting any sensitive information into an AI tool.
- AI-generated content should always be reviewed by qualified personnel before being used in organizational decision-making. Users are responsible for validating applicability, scope, inputs, assumptions, and outputs of AI systems.
- Consider using multiple AI tools and comparing results when validating AI output.
- Consult NIST's AI resources for safe and effective use of AI.
- All individuals, organizations, and records included within this quick-start guide and in the accompanying supplemental materials are products of fiction for illustrative purposes.

The following pages provide descriptions of three use cases and sample executable notional prompts using the CO-STAR framework to illustrate the use of AI for planning and implementing CSF activities. Appendix A provides high-level process for using the guide's supplemental contents to explore using AI for CSF analysis.

Use Case 1: Review Organizational Policies and Strategy to Align with CSF Governance

WHY IT'S USEFUL

CSF 2.0 Governance Alignment View

Provides a structured way to evaluate whether cybersecurity policy, strategy, and risk governance are aligned with CSF 2.0 GOVERN (GV), with emphasis on accountability, oversight, and risk decision-making.

EXAMPLE SOURCE MATERIAL (INPUTS)

- Governance artifacts (e.g., policies, standards, procedures)
- Cybersecurity risk governance context (e.g., risk appetite, tolerance, decisions)
- Organizational strategy documents
- Organizational context (e.g., industry, regulatory environment)
- Supply chain and third-party risk context

EXAMPLE WAYS AI CAN HELP (OUTCOMES)

- CSF 2.0 GV evaluation across OC / RM / RR / PO / OV / SC
- Cybersecurity risk governance consistency analysis across GV domains
- Identification of governance deficiencies and inconsistencies

SIMPLIFIED EXAMPLE OF EXECUTABLE CO-STAR PROMPT*

Context: Evaluate cybersecurity governance using NIST CSF 2.0 GOVERN (GV); inputs include the attached cybersecurity policies and requirements.

Objective: Evaluate consistency of cybersecurity risk governance across CSF 2.0 Govern function, focused on intent, accountability, and risk decision structure

Style: Executive, evidence-based; no inference beyond provided artifacts; no maturity scoring; no benchmarking unless explicitly provided

Tone: Neutral, professional, CISO-level cybersecurity risk governance review; structured and non-prescriptive

Audience: Cybersecurity governance, risk management, and assurance stakeholders

Response: Executive summary of GV alignment posture and key misalignments; GV evaluation across each CSF 2.0 Govern category with status (Aligned / Partial / Misaligned / Not addressed) and supporting evidence; governance breakdowns (policy vs strategy, risk appetite, accountability, oversight, supply chain); governance risk implications (decision authority, regulatory exposure, oversight, risk acceptance)

View the Supplemental Files for Use Case 1

**The prompt on this slide is abbreviated for illustrative purposes. The supplemental files provide a more comprehensive prompt for this use case*

Use Case 2: CSF Current State Profile Development Based on Organizational Artifacts

WHY IT'S USEFUL

Analysis of Current Cybersecurity State

Facilitate mapping of existing practices to the CSF Core's Subcategories to improve efficiency and help rapidly develop a draft Current State Profile.

Automate labor-intensive tasks of reviewing, correlating, and aligning numerous source documents, many of which can be large. Create an effective draft CSF profile with consistent language, reduced variance, and comprehensive outcomes.

EXAMPLE SOURCE MATERIAL (INPUTS)

- Governance and policy documents (e.g., security policies, standards, procedures).
- Existing control frameworks in use (e.g., ISO 27001, SOC 2, PCI DSS, CIS).
- Assessment artifacts (e.g., prior risk assessments, internal/external audit findings, penetration test reports, and vulnerability scan results).
- Other technical and operational process & procedure documentation.
- Organizational strategy documents.
- Organizational context (e.g., business objectives, risk tolerance, regulatory/contractual requirements, and information about practices not otherwise documented).

EXAMPLE WAYS AI CAN HELP (OUTCOMES)

- Compressing the initial drafting from weeks to hours, rapidly ingesting and correlating large document sets.
- Applying uniform language and interpretation that can then be reviewed and refined by human subject matter experts.
- Surfacing cross-references—linking a policy statement to a scan result to an audit finding—that a manual reviewer might overlook.

SIMPLIFIED EXAMPLE OF EXECUTABLE CO-STAR PROMPT*

Context: You are supporting a NIST CSF 2.0 assessment. [Describe any particular sector overlay (e.g., Community Profile) used for an assessment lens.] Work from the NIST CSF 2.0 Organizational Profile template. Use only the attached source materials for your analysis.

Objective: For every outcome in column A (the CSF 2.0 core) populate Column G with current policies and Column H with current practices. Map each CSF outcome to the specific requirement(s) and/or risk response(s) that drive it and cite them.

Style: Source-grounded and traceable. No fabrication. If an outcome is not addressed in the sources, say so plainly. Be Concise and specific with consistent structure across all rows so the columns read as a coherent set.

Tone: Technical but plainly readable. Use precise cybersecurity terminology.

Audience: The CISO and cybersecurity team members, who will use this as the evidence-based Current Profile baseline for gap analysis, prioritization, and (later) target-state planning.

Response: Complete the two columns for all outcomes in a form that maps cleanly back to the template, one row per outcome in column A order. Use the Strength / Partial / Gap prefix in column H wherever the interviews provided a read. After the table, add a short "Assumptions & Evidence Gaps" note listing: (a) any outcomes where the sources were silent or thin, and (b) any places where column H rests on documented process rather than observed practice.

View the Supplemental Files for Use Case 2

**The prompt on this slide is abbreviated for illustrative purposes. The supplemental files provide a more comprehensive prompt for this use case.*

Use Case 3: CSF Target State Profile Development Based on Organizational Artifacts

WHY IT'S USEFUL

Risk-Based Definition of Target State Profile

A CSF Organizational Profile draws from complex and interrelated factors to describe desired outcomes to meet mission objectives, stakeholder expectations, address the risk landscape, and fulfill requirements.

EXAMPLE SOURCE MATERIAL (INPUTS)

- CSF 2.0 Community Profiles created and published to address shared interests and goals among several organizations.
- Risk governance and management artifacts (e.g., policies, senior leaders' risk guidance and direction, risk registers).
- Organizational strategy documents.
- Industry examples and standards of good practices and recommended activities to address known risks and requirements.

EXAMPLE WAYS AI CAN HELP (OUTCOMES)

- Mapping the organization's context against the NIST CSF 2.0 Functions to produce an initial target profile draft, reducing manual research and writing time typically needed to get from a blank page to a workable starting point.
- Ensuring all relevant Categories and Subcategories are considered systematically, flagging gaps or inconsistencies between stated risk priorities and the selected target tiers.
- Translating technical outcomes into plain-language descriptions suited to different audiences (e.g., executives, auditors, technical teams).

SIMPLIFIED EXAMPLE OF EXECUTABLE CO-STAR PROMPT*

Context: You are supporting a NIST CSF 2.0 assessment. [Describe any particular sector overlay (e.g., CRI Profile) used for an assessment lens.] Work from the NIST CSF 2.0 Organizational Profile template. Use only the attached source materials for your analysis.

Objective: Define a target state that will support an acceptable level of risk and fulfill applicable requirements. For every outcome in column A (the CSF 2.0 core) populate Column N with necessary policies & procedures, and Column O with desired practices. Map each CSF outcome to the specific requirement(s) and/or risk response(s) that drive it and cite them.

Style: Source-grounded and traceable. No fabrication. If an outcome is not addressed in the sources, say so plainly. Be Concise and specific with consistent structure across all rows so the columns read as a coherent set.

Tone: Technical but plainly readable. Use precise cybersecurity terminology.

Audience: The CISO and cybersecurity team members, who will use this as the risk-based Target Profile baseline for risk analysis, prioritization, and target-state planning.

Response: Complete the columns for all outcomes in a form that maps cleanly back to the template, one row per outcome in column A order. Cite driver(s) inline in column N (e.g., requirement ID and/or risk ID; best-practice standard where applicable). After the table, add an "Assumptions & Evidence Gaps" note listing: (a) outcomes where the requirements/register defined no explicit target and you applied best practice, (b) targets that imply new tooling, budget, or staffing not confirmed in the sources, and (c) dependencies or sequencing the team should validate before creating a roadmap.

View the Supplemental Files for Use Case 3

**The prompt on this slide is abbreviated for illustrative purposes. The supplemental files provide a more comprehensive prompt for this use case.*

PUTTING IT INTO PRACTICE

Tips for Getting Started

The sample prompts within this quick-start guide provide organizations with a starting point for using AI in planning and implementing CSF 2.0.
Choose a prompt format that best meets the needs of the organization. CO-STAR is used for this quick-start guide, but organizations may want or need to use an alternative prompt format.
Choose an AI tool (or tools) authorized for use by the organization's security and privacy team.
Convene a stakeholder group to discuss scoping, resourcing, roles and responsibilities, etc.
Collect the necessary artifacts (e.g., system security plans, staff interviews, policy handbook, risk register, current CSF organizational profile) that are approved to be ingested into an approved AI tool.
Continuously review and refine inputs and outputs to achieve desired results.

Important Reminders

- Always review an AI tool's privacy and security settings (e.g., data retention, training, access privileges, and confidentiality terms) and company data policy before inputting sensitive information into an AI tool.
- AI-generated content should always be reviewed by qualified personnel before being used in organizational decision-making. Users are responsible for validating applicability, scope, inputs, assumptions, and outputs of AI systems.
- Consider using multiple AI tools and comparing results when validating AI output.

Reference Data and AI

Informative References identify relationships between elements of a focal document such as the CSF 2.0 Core and corresponding elements in other authoritative sources, including frameworks, standards, guidelines, and best practices. The Cybersecurity and Privacy Reference Tool (CPRT), CSF 2.0 Reference Tool, and Online Informative References Program (OLIR) enable organizations to visualize and understand relationships among reference data within various frameworks and guidelines in support of effective implementation.

- Data can be exported from these tools into a structured format. Practitioners can then leverage AI tools to analyze, compare, or crosswalk frameworks.
- AI-assisted mappings should be reviewed, especially where explicit reference data is limited.
- AI-assisted mappings or crosswalks should retain identifiers, source context, provenance, and statuses to avoid unsupported mappings being used without appropriate review or context by a subject-matter expert. Confirm that reference data and other artifacts used reflect the most current published version before using it as input for automated analysis.

Learn More

- AI at NIST
- NIST AI Resource Center (AIRC)
- AI Risk Management Framework (AI RMF)
- Cyber AI Profile

ABBREVIATIONS AND GLOSSARY

Abbreviations

- **AI** — Artificial Intelligence
- **CPRT** — Cybersecurity & Privacy Reference Tool
- **CSF** — Cybersecurity Framework
- **CSRC** — Cybersecurity Resource Center
- **GRC** — Governance, Risk, & Compliance
- **ICT** — Information & Communications Technology
- **ISO** — International Organization for Standardization
- **OLIR** — Online Informative References
- **PCI DSS** — Payment Card Industry Data Security Standard
- **RMF** — Risk Management Framework
- **SOC** — System and Organization Controls
- **SSP** — System Security Plan

Additional NIST CSF and AI Resources

- CSF 2.0 Resource Library
- CSF 2.0 Informative References page
- AI at NIST
- NIST AI Resource Center (AIRC)
- NIST AI Risk Management Framework (AI RMF)

Glossary*

- **CSF Core** — A taxonomy of high-level cybersecurity outcomes to help manage cybersecurity risks. Its components are a hierarchy of Functions, Categories, and Subcategories that detail each outcome.
- **Function** — The highest level of organization for cybersecurity outcomes within the CSF Core. There are six CSF Functions: Govern, Identify, Protect, Detect, Respond, and Recover.
- **Category** — Related cybersecurity outcomes that collectively comprise a CSF Function.
- **Subcategory** — Further divides each Category into more specific outcomes of technical and management activities (e.g., PR.AA-01).
- **Outcome** — A desired cybersecurity result described by a CSF Subcategory.
- **Organizational Profile** — A mechanism for describing an organization's current and/or target cybersecurity posture in terms of the CSF Core's outcomes.
- **Community Profile** — A baseline of CSF outcomes that is created and published to address shared interests and goals among a number of organizations (e.g., a sector or use case).
- **Informative Reference** — Mappings that indicate relationships between the CSF Core and various standards, guidelines, regulations, and other content, helping inform how an organization may achieve Core outcomes.
- **Generative AI** — AI systems that produce novel text or artifacts in response to prompts.
- **Hallucination** — Plausible but inaccurate AI output; requires expert review before use.
- **Prompt engineering** — Iteratively refining prompts to obtain reliable, on-target AI output.
- **Proposed / Derived mapping** — Status labels for AI-assisted mappings pending expert validation.

**Definitions provided are intended as plain language. Please see the NIST Glossary for official NIST definitions.*

APPENDIX A

Use Case Supplemental Contents:

- This quick-start guide is accompanied by a portfolio of supplemental files. They include:
  - **A folder labeled "Organizational Documents*,"** which includes simulated records for a fictitious company called "Halverston Community Bank."
    - These simulated "Organizational Documents" were created with generative AI for illustrative purposes and should not be used as templates for actual use.
  - **A folder for each use case** including:
    - The abbreviated sample prompt from pages 4-6.
    - An expanded version of the sample prompts
    - The provided CO-STAR prompts deliberately reflect stylistic differences (e.g., tone, voice) to provide variety.

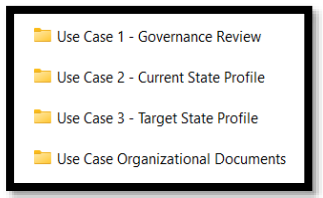

**Any reference to non-NIST sites or resources are included for illustrative purposes. NIST does not necessarily endorse the views expressed or the facts presented on those sites. Further, NIST does not endorse any commercial products that may be advertised or available on these sites.*

Sample Workflow for Use Case Exercise

The following workflow provides a high-level process for using the supplemental contents to begin exploring using AI for CSF analysis.

Open the AI tool of choice.

Choose an AI tool (or tools) authorized for use by the organization's security and privacy team.

Upload files from the "Organizational Documents" folder into the AI tool.

Simulated documents provided:
- Bank security requirements
- Staff interview notes
- Security policy handbook
- Risk Register
- Organizational Pro file template

Copy the entire CO-STAR formatted prompt from one of the use cases and paste into the AI tool.

When presented with options, use a more advanced processing option to reduce likelihood of hallucinations (though they can still occur).

Generate output in the preferred format for review.

AI-generated content should always be reviewed by qualified personnel. The content generated in this exercise reflects a point-in-time output. Results can change over time or may vary.

Next steps might include: *repeat the process with a new use case, try a different AI tool and compare results, or explore how refining the prompts impacts outputs.*