Kyūdō

SOVEREIGNTY-GRADE AI · GRC

# NIST CSF v2.0 Mapping Reference

Vigilance with Purpose. Security with Control.


PRESENTED BY

**Kyūdō — a KMicro Technologies platform**

3525-265 Hyland Avenue

Costa Mesa, CA 92626 · kyudo.ai

hello@kyudo.ai



kyudo.ai · April 2026

CLASSIFICATION: PUBLIC

00 · READ THIS FIRST

# A reference, not an implementation playbook.

This document is a mapping reference for the NIST Cybersecurity Framework version 2.0, published by the National Institute of Standards and Technology on February 26, 2024. It enumerates every Function, Category, and Subcategory in the CSF Core, maps each Subcategory to specific Microsoft Security stack signal sources, and crosswalks the framework against the other standards regulated organizations operate against — SOC 2 Trust Services Criteria, ISO/IEC 27001:2022 Annex A, CMMC v2 Level 2, and HIPAA Security Rule.

Unlike SOC 2 or ISO 27001, NIST CSF is not a certifiable standard. It is a voluntary framework. There is no ‘CSF certification’ from NIST. The framework's value is as a structured taxonomy for assessing posture, defining target outcomes, communicating across executives and practitioners, and aligning a security program against a common language. CSF Profiles — Current and Target — are how organizations operationalize that alignment. CSF Tiers — Partial, Risk Informed, Repeatable, Adaptive — are how organizations characterize the rigor of their cybersecurity risk governance and management.

This reference is built for three audiences. Security leaders who use CSF to communicate posture to boards, executives, regulators, and customers. GRC practitioners who need to translate CSF outcomes into specific control implementations and evidence sources. Compliance teams who carry the cross-framework burden — the same Microsoft signal source, the same access review, the same incident record needs to satisfy CSF, SOC 2, ISO 27001, CMMC, and HIPAA simultaneously. The crosswalks in Section 6 are how that burden becomes manageable.

Section 1 establishes CSF v2.0 in 2026 — the six Functions, what changed from v1.1, and why the new Govern Function rewrites how organizations approach cybersecurity governance. Section 2 covers Tiers. Section 3 covers Profiles. Section 4 is the complete Core enumeration: every Function, every Category, every Subcategory with the exact NIST language. Section 5 maps each Subcategory to its Microsoft Security stack signal source. Section 6 crosswalks every Function to SOC 2, ISO 27001:2022, CMMC L2, and HIPAA. Section 7 is the Kyūdō continuous-readiness model applied to multi-framework operation.

> **This guide reflects NIST CSF v2.0 as published February 26, 2024**
>
> CSF v2.0 expanded the framework from five to six core Functions by adding Govern, expanded scope from critical infrastructure to all organizations, and reorganized into 22 Categories and 106 Subcategories. CSF v1.1 (2018) is superseded; new CSF adoption should be against v2.0.
>
> NIST CSF is a voluntary framework. There is no certification. This reference cites NIST CSWP 29 (February 2024) as its authoritative source. Always defer to the source publication for the latest guidance.

01 · FOUNDATIONS

# CSF v2.0 in 2026.

## The framework, in one paragraph

The NIST Cybersecurity Framework v2.0 is a voluntary framework that provides a taxonomy of high-level cybersecurity outcomes — organized into six Functions, 22 Categories, and 106 Subcategories — that any organization can use to assess its cybersecurity posture, prioritize improvements, and communicate cybersecurity risk to leadership. CSF describes desired outcomes; it does not prescribe how those outcomes must be achieved. Implementation is informed by Informative References (which point to controls in other frameworks like SP 800-53, ISO 27001, and CIS Critical Security Controls) and Implementation Examples (which suggest concrete action steps). The CSF Core, Profiles, and Tiers operate together: Profiles describe current and target posture in CSF terms, Tiers characterize the rigor of governance and management, and the Core provides the outcomes that Profiles measure against.

## What changed in v2.0

CSF v2.0, published February 26, 2024, was the first major revision since the framework's creation in 2014 (v1.0) and its incremental update in 2018 (v1.1). The changes are substantive but the core philosophy is unchanged: outcome-focused, voluntary, technology-neutral, scalable to any organization.

| Dimension | CSF v1.1 (2018) | CSF v2.0 (current) |
|---|---|---|
| Functions | 5 — Identify, Protect, Detect, Respond, Recover | 6 — Govern (NEW), Identify, Protect, Detect, Respond, Recover |
| Categories | 23 | 22 |
| Subcategories | 108 | 106 (16 conceptually new; many merged or relocated) |
| Audience | Critical infrastructure (per the original 2014 Executive Order) | All organizations regardless of size, sector, or maturity |
| Title | Framework for Improving Critical Infrastructure Cybersecurity | The NIST Cybersecurity Framework (CSF) 2.0 |
| Implementation Examples | Limited | Comprehensive online catalog with action-oriented examples per Subcategory |

| Dimension | CSF v1.1 (2018) | CSF v2.0 (current) |
|---|---|---|
| Informative References | Static; updated with framework releases | Online and continuously updated; mappings to SP 800-53, ISO 27001, CIS Controls, and others |

*Source — NIST CSWP 29, The NIST Cybersecurity Framework (CSF) 2.0, February 26, 2024.*

## The six Functions

Each Function is named after a verb that summarizes its outcomes. The Functions are addressed concurrently — they are not a sequential workflow. Govern, Identify, Protect, and Detect operate continuously. Respond and Recover stay ready and engage when incidents occur. Govern sits at the center because it informs how an organization implements the other five.

| Function | Categories | What it covers |
|---|---|---|
| Govern (GV) | 6 | The organization's cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored. New in v2.0. Addresses organizational context, risk management strategy, roles and responsibilities, policy, oversight, and — critically — cybersecurity supply chain risk management as a discrete category. |
| Identify (ID) | 3 | The organization's current cybersecurity risks are understood. Asset inventory across data, hardware, software, services, and people; risk assessment; and improvement processes that span all six Functions. |
| Protect (PR) | 5 | Safeguards to manage the organization's cybersecurity risks are used. Identity, authentication, and access; awareness and training; data security; platform security; and resilience of technology infrastructure. |
| Detect (DE) | 2 | Possible cybersecurity attacks and compromises are found and analyzed. Continuous monitoring and adverse event analysis. |
| Respond (RS) | 4 | Actions regarding a detected cybersecurity incident are taken. Incident management, analysis, reporting and communication, and mitigation. |
| Recover (RC) | 2 | Assets and operations affected by a cybersecurity incident are restored. Recovery plan execution and recovery communication. |

## Why Govern matters

Govern was added in v2.0 because governance had been spread across the Identify Function and other locations in v1.1 — implicitly present, but not centrally addressed. The 2024 update made governance explicit and central. Govern now contains nearly 30 percent of all CSF Subcategories, and Cybersecurity Supply Chain Risk Management (GV.SC) doubled in subcategory count compared to its v1.1 location. The signal is unambiguous: NIST sees governance and supply chain as the two areas where most organizations have the largest gaps.

For organizations adopting CSF v2.0 from a CSF v1.1 starting position, the practical implication is that the Govern Function is where the most new work concentrates: documenting organizational context, formalizing risk management strategy as policy, establishing risk appetite and tolerance, structuring oversight, and operationalizing supply chain risk management as a discrete program.

### Sixteen Subcategories in v2.0 are conceptually new

NIST mapped CSF v2.0 Subcategories back to v1.1 and identified 16 that have no v1.1 equivalent. Most cluster in Govern (organizational context, risk strategy, policy, oversight, supply chain), with additions in Identify (improvement) and Detect (adverse event analysis). Mature programs have likely been performing many of these activities informally; v2.0 names them explicitly so they can be tracked, measured, and improved.

02 · TIERS

# Partial, Risk Informed, Repeatable, Adaptive.

CSF Tiers characterize the rigor of an organization's cybersecurity risk governance and management practices. Tiers are not maturity levels for individual controls — they describe the organization's overall approach to managing cybersecurity risk. An organization can choose to use Tiers to inform its Current and Target Profiles, communicate internally about where it is and where it wants to be, or set a benchmark for how cybersecurity risks are managed enterprise-wide.

The four Tiers describe a progression from informal, ad hoc responses to approaches that are agile, risk-informed, and continuously improving. Higher Tiers are not always the goal — they cost more and require more organizational investment. Selection should be driven by risk exposure, regulatory mandate, and cost-benefit analysis.

| Tier | Cybersecurity risk governance | Cybersecurity risk management |
|---|---|---|
| Tier 1: Partial | Application of the cybersecurity risk strategy is managed in an ad hoc manner. Prioritization is ad hoc and not formally based on objectives or threat environment. | Limited awareness of cybersecurity risks at the organizational level. Implementation occurs irregularly, case by case. The organization may not have processes that enable cybersecurity information to be shared internally. Generally unaware of cybersecurity risks associated with suppliers and their products and services. |
| Tier 2: Risk Informed | Risk management practices are approved by management but may not be established as organization-wide policy. Prioritization is directly informed by organizational risk objectives, the threat environment, or business/mission requirements. | Awareness of cybersecurity risks at the organizational level, but no organization-wide approach to managing them. Cyber risk assessment of organizational and external assets occurs but is not typically repeatable. Cybersecurity information is shared informally. Aware of supplier risks but does not act consistently or formally in response. |
| Tier 3: Repeatable | Risk management practices are formally approved and expressed as policy. Risk-informed policies, | Organization-wide approach to managing cybersecurity risks. Cybersecurity information is |

| Tier | Cybersecurity risk governance | Cybersecurity risk management |
|---|---|---|
| | processes, and procedures are defined, implemented, and reviewed. Cybersecurity practices are regularly updated based on risk management process applied to changes in business/mission, threats, and technology. | routinely shared throughout the organization. Consistent methods are in place to respond effectively to changes in risk. Personnel possess the knowledge and skills to perform appointed roles. The organization consistently and accurately monitors cybersecurity risks of assets. Senior cyber and non-cyber executives communicate regularly. Risk strategy is informed by supplier risks, with formal mechanisms (written agreements, governance structures, policy enforcement). |
| Tier 4: Adaptive | Organization-wide approach to managing cybersecurity risks using risk-informed policies, processes, and procedures. Cybersecurity-risk-to-organizational-objectives relationship is clearly understood and considered when making decisions. Executives monitor cybersecurity risks alongside financial and other organizational risks. Budget reflects current and predicted risk environment and risk tolerance. | Cybersecurity risk management is part of organizational culture. The organization adapts practices based on previous and current cybersecurity activities, including lessons learned and predictive indicators. Continuous improvement, with rapid and efficient adaptation to changing business or mission objectives. |

*Source — NIST CSWP 29, Appendix B. Notional Illustration of the CSF Tiers.*

## How to use Tiers in practice

Tiers work best as a communication tool, not a scoring tool. A board conversation is materially clearer when the CISO can state ‘our current cybersecurity risk management is at Tier 2 with elements of Tier 3, and our target is Tier 3 within 18 months’ than when the conversation is conducted in raw subcategory counts. Tiers translate the technical content of CSF into a register that executives, regulators, and insurers can act on.

That said, Tiers should not be confused with maturity levels for individual controls. A single subcategory can be implemented at high quality (e.g., automated, monitored, drift-detected) inside an organization that operates overall at Tier 2. Tiers describe the organization's approach to managing cybersecurity risk; control-level maturity is assessed separately, often using NIST's Cybersecurity & Privacy Capability Maturity Model (C|P-CMM) or similar frameworks.

03 · PROFILES

# Current Profile, Target Profile, Community Profile.

A CSF Organizational Profile describes an organization's current and/or target cybersecurity posture in terms of the Core's outcomes. Profiles are how organizations operationalize CSF — they convert the abstract framework into something specific, measurable, and aligned to a particular scope. An organization can have as many Profiles as desired, each scoped to different parts of the business: an enterprise-wide Profile, a Profile scoped to a financial system, a Profile scoped to ransomware preparedness, a Profile for a specific subsidiary or product line.

## The two Profile types

| Profile type | What it specifies | Primary use |
|---|---|---|
| Current Profile | The Core outcomes the organization is currently achieving (or attempting to achieve), and how or to what extent each outcome is being achieved. | Documenting today's posture; communicating capabilities and known gaps to stakeholders, partners, customers, regulators. |
| Target Profile | The desired outcomes the organization has selected and prioritized for achieving its risk management objectives. Considers anticipated changes to posture (new requirements, new technology, threat trends). | Setting direction; expressing requirements to suppliers and partners; driving investment and remediation plans. |

Most organizations maintain both. The gap between Current and Target becomes the action plan: a prioritized list of subcategory-level gaps with owners, deadlines, and expected outcomes. The action plan is typically captured in a risk register, a Plan of Action and Milestones (POA&M), or an equivalent tracking artifact.

## Community Profiles

A Community Profile is a baseline of CSF outcomes that addresses shared interests across multiple organizations — typically a sector, subsector, technology category, or threat type. NIST and other organizations publish Community Profiles for use cases including ransomware, manufacturing, healthcare, financial services, and election security. An organization can adopt a Community Profile as the basis for its own Target Profile, which materially reduces the effort of building a Target from scratch and aligns the organization with peers and regulators.

As of 2026, NIST hosts a growing repository of Community Profiles on the CSF website. Adopting a relevant Community Profile is the fastest path to a credible Target Profile for organizations new to CSF.

## Building a Profile — the five steps

1. Scope the Profile. Document the high-level facts and assumptions on which the Profile is based. The scope determines what is in and out: an entire enterprise, a specific business unit, a particular threat scenario, a specific technology environment.
2. Gather the inputs. Organizational policies, risk management priorities, enterprise risk profiles, business impact analyses, cybersecurity requirements and standards followed, practices and tools, and work roles.
3. Create the Profile. For each in-scope CSF outcome, document the relevant information — what is being achieved, by what mechanism, to what extent. Consider using a Community Profile as the basis for the Target.
4. Analyze the gap and create an action plan. Compare Current to Target. Develop a prioritized action plan to close the gaps. Capture in a risk register, risk detail report, or POA&M.
5. Implement the action plan and update the Profile. Execute against the action plan. Update the Current Profile as outcomes are achieved. Repeat the cycle as often as needed for continuous improvement.

04 · THE COMPLETE CSF CORE

# Every Function, every Category, every Subcategory.

This section enumerates the complete CSF v2.0 Core as published in NIST CSWP 29, Appendix A. The language reflects the official NIST text. Subcategory numbering is intentionally non-sequential — gaps indicate v1.1 Subcategories that were relocated, merged, or renumbered in v2.0. Organizations using CSF v2.0 should reference Subcategories by their identifier (e.g., GV.OC-01) and the official text from this enumeration; paraphrasing is fine for communication, but reporting and assessment should anchor to the exact NIST language.

## Govern (GV)

The organization's cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored.

### GV.OC — Organizational Context

The circumstances — mission, stakeholder expectations, dependencies, and legal, regulatory, and contractual requirements — surrounding the organization's cybersecurity risk management decisions are understood.

| ID | Subcategory outcome |
|---|---|
| GV.OC-01 | The organizational mission is understood and informs cybersecurity risk management. |
| GV.OC-02 | Internal and external stakeholders are understood, and their needs and expectations regarding cybersecurity risk management are understood and considered. |
| GV.OC-03 | Legal, regulatory, and contractual requirements regarding cybersecurity — including privacy and civil liberties obligations — are understood and managed. |
| GV.OC-04 | Critical objectives, capabilities, and services that external stakeholders depend on or expect from the organization are understood and communicated. |
| GV.OC-05 | Outcomes, capabilities, and services that the organization depends on are understood and communicated. |

### GV.RM — Risk Management Strategy

The organization's priorities, constraints, risk tolerance and appetite statements, and assumptions are established, communicated, and used to support operational risk decisions.

| ID | Subcategory outcome |
|---|---|
| GV.RM-01 | Risk management objectives are established and agreed to by organizational stakeholders. |
| GV.RM-02 | Risk appetite and risk tolerance statements are established, communicated, and maintained. |
| GV.RM-03 | Cybersecurity risk management activities and outcomes are included in enterprise risk management processes. |
| GV.RM-04 | Strategic direction that describes appropriate risk response options is established and communicated. |
| GV.RM-05 | Lines of communication across the organization are established for cybersecurity risks, including risks from suppliers and other third parties. |
| GV.RM-06 | A standardized method for calculating, documenting, categorizing, and prioritizing cybersecurity risks is established and communicated. |
| GV.RM-07 | Strategic opportunities (i.e., positive risks) are characterized and are included in organizational cybersecurity risk discussions. |

## GV.RR — Roles, Responsibilities, and Authorities

Cybersecurity roles, responsibilities, and authorities to foster accountability, performance assessment, and continuous improvement are established and communicated.

| ID | Subcategory outcome |
|---|---|
| GV.RR-01 | Organizational leadership is responsible and accountable for cybersecurity risk and fosters a culture that is risk-aware, ethical, and continually improving. |
| GV.RR-02 | Roles, responsibilities, and authorities related to cybersecurity risk management are established, communicated, understood, and enforced. |
| GV.RR-03 | Adequate resources are allocated commensurate with the cybersecurity risk strategy, roles, responsibilities, and policies. |
| GV.RR-04 | Cybersecurity is included in human resources practices. |

## GV.PO — Policy

Organizational cybersecurity policy is established, communicated, and enforced.

| ID | Subcategory outcome |
|---|---|
| GV.PO-01 | Policy for managing cybersecurity risks is established based on organizational context, cybersecurity strategy, and priorities and is communicated and |

| ID | Subcategory outcome |
|---|---|
| | enforced. |
| GV.PO-02 | Policy for managing cybersecurity risks is reviewed, updated, communicated, and enforced to reflect changes in requirements, threats, technology, and organizational mission. |

## GV.OV — Oversight

Results of organization-wide cybersecurity risk management activities and performance are used to inform, improve, and adjust the risk management strategy.

| ID | Subcategory outcome |
|---|---|
| GV.OV-01 | Cybersecurity risk management strategy outcomes are reviewed to inform and adjust strategy and direction. |
| GV.OV-02 | The cybersecurity risk management strategy is reviewed and adjusted to ensure coverage of organizational requirements and risks. |
| GV.OV-03 | Organizational cybersecurity risk management performance is evaluated and reviewed for adjustments needed. |

## GV.SC — Cybersecurity Supply Chain Risk Management

Cyber supply chain risk management processes are identified, established, managed, monitored, and improved by organizational stakeholders. The category doubled in subcategory count from CSF v1.1 — NIST's strongest signal that supply chain is now a first-class concern, not a footnote within Identify.

| ID | Subcategory outcome |
|---|---|
| GV.SC-01 | A cybersecurity supply chain risk management program, strategy, objectives, policies, and processes are established and agreed to by organizational stakeholders. |
| GV.SC-02 | Cybersecurity roles and responsibilities for suppliers, customers, and partners are established, communicated, and coordinated internally and externally. |
| GV.SC-03 | Cybersecurity supply chain risk management is integrated into cybersecurity and enterprise risk management, risk assessment, and improvement processes. |
| GV.SC-04 | Suppliers are known and prioritized by criticality. |
| GV.SC-05 | Requirements to address cybersecurity risks in supply chains are established, prioritized, and integrated into contracts and other types of agreements with |

| ID | Subcategory outcome |
|---|---|
| | suppliers and other relevant third parties. |
| GV.SC-06 | Planning and due diligence are performed to reduce risks before entering into formal supplier or other third-party relationships. |
| GV.SC-07 | The risks posed by a supplier, their products and services, and other third parties are understood, recorded, prioritized, assessed, responded to, and monitored over the course of the relationship. |
| GV.SC-08 | Relevant suppliers and other third parties are included in incident planning, response, and recovery activities. |
| GV.SC-09 | Supply chain security practices are integrated into cybersecurity and enterprise risk management programs, and their performance is monitored throughout the technology product and service life cycle. |
| GV.SC-10 | Cybersecurity supply chain risk management plans include provisions for activities that occur after the conclusion of a partnership or service agreement. |

# Identify (ID)

The organization's current cybersecurity risks are understood.

## ID.AM — Asset Management

Assets (e.g., data, hardware, software, systems, facilities, services, people) that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to organizational objectives and the organization's risk strategy.

| ID | Subcategory outcome |
|---|---|
| ID.AM-01 | Inventories of hardware managed by the organization are maintained. |
| ID.AM-02 | Inventories of software, services, and systems managed by the organization are maintained. |
| ID.AM-03 | Representations of the organization's authorized network communication and internal and external network data flows are maintained. |
| ID.AM-04 | Inventories of services provided by suppliers are maintained. |
| ID.AM-05 | Assets are prioritized based on classification, criticality, resources, and impact on the mission. |
| ID.AM-07 | Inventories of data and corresponding metadata for designated data types are maintained. |
| ID.AM-08 | Systems, hardware, software, services, and data are managed throughout their life cycles. |

## ID.RA — Risk Assessment

The cybersecurity risk to the organization, assets, and individuals is understood by the organization.

| ID | Subcategory outcome |
|---|---|
| ID.RA-01 | Vulnerabilities in assets are identified, validated, and recorded. |
| ID.RA-02 | Cyber threat intelligence is received from information sharing forums and sources. |
| ID.RA-03 | Internal and external threats to the organization are identified and recorded. |
| ID.RA-04 | Potential impacts and likelihoods of threats exploiting vulnerabilities are identified and recorded. |
| ID.RA-05 | Threats, vulnerabilities, likelihoods, and impacts are used to understand inherent risk and inform risk response prioritization. |

| ID | Subcategory outcome |
|---|---|
| ID.RA-06 | Risk responses are chosen, prioritized, planned, tracked, and communicated. |
| ID.RA-07 | Changes and exceptions are managed, assessed for risk impact, recorded, and tracked. |
| ID.RA-08 | Processes for receiving, analyzing, and responding to vulnerability disclosures are established. |
| ID.RA-09 | The authenticity and integrity of hardware and software are assessed prior to acquisition and use. |
| ID.RA-10 | Critical suppliers are assessed prior to acquisition. |

## ID.IM — Improvement

Improvements to organizational cybersecurity risk management processes, procedures, and activities are identified across all CSF Functions.

| ID | Subcategory outcome |
|---|---|
| ID.IM-01 | Improvements are identified from evaluations. |
| ID.IM-02 | Improvements are identified from security tests and exercises, including those done in coordination with suppliers and relevant third parties. |
| ID.IM-03 | Improvements are identified from execution of operational processes, procedures, and activities. |
| ID.IM-04 | Incident response plans and other cybersecurity plans that affect operations are established, communicated, maintained, and improved. |

# Protect (PR)

Safeguards to manage the organization's cybersecurity risks are used.

## PR.AA — Identity Management, Authentication, and Access Control

Access to physical and logical assets is limited to authorized users, services, and hardware and managed commensurate with the assessed risk of unauthorized access.

| ID | Subcategory outcome |
|---|---|
| PR.AA-01 | Identities and credentials for authorized users, services, and hardware are managed by the organization. |
| PR.AA-02 | Identities are proofed and bound to credentials based on the context of interactions. |
| PR.AA-03 | Users, services, and hardware are authenticated. |
| PR.AA-04 | Identity assertions are protected, conveyed, and verified. |
| PR.AA-05 | Access permissions, entitlements, and authorizations are defined in a policy, managed, enforced, and reviewed, and incorporate the principles of least privilege and separation of duties. |
| PR.AA-06 | Physical access to assets is managed, monitored, and enforced commensurate with risk. |

## PR.AT — Awareness and Training

The organization's personnel are provided with cybersecurity awareness and training so that they can perform their cybersecurity-related tasks.

| ID | Subcategory outcome |
|---|---|
| PR.AT-01 | Personnel are provided with awareness and training so that they possess the knowledge and skills to perform general tasks with cybersecurity risks in mind. |
| PR.AT-02 | Individuals in specialized roles are provided with awareness and training so that they possess the knowledge and skills to perform relevant tasks with cybersecurity risks in mind. |

## PR.DS — Data Security

Data are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information. PR.DS-10 (data-in-use) is conceptually new in v2.0 — reflecting modern concerns about memory-resident data, side-channel attacks, and confidential computing.

| ID | Subcategory outcome |
|---|---|
| PR.DS-01 | The confidentiality, integrity, and availability of data-at-rest are protected. |
| PR.DS-02 | The confidentiality, integrity, and availability of data-in-transit are protected. |
| PR.DS-10 | The confidentiality, integrity, and availability of data-in-use are protected. |
| PR.DS-11 | Backups of data are created, protected, maintained, and tested. |

## PR.PS — Platform Security

The hardware, software (e.g., firmware, operating systems, applications), and services of physical and virtual platforms are managed consistent with the organization's risk strategy to protect their confidentiality, integrity, and availability.

| ID | Subcategory outcome |
|---|---|
| PR.PS-01 | Configuration management practices are established and applied. |
| PR.PS-02 | Software is maintained, replaced, and removed commensurate with risk. |
| PR.PS-03 | Hardware is maintained, replaced, and removed commensurate with risk. |
| PR.PS-04 | Log records are generated and made available for continuous monitoring. |
| PR.PS-05 | Installation and execution of unauthorized software are prevented. |
| PR.PS-06 | Secure software development practices are integrated, and their performance is monitored throughout the software development life cycle. |

## PR.IR — Technology Infrastructure Resilience

Security architectures are managed with the organization's risk strategy to protect asset confidentiality, integrity, and availability, and organizational resilience.

| ID | Subcategory outcome |
|---|---|
| PR.IR-01 | Networks and environments are protected from unauthorized logical access and usage. |
| PR.IR-02 | The organization's technology assets are protected from environmental threats. |
| PR.IR-03 | Mechanisms are implemented to achieve resilience requirements in normal and adverse situations. |
| PR.IR-04 | Adequate resource capacity to ensure availability is maintained. |

# Detect (DE)

Possible cybersecurity attacks and compromises are found and analyzed.

## DE.CM — Continuous Monitoring

Assets are monitored to find anomalies, indicators of compromise, and other potentially adverse events.

| ID | Subcategory outcome |
|---|---|
| DE.CM-01 | Networks and network services are monitored to find potentially adverse events. |
| DE.CM-02 | The physical environment is monitored to find potentially adverse events. |
| DE.CM-03 | Personnel activity and technology usage are monitored to find potentially adverse events. |
| DE.CM-06 | External service provider activities and services are monitored to find potentially adverse events. |
| DE.CM-09 | Computing hardware and software, runtime environments, and their data are monitored to find potentially adverse events. |

## DE.AE — Adverse Event Analysis

Anomalies, indicators of compromise, and other potentially adverse events are analyzed to characterize the events and detect cybersecurity incidents.

| ID | Subcategory outcome |
|---|---|
| DE.AE-02 | Potentially adverse events are analyzed to better understand associated activities. |
| DE.AE-03 | Information is correlated from multiple sources. |
| DE.AE-04 | The estimated impact and scope of adverse events are understood. |
| DE.AE-06 | Information on adverse events is provided to authorized staff and tools. |
| DE.AE-07 | Cyber threat intelligence and other contextual information are integrated into the analysis. |
| DE.AE-08 | Incidents are declared when adverse events meet the defined incident criteria. |

# Respond (RS)

Actions regarding a detected cybersecurity incident are taken.

## RS.MA — Incident Management

Responses to detected cybersecurity incidents are managed.

| ID | Subcategory outcome |
|---|---|
| RS.MA-01 | The incident response plan is executed in coordination with relevant third parties once an incident is declared. |
| RS.MA-02 | Incident reports are triaged and validated. |
| RS.MA-03 | Incidents are categorized and prioritized. |
| RS.MA-04 | Incidents are escalated or elevated as needed. |
| RS.MA-05 | The criteria for initiating incident recovery are applied. |

## RS.AN — Incident Analysis

Investigations are conducted to ensure effective response and support forensics and recovery activities.

| ID | Subcategory outcome |
|---|---|
| RS.AN-03 | Analysis is performed to establish what has taken place during an incident and the root cause of the incident. |
| RS.AN-06 | Actions performed during an investigation are recorded, and the records' integrity and provenance are preserved. |
| RS.AN-07 | Incident data and metadata are collected, and their integrity and provenance are preserved. |
| RS.AN-08 | An incident's magnitude is estimated and validated. |

## RS.CO — Incident Response Reporting and Communication

Response activities are coordinated with internal and external stakeholders as required by laws, regulations, or policies.

| ID | Subcategory outcome |
|---|---|
| RS.CO-02 | Internal and external stakeholders are notified of incidents. |
| RS.CO-03 | Information is shared with designated internal and external stakeholders. |

## RS.MI — Incident Mitigation

Activities are performed to prevent expansion of an event and mitigate its effects.

| ID | Subcategory outcome |
|---|---|
| RS.MI-01 | Incidents are contained. |
| RS.MI-02 | Incidents are eradicated. |

## Recover (RC)

Assets and operations affected by a cybersecurity incident are restored.

### RC.RP — Incident Recovery Plan Execution

Restoration activities are performed to ensure operational availability of systems and services affected by cybersecurity incidents.

| ID | Subcategory outcome |
|---|---|
| RC.RP-01 | The recovery portion of the incident response plan is executed once initiated from the incident response process. |
| RC.RP-02 | Recovery actions are selected, scoped, prioritized, and performed. |
| RC.RP-03 | The integrity of backups and other restoration assets is verified before using them for restoration. |
| RC.RP-04 | Critical mission functions and cybersecurity risk management are considered to establish post-incident operational norms. |
| RC.RP-05 | The integrity of restored assets is verified, systems and services are restored, and normal operating status is confirmed. |
| RC.RP-06 | The end of incident recovery is declared based on criteria, and incident-related documentation is completed. |

### RC.CO — Incident Recovery Communication

Restoration activities are coordinated with internal and external parties.

| ID | Subcategory outcome |
|---|---|
| RC.CO-03 | Recovery activities and progress in restoring operational capabilities are communicated to designated internal and external stakeholders. |
| RC.CO-04 | Public updates on incident recovery are shared using approved methods and messaging. |

*Source — NIST CSWP 29, Appendix A, The NIST Cybersecurity Framework (CSF) 2.0, February 26, 2024. Subcategory text reproduced from the official NIST publication.*

05 · CSF V2.0 × MICROSOFT SECURITY STACK

# From Subcategory outcome to native Microsoft signal.

This section maps high-frequency CSF v2.0 Subcategories to specific Microsoft Security stack signal sources. The mapping is not exhaustive — it covers the Subcategories where Microsoft-native evidence is most direct and most defensible. Where additional sources exist in your environment (third-party MDM, code-scanning tools, separate ticketing systems), they become supplementary evidence for the same outcomes.

## Govern — GV

Govern outcomes are largely governance and process activities that produce documentary evidence rather than telemetry. Microsoft signal applies most directly where governance intersects with technical configuration: privileged role assignments, policy enforcement, supplier app governance.

| Subcategory | Primary Microsoft signal source |
|---|---|
| GV.OC-03 (legal/regulatory) | Microsoft Purview Compliance Manager assessments mapped to applicable regulations; tenant-level data residency configuration evidence. |
| GV.RR-02 (roles/responsibilities) | Microsoft Entra ID role assignments, Privileged Identity Management role catalog, administrative-unit assignments. |
| GV.RR-03 (resources) | Microsoft 365 license inventory, Azure subscription-level cost and resource allocation. |
| GV.PO-01–02 (policy) | SharePoint policy library with versioning enabled; Microsoft Purview retention labels enforcing policy lifecycle. |
| GV.SC-04 (suppliers known and prioritized) | Microsoft Defender for Cloud Apps cloud-app inventory; Entra ID third-party application consent inventory. |
| GV.SC-07 (supplier risks monitored) | Defender for Cloud Apps risk scoring per discovered cloud app; Entra ID third-party application activity logs. |
| GV.SC-09 (supply chain integrated) | Microsoft Service Trust Portal reports for inherited Microsoft subservice posture; documented Microsoft as material supplier in vendor inventory. |

## Identify — ID

| Subcategory | Primary Microsoft signal source |
|---|---|
| ID.AM-01 (hardware inventory) | Microsoft Intune device inventory; Microsoft Defender for Endpoint device list; Azure Resource Graph queries for compute resources. |
| ID.AM-02 (software/services/systems inventory) | Microsoft Defender Vulnerability Management software inventory; Microsoft 365 admin center service inventory; Azure Resource Graph queries. |
| ID.AM-03 (network communication/data flows) | Azure Network Watcher topology, NSG flow logs, VNet topology; Microsoft Sentinel network analytics. |
| ID.AM-04 (supplier services inventory) | Defender for Cloud Apps cloud-app discovery; Azure Marketplace consumption; M365 connected apps inventory. |
| ID.AM-05 (asset prioritization) | Microsoft Purview data classification driving asset criticality; Azure resource tags. |
| ID.AM-07 (data inventory) | Microsoft Purview Data Map and Data Catalog; sensitivity-label inventory across M365. |
| ID.AM-08 (asset lifecycle) | Azure resource lifecycle policies; Intune device lifecycle and retirement workflows. |
| ID.RA-01 (vulnerabilities identified) | Microsoft Defender Vulnerability Management findings; Defender for Cloud recommendations; third-party scanner findings exported to Sentinel. |
| ID.RA-02 (threat intelligence received) | Microsoft Defender Threat Intelligence; Sentinel threat-intelligence connectors; ISAC/ISAO feeds integrated to Sentinel. |
| ID.RA-05 (inherent risk understood) | Microsoft Defender for Cloud Secure Score and trend; Microsoft Purview Compliance Manager improvement actions. |
| ID.RA-08 (vulnerability disclosure) | Documented coordinated vulnerability disclosure process; Microsoft Security Response Center engagement record where applicable. |
| ID.RA-09 (authenticity/inte | Microsoft Defender Vulnerability Management software baseline; signed software requirements in Intune. |

| Subcategory | Primary Microsoft signal source |
|---|---|
| grity prior to acquisition) | |
| ID.RA-10 (critical suppliers assessed) | Documented vendor security review prior to procurement; Microsoft as supplier with Service Trust Portal evidence. |
| ID.IM-01–04 (improvements identified) | Sentinel incident retrospectives, Defender for Cloud Secure Score deltas over time, Microsoft Purview Compliance Manager improvement-action history. |

## Protect — PR

The Protect Function is where Microsoft-native evidence is densest. Identity, data, platform, and infrastructure controls all have direct API and audit-log evidence.

| Subcategory | Primary Microsoft signal source |
|---|---|
| PR.AA-01 (identities/credentials managed) | Microsoft Entra ID user inventory, service principal inventory, managed identity inventory. |
| PR.AA-02 (identity proofing) | Entra Verified ID (where deployed); Entra ID identity verification policies; HRIS-to-Entra provisioning evidence. |
| PR.AA-03 (authentication) | Entra ID sign-in logs; Conditional Access policy enforcement records; Authentication Methods activity report. |
| PR.AA-04 (identity assertions) | Entra ID token issuance logs; SAML/OIDC federation configuration; certificate inventory in Azure Key Vault. |
| PR.AA-05 (least privilege/SoD) | Entra Privileged Identity Management role catalog and activation logs; Entra ID access reviews; conflicting-role detections. |
| PR.AA-06 (physical access) | Inherited from Microsoft Azure subservice (Service Trust Portal); on-premises facility access records where applicable. |
| PR.AT-01–02 (awareness/training) | Microsoft Defender for Office 365 attack-simulation training records; Microsoft Viva Learning completion records; LMS exports. |
| PR.DS-01 (data-at-rest) | Azure Storage encryption settings; Azure SQL Transparent Data Encryption status; customer-managed keys in Azure Key Vault; Microsoft Purview encryption labels. |

| Subcategory | Primary Microsoft signal source |
|---|---|
| PR.DS-02 (data-in-transit) | TLS configuration evidence (Azure Front Door, Application Gateway, App Service); certificate inventory and expiry monitoring. |
| PR.DS-10 (data-in-use) | Azure Confidential Computing workloads; Microsoft 365 customer key (Service Encryption); Azure SQL Always Encrypted with secure enclaves. |
| PR.DS-11 (backups) | Azure Backup recovery point success metrics; restore-test reports; retention schedule alignment with policy. |
| PR.PS-01 (configuration management) | Azure Policy compliance state over time; Defender for Cloud regulatory compliance dashboard; Intune compliance and configuration policies. |
| PR.PS-02 (software lifecycle) | Defender Vulnerability Management software inventory; Intune software update policies; M365 service health and feature lifecycle. |
| PR.PS-04 (logs generated) | Entra ID audit and sign-in logs; M365 Unified Audit Log; Azure Activity Log; Sentinel ingestion connectors. |
| PR.PS-05 (unauthorized software prevention) | Microsoft Defender Application Control; AppLocker; Intune app-control policies. |
| PR.PS-06 (secure development) | GitHub Advanced Security findings; Azure DevOps pipeline security checks; SAST/DAST integration evidence. |
| PR.IR-01 (network logical access) | Azure Firewall logs; NSG rules; private endpoint topology; Defender for Cloud network recommendations. |
| PR.IR-02 (environmental threats) | Inherited from Azure subservice for cloud workloads; on-premises environmental monitoring evidence where applicable. |
| PR.IR-03 (resilience mechanisms) | Azure availability zone deployment; paired-region replication; Azure Site Recovery test failovers. |
| PR.IR-04 (capacity) | Azure Monitor capacity metrics; auto-scaling policy and history; capacity-planning documentation. |

## Detect — DE

Detect is the function where Microsoft Sentinel and the Defender XDR family produce the most direct evidence. Continuous monitoring and adverse event analysis are operational SOC concerns mapped to specific telemetry pipelines.

| Subcategory | Primary Microsoft signal source |
|---|---|
| DE.CM-01 (network monitoring) | Microsoft Sentinel analytic rules over network signal; Defender for Cloud network alerts; Azure Firewall threat intel filtering. |
| DE.CM-02 (physical environment) | Inherited from Azure subservice for cloud-hosted workloads; on-premises CCTV/access systems integrated to SIEM where applicable. |
| DE.CM-03 (personnel/technology usage) | Microsoft Defender for Cloud Apps user activity; Insider Risk Management signals; Entra ID Identity Protection risk detections. |
| DE.CM-06 (external service provider monitoring) | Defender for Cloud Apps third-party app activity; Entra ID guest user activity logs; service-principal sign-in logs. |
| DE.CM-09 (computing hardware/software/runtime) | Microsoft Defender for Endpoint EDR telemetry; Defender for Cloud workload protection; Defender for Identity ATP signals. |
| DE.AE-02 (events analyzed) | Sentinel investigation graph; Defender XDR automated investigations; threat-hunting queries history. |
| DE.AE-03 (correlation across sources) | Sentinel cross-source correlation rules; Defender XDR unified incidents; Sentinel Fusion ML-based correlation. |
| DE.AE-04 (impact/scope) | Sentinel incident scoping (entities involved, severity, blast radius); Defender XDR incident impact assessment. |
| DE.AE-06 (information to staff/tools) | Sentinel incident assignment workflows; Microsoft Teams notifications; ServiceNow / Jira integration via Sentinel playbooks. |
| DE.AE-07 (threat intelligence integrated) | Sentinel threat-intelligence indicators applied to detection rules; Defender Threat Intelligence enrichment. |
| DE.AE-08 (incidents declared) | Sentinel incident creation criteria; Defender XDR incident generation logic; documented incident-declaration thresholds. |

## Respond — RS

| Subcategory | Primary Microsoft signal source |
|---|---|
| RS.MA-01 (IR plan execution) | Sentinel incident workflow; documented runbooks attached to incidents; SOAR playbook execution history. |
| RS.MA-02 (triage/validation) | Sentinel incident triage records; Defender XDR investigation outcomes. |
| RS.MA-03 (categorize/prioritize) | Sentinel incident severity assignment; Defender XDR incident classification. |
| RS.MA-04 (escalation) | Sentinel automation rules; Microsoft Teams escalation channels; documented escalation matrix. |
| RS.AN-03 (root cause analysis) | Sentinel investigation graph and entity timeline; Defender XDR attack timeline. |
| RS.AN-06–07 (records preserved) | Sentinel incident artifacts with immutable retention; Microsoft Purview eDiscovery for forensic preservation; Azure Storage immutable blob retention. |
| RS.AN-08 (magnitude estimated) | Sentinel incident impact metrics; Defender XDR scope of compromise; affected-entity counts. |
| RS.CO-02–03 (notification/sharing) | Sentinel SOAR playbooks for stakeholder notification; documented notification matrix; secure communication channels. |
| RS.MI-01 (containment) | Defender for Endpoint device isolation; Entra ID emergency access revocation; Conditional Access break-glass policies; Sentinel automated containment playbooks. |
| RS.MI-02 (eradication) | Defender for Endpoint malware removal; automated investigation remediation; Entra ID compromised account remediation. |

## Recover — RC

| Subcategory | Primary Microsoft signal source |
|---|---|
| RC.RP-01 (recovery executed) | Documented recovery runbook execution; Azure Site Recovery failover records; Azure Backup restore jobs. |
| RC.RP-02 (recovery actions | Sentinel incident recovery tasks; documented recovery decision artifacts. |

| Subcategory | Primary Microsoft signal source |
|---|---|
| selected/scoped) | |
| RC.RP-03 (backup integrity verified) | Azure Backup recovery-point validation; restore-test reports; backup checksum verification logs. |
| RC.RP-04 (post-incident operational norms) | Documented after-action review; Sentinel incident lessons-learned attachments. |
| RC.RP-05 (restoration verified) | Azure Site Recovery post-failover validation; system-health checks; documented service-restoration evidence. |
| RC.RP-06 (end of recovery declared) | Documented recovery-completion criteria and sign-off; Sentinel incident closure with completion narrative. |
| RC.CO-03–04 (communication) | Documented stakeholder communication; M365 distribution lists for status; public status page updates. |

# CSF as a Rosetta Stone.

CSF v2.0 is purposefully designed as a translation layer. Its Functions and Categories were structured to map cleanly to other frameworks regulated organizations operate against — SOC 2 Trust Services Criteria, ISO/IEC 27001:2022 Annex A, CMMC v2 Level 2, and HIPAA Security Rule are the most common in mid-market and enterprise contexts. NIST publishes Informative References that map specific Subcategories to specific controls in other frameworks; this section provides function-level crosswalks for navigation and high-density category-level crosswalks for the most consequential mappings.

These crosswalks are practitioner-grade. They enable an organization to ask: if we invest in implementing PR.AA (identity management, authentication, and access control) at high quality, what does that buy us across the rest of our framework portfolio? The answer is most of SOC 2 CC6, most of ISO 27001 A.5.15 + A.8.2 through A.8.5, CMMC L2 AC and IA control families, and HIPAA Security Rule §164.312(a) and (d). The same evidence satisfies multiple obligations — a properly architected control program tests once and reports many times.

## Function-level crosswalk

This high-altitude view shows where each CSF Function lands in the other major frameworks. Use it to set expectations during program planning; use the category-level crosswalks below for operational mapping.

| CSF v2.0 Function | SOC 2 TSC | ISO 27001:2022 | CMMC v2 L2 | HIPAA Security Rule |
|---|---|---|---|---|
| Govern (GV) | CC1 (Control Environment), CC2 (Communication), CC3 (Risk Assessment), CC9 (Risk Mitigation — vendors) | Clauses 4–6 (Context, Leadership, Planning); A.5.1–5.7 (policies); A.5.19–5.23 (suppliers, cloud) | RA, CA (Risk Assessment, Security Assessment) families | §164.308(a)(1)(i) (Security Management Process) |
| Identify (ID) | CC3 (Risk Assessment), CC4 (Monitoring — partial) | A.5.9 (asset inventory); A.5.12 (classification); 6.1.2 (risk assessment) | CM, RA (Configuration Management, Risk Assessment) | §164.308(a)(1)(ii)(A) (Risk Analysis); §164.310 (Physical Safeguards — |

| CSF v2.0 Function | SOC 2 TSC | ISO 27001:2022 | CMMC v2 L2 | HIPAA Security Rule |
|---|---|---|---|---|
| | | | | partial) |
| Protect (PR) | CC6 (Logical and Physical Access), CC8 (Change Management), CC5 (Control Activities) | A.5.15, A.5.18; A.6.1–6.8; A.7.1–7.14; A.8.1–8.34 | AC, IA, MA, MP, PE, PS, SC, SI (most technical and people families) | §164.308(a)(3)–(a)(7); §164.310; §164.312 |
| Detect (DE) | CC7 (System Operations), CC4 (Monitoring) | A.8.15 (logging), A.8.16 (monitoring); A.5.24 (incident detection) | AU, SI (Audit, System and Information Integrity) | §164.308(a)(1)(ii)(D) (Information System Activity Review); §164.312(b) (Audit Controls) |
| Respond (RS) | CC7 (System Operations — incident handling) | A.5.24–5.28 (incident management) | IR (Incident Response) | §164.308(a)(6) (Security Incident Procedures); §164.404 (Notification) |
| Recover (RC) | A1 (Availability) where in scope; CC7 (System Operations) | A.5.30 (ICT readiness); A.8.13 (backup); A.8.14 (redundancy) | CP (Contingency Planning) | §164.308(a)(7) (Contingency Plan) |

*Crosswalk reflects principal mappings. Specific Subcategory-to-control mappings vary by Informative Reference and framework version. NIST Informative References are the authoritative source for granular crosswalks.*

## Govern (GV) × SOC 2 + ISO 27001:2022 — detail

| CSF Subcategory | SOC 2 Common Criteria | ISO 27001:2022 Annex A / Clause |
|---|---|---|
| GV.OC-01 (mission) | CC1.1, CC1.5 | Clause 4.1 (Context); A.5.1 |
| GV.OC-02 (stakeholders) | CC1.2, CC2.3 | Clause 4.2 (Interested parties) |

| CSF Subcategory | SOC 2 Common Criteria | ISO 27001:2022 Annex A / Clause |
|---|---|---|
| GV.OC-03 (legal/regulatory) | CC2.3, CC9.1 | A.5.31, A.5.32, A.5.33, A.5.34 (legal, IP, records, privacy) |
| GV.RM-01–07 (risk management strategy) | CC3.1, CC3.2, CC3.3, CC3.4 | Clause 6.1 (Risk planning); A.5.4 (management responsibility) |
| GV.RR-01–04 (roles) | CC1.3, CC1.4, CC1.5 | Clause 5.3 (Roles); A.5.2, A.5.3 (segregation of duties) |
| GV.PO-01–02 (policy) | CC2.2, CC5.3 | Clause 5.2 (Policy); A.5.1 (information security policy) |
| GV.OV-01–03 (oversight) | CC4.1, CC4.2 | Clause 9.3 (Management review); Clause 10 (Improvement) |
| GV.SC-01–10 (supply chain) | CC9.2 (vendor risk) | A.5.19, A.5.20, A.5.21, A.5.22, A.5.23 (supplier relationships, cloud) |

## Identify (ID) × SOC 2 + ISO 27001:2022 — detail

| CSF Subcategory | SOC 2 Common Criteria | ISO 27001:2022 Annex A / Clause |
|---|---|---|
| ID.AM-01–05 (asset inventory) | CC3.2, CC6.1 | A.5.9, A.5.10, A.5.11, A.5.12, A.5.13 |
| ID.AM-07 (data inventory) | CC3.2, CC6.1 | A.5.12 (classification), A.5.13 (labelling) |
| ID.AM-08 (asset lifecycle) | CC6.1, CC6.5 | A.5.14, A.7.10 (storage media), A.7.14 (secure disposal) |
| ID.RA-01–10 (risk assessment) | CC3.1–3.4, CC9.1 | Clause 6.1.2; A.5.7 (threat intelligence); A.5.31 (legal/regulatory) |
| ID.IM-01–04 (improvement) | CC4.1, CC4.2 | Clause 10 (Improvement); Clause 9.2 (Internal audit) |

## Protect (PR) × SOC 2 + ISO 27001:2022 — detail

| CSF Subcategory | SOC 2 Common Criteria | ISO 27001:2022 Annex A |
|---|---|---|
| PR.AA-01–06 (identity/auth/access) | CC6.1, CC6.2, CC6.3, CC6.6 | A.5.15, A.5.16, A.5.17, A.5.18, A.8.2, A.8.3, A.8.5 |
| PR.AT-01–02 (training) | CC2.2, CC1.4 | A.6.3 (awareness, education, training) |
| PR.DS-01 (data-at-rest) | CC6.1, CC6.7 | A.8.24 (cryptography), A.5.13 (labelling), A.5.14 (information transfer) |
| PR.DS-02 (data-in-transit) | CC6.1, CC6.7 | A.8.24, A.8.20 (network controls) |
| PR.DS-10 (data-in-use) | CC6.1, CC6.7 | A.8.24 (cryptography — with confidential computing); A.8.11 (data masking) |
| PR.DS-11 (backups) | A1.2 (Availability) | A.8.13 (information backup) |
| PR.PS-01 (config mgmt) | CC8.1 | A.8.9 (configuration management — NEW) |
| PR.PS-02–03 (sw/hw lifecycle) | CC6.1, CC8.1 | A.5.14, A.7.13, A.7.14, A.8.6 |
| PR.PS-04 (logs) | CC7.1, CC7.2 | A.8.15 (logging) |
| PR.PS-05 (unauthorized sw) | CC6.1, CC6.6 | A.8.7 (malware protection), A.8.19 (installation) |
| PR.PS-06 (secure SDLC) | CC8.1 | A.8.25, A.8.27, A.8.28 (secure coding — NEW), A.8.29, A.8.30 |
| PR.IR-01 (network) | CC6.6, CC6.7 | A.8.20, A.8.21, A.8.22 (network controls and segregation) |
| PR.IR-02 (environmental) | CC6.4 (physical access) | A.7.1, A.7.5, A.7.8, A.7.11, A.7.12 |
| PR.IR-03–04 (resilience/capacity) | A1.2 (Availability), CC9.1 | A.5.30 (ICT readiness — NEW), A.8.6 (capacity), A.8.14 (redundancy) |

## Detect, Respond, Recover × SOC 2 + ISO 27001:2022 — detail

| CSF Subcategory | SOC 2 Common Criteria | ISO 27001:2022 Annex A |
|---|---|---|
| DE.CM-01–09 (continuous monitoring) | CC7.1, CC7.2, CC7.3 | A.8.15 (logging), A.8.16 (monitoring — NEW), A.8.17 (clock sync) |
| DE.AE-02–08 (event analysis) | CC7.2, CC7.3, CC7.4 | A.5.24, A.5.25 (incident assessment, decisions), A.8.16 |
| RS.MA-01–05 (incident management) | CC7.3, CC7.4, CC7.5 | A.5.24, A.5.25, A.5.26 (incident response) |
| RS.AN-03–08 (incident analysis) | CC7.4 | A.5.27 (lessons learnt), A.5.28 (evidence collection) |
| RS.CO-02–03 (notification) | CC7.4, CC2.3 | A.5.5 (contact with authorities), A.5.6 (special interest groups) |
| RS.MI-01–02 (containment, eradication) | CC7.4, CC7.5 | A.5.26 (response to incidents) |
| RC.RP-01–06 (recovery) | A1.2, A1.3, CC7.5 | A.5.29, A.5.30, A.8.13, A.8.14 |
| RC.CO-03–04 (recovery comms) | CC2.3, A1.3 | A.5.5, A.5.6, A.5.29 |

## Where the cross-framework dividend is largest

Three categories produce the highest cross-framework leverage — invest in these and the same evidence satisfies the largest number of obligations across the portfolio.

- PR.AA (Identity, Authentication, Access Control) — satisfies SOC 2 CC6.1–6.3 and CC6.6; ISO 27001 A.5.15, A.8.2–8.5; CMMC L2 AC and IA families; HIPAA §164.312(a) and (d). Highest evidence reuse of any category.
- PR.PS (Platform Security) — configuration, logging, secure SDLC. Satisfies SOC 2 CC8 and CC7; ISO 27001 A.8.9, A.8.15, A.8.25–8.28; CMMC L2 CM and AU; HIPAA §164.312(b) and §164.310(d).
- DE.CM + DE.AE (Detection) — a tuned SIEM with analytic rules and incident records satisfies SOC 2 CC7, ISO 27001 A.8.16, CMMC L2 AU and SI, HIPAA §164.308(a)(1)(ii)(D).

# CSF as queryable graph, not static reference.

This guide treats CSF v2.0 the way most organizations actually need to consume it: as a reference catalog of outcomes, mapped to operational signal, crosswalked to the other frameworks the organization answers to. That treatment is correct as far as it goes — but it stops short of where mature programs land. The mature treatment makes CSF a queryable substrate, not a static reference document.

## Pattern 1 — Subcategories as graph nodes, not spreadsheet rows

In the conventional pattern, CSF is a spreadsheet: 106 rows, one per Subcategory, with columns for Current state, Target state, owner, evidence, and notes. The spreadsheet is updated quarterly or before assessments and never queried in the operational sense. Every cross-framework crosswalk is its own spreadsheet. Every Profile is its own spreadsheet.

In the continuous-readiness pattern, every CSF Subcategory is a node in a Knowledge Graph. The same node connects to the SOC 2 criterion it crosswalks to, the ISO 27001 control it maps to, the CMMC practice it satisfies, the Microsoft signal source that produces evidence for it, the risk register entries it treats, the policy that governs it, and the implementation maturity assessment that scores it. A single update to one node — a new Conditional Access policy that strengthens PR.AA-03 — propagates simultaneously to SOC 2 CC6.1, ISO 27001 A.8.5, CMMC IA.L2-3.5.3, and HIPAA §164.312(d) without requiring four spreadsheet edits.

## Pattern 2 — Profiles as graph traversals

Current Profile, Target Profile, ransomware Community Profile, manufacturing Community Profile — each is a different traversal of the same underlying graph. A Current Profile is the projection of nodes onto their current implementation state. A Target Profile is the projection onto desired state. The gap between them is a graph diff. The action plan is the sequence of node updates required to close the diff.

This is how organizations operating at Tier 4 (Adaptive) actually function. They do not maintain Profiles as static documents; they maintain a single source of truth and project Profiles dynamically as needed. Kyūdō's Knowledge Graph operationalizes this pattern — Profile rendering is a query, not a deliverable.

## Pattern 3 — Microsoft signal at the subcategory level

Section 5 of this guide maps each high-frequency CSF Subcategory to a primary Microsoft signal source. In the conventional pattern, those signals are read at audit time, formatted into evidence packets, and stored. In the continuous-readiness pattern, those signals are read continuously — the state of every Conditional Access policy, every Defender for Endpoint alert, every Sentinel detection, every Purview DLP policy execution — and the affected Subcategories update their

state in real time. PR.AA-03 (‘users, services, and hardware are authenticated’) is not assessed quarterly; it is true or false right now, with the supporting evidence already collected.

This is the architectural property that enables ‘evidence is already true between audits’ — the Manifesto's central claim. The evidence does not need to be assembled because it was never disassembled.

> **One control set, every framework**
> The Secure Controls Framework (SCF) anchors the Kyūdō Knowledge Graph as the meta-framework substrate. SCF includes 1,470+ controls across 80+ frameworks. NIST CSF v2.0 maps to SCF; SCF maps to SOC 2, ISO 27001, CMMC, HIPAA, GDPR, EU AI Act, and the rest. A single Microsoft Defender for Cloud configuration baseline, a single Microsoft Sentinel detection record, a single Microsoft Purview DLP policy execution attests against every framework where the SCF crosswalk holds. Adding one more framework to the portfolio is a graph operation, not a program rebuild.

## Pattern 4 — Sovereignty as architecture

The conventional GRC architecture deploys a SaaS governance platform that ingests evidence from the customer's environment, processes it externally, and presents posture in the vendor's cloud. This requires data egress: control-state telemetry, risk register contents, and sometimes raw logs leave the customer's environment to be governed.

The continuous-readiness pattern inverts this. The governance layer deploys inside the customer's own security boundary — in regulated organizations on the Microsoft stack, this means inside the customer's Azure tenant. Microservices run in customer-owned AKS clusters with private endpoints, system-assigned managed identities, and customer-managed encryption keys. No governance data crosses the tenant boundary. The deployment model is the moat — no SaaS-first competitor can replicate it without re-architecting their platform.

For CSF v2.0 specifically, this matters because GV.SC (Cybersecurity Supply Chain Risk Management) becomes substantially easier to satisfy when the governance platform itself is not a third-party processor of in-scope data. The exclusion of governance data from the supplier-risk calculus is an architectural property, not a contractual claim.

## Pattern 5 — Auditor-defensible AI

AI in GRC is now a category-saturated claim. Most platforms position AI as a chat layer over documents. The continuous-readiness pattern requires AI that survives the auditor's next question: every AI-produced explanation, mapping, or recommendation must have a source, a confidence level, and a re-performable result.

In Kyūdō, AI is layered. Deterministic functions handle scoring, state transitions, Profile rendering, and crosswalk traversal. AI functions handle explanation, draft policy generation, control-mapping suggestions, and natural-language traversal of the Knowledge Graph. The two

layers never share a trust contract: the deterministic engine produces the answer, AI produces the prose. Where an auditor asks ‘how does PR.AA-03 trace to specific authentication events?’ the answer is a graph traversal — not a model output.

## Where this leaves you

If you are using CSF v2.0 as a communication framework with executives, regulators, or insurers, this Mapping Reference is sufficient. The Functions, Categories, Subcategories, Tiers, and Profiles enable the conversation. The crosswalks let you connect the conversation to your existing framework portfolio.

If you are operating CSF v2.0 in production — maintaining Current and Target Profiles, tracking subcategory-level posture continuously, and reporting CSF outcomes alongside SOC 2, ISO 27001, CMMC, and HIPAA — the architecture this section describes is the direction the practice is moving. The marginal cost of adding one more framework, one more Subcategory, or one more Microsoft signal source should approach zero. If it does not, the bottleneck is the architecture.

Kyūdō is the platform that makes that architecture available to regulated organizations running Microsoft 365 and Azure. The next step, if useful, is a deployment workshop in your tenant. The architecture brief is one click. The conversation is one email.

—

### If this is useful, the next step is concrete

Architecture briefing — a 30-minute walkthrough of the Kyūdō deployment in your Azure tenant: CSF Profile rendering, multi-framework crosswalk, evidence flow, and the sovereignty model. → hello@kyudo.ai

Controls workshop — 90 minutes mapping your current controls to CSF v2.0 with side-by-side SOC 2, ISO 27001, CMMC, and HIPAA views from a single control set. → kyudo.ai/workshop

Trust packet — our SOC 2 posture, ISO 27001 architecture commitments, NIST CSF self-attestation, data-residency statement, and the Microsoft estate dependency map. Available on request.

APPENDIX A · TIER INDICATORS

# Practical signals that your organization is at each Tier.

CSF Tiers describe the rigor of cybersecurity risk governance and management. The table below provides practical indicators — things you will observe in an organization at each Tier — to support self-assessment. These indicators are not from NIST; they are practitioner observations consistent with the Tier descriptions in NIST CSWP 29 Appendix B.

| Tier | Indicators you will observe |
|---|---|
| Tier 1: Partial | Risk decisions made by individuals on the spot. No risk register, or one that is years out of date. Cybersecurity policy may exist but is not enforced or known by personnel. Incidents handled by whoever is available, with no documented runbook. No vendor risk program. Awareness training is annual at best, often skipped. |
| Tier 2: Risk Informed | Risk register exists and is reviewed periodically by the security team. Cybersecurity policy is approved by senior management but not yet board-level. Incident response runbook exists; tabletop exercises occur but inconsistently. Vendor risk program exists at onboarding but not for ongoing monitoring. Cybersecurity is discussed at the executive level but not integrated with enterprise risk management. |
| Tier 3: Repeatable | Risk register is live, reviewed monthly, and connected to controls and exceptions. Cybersecurity policy is board-approved, reviewed annually, and enforced through automation where possible. Incident response runbook is tested at least annually with documented after-action reports. Vendor risk program covers onboarding, ongoing monitoring, and offboarding with evidence. Cybersecurity metrics flow up to the board through a defined cadence. The CISO is empowered and adequately resourced. |
| Tier 4: Adaptive | Risk management is dynamic and predictive — risks are identified, prioritized, and mitigated before they manifest as incidents. Cybersecurity is integrated into business decision-making, not consulted as a checkpoint. Incident response is exercised through realistic, recurring exercises (live-fire scenarios, red team) with continuous lessons-learned integration. Vendor risk monitoring is continuous and automated. Cybersecurity is treated by the board as a peer of financial, operational, and reputational risk. The organization is recognized as a leader in its sector or vertical for cybersecurity practice. |

APPENDIX B · AUTHORITATIVE SOURCES

# Where to verify and go deeper.

This guide reflects the official NIST publication of CSF v2.0 and current 2025–2026 practice. The references below are organized by primary, supplemental, and cross-framework sources.

## Primary — NIST CSF v2.0

- NIST CSWP 29, The NIST Cybersecurity Framework (CSF) 2.0, February 26, 2024. The authoritative publication. Available at nvlpubs.nist.gov.
- NIST SP 1299, NIST Cybersecurity Framework 2.0: Resource & Overview Guide. Companion document with introductory guidance.
- NIST CSF 2.0 Reference Tool — the interactive tool that publishes the Core, Informative References, and Implementation Examples in machine-readable formats. Hosted at csrc.nist.gov.
- NIST CSF Quick-Start Guides — brief actionable documents on specific CSF-related topics, updated independently of the main publication.

## Supplemental — NIST

- NIST SP 800-53 Rev 5.2.0, Security and Privacy Controls. Primary Informative Reference for CSF Subcategories.
- NIST SP 800-171 Rev 3.0, Protecting Controlled Unclassified Information.
- NIST SP 800-30 Rev 1, Guide for Conducting Risk Assessments.
- NIST SP 800-37 Rev 2, Risk Management Framework for Information Systems and Organizations.
- NIST SP 800-161 Rev 1, Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations — the deep reference for GV.SC.
- NIST IR 8286 series, Integrating Cybersecurity and Enterprise Risk Management.
- NIST AI Risk Management Framework (AI RMF), for organizations using AI within the CSF scope.
- NIST Privacy Framework, used together with CSF for privacy-relevant outcomes.

## Cross-framework

- AICPA TSP Section 100, 2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (with Revised Points of Focus — 2022). The authoritative SOC 2 criteria.
- ISO/IEC 27001:2022 — the certifiable ISMS standard.
- ISO/IEC 27002:2022 — implementation guidance for ISO 27001 Annex A controls.
- CMMC v2 Level 2 — NIST SP 800-171-derived controls for the U.S. defense industrial base.
- HIPAA Security Rule — 45 CFR Part 164, Subpart C.

- Secure Controls Framework (SCF) v2025.x — the meta-framework crosswalking 80+ frameworks including all of the above.

## Microsoft documentation

- Microsoft Service Trust Portal — published assessments, audit reports, and Customer Assurance Service documents for Microsoft cloud services.
- Microsoft Purview Compliance Manager — NIST CSF v2.0 assessment template with improvement actions mapped to Microsoft 365 and Azure.
- Microsoft Defender for Cloud regulatory compliance dashboard — includes NIST CSF baseline.
- Microsoft Cloud Adoption Framework — security baseline guidance aligned to NIST CSF outcomes.

APPENDIX C · GLOSSARY

# Terms used in this reference.

| Term | Definition |
|---|---|
| Adaptive (Tier 4) | The highest CSF Tier. The organization adapts cybersecurity practices based on previous and current activities, including lessons learned and predictive indicators. Cybersecurity is part of organizational culture. |
| Category | A group of related cybersecurity outcomes that collectively comprise a CSF Function. CSF v2.0 has 22 Categories. |
| Community Profile | A baseline of CSF outcomes addressing shared interests across multiple organizations — typically a sector, technology, or threat type. |
| Core | The hierarchy of Functions, Categories, and Subcategories that make up CSF v2.0. The Core describes desired outcomes; it does not prescribe how to achieve them. |
| Current Profile | An Organizational Profile specifying the Core outcomes the organization is currently achieving. |
| CSF | Cybersecurity Framework. The shorthand for the NIST Cybersecurity Framework. |
| Function | The highest-level grouping of cybersecurity outcomes in CSF. v2.0 has six Functions: Govern, Identify, Protect, Detect, Respond, Recover. |
| Govern (GV) | The CSF Function added in v2.0. Establishes, communicates, and monitors the cybersecurity risk management strategy, expectations, and policy. |
| Implementation Example | A notional, action-oriented step that helps an organization achieve a Subcategory outcome. Examples are illustrative, not prescriptive or exhaustive. |
| Informative Reference | A mapping that indicates relationships between CSF outcomes and specific controls, requirements, or guidance in other publications (SP 800-53, ISO 27001, CIS Controls, etc.). |
| NIST | National Institute of Standards and Technology. The U.S. federal agency that publishes the CSF and related cybersecurity standards. |
| NIST CSWP 29 | The official publication of CSF v2.0, dated February 26, 2024. |

| Term | Definition |
|---|---|
| Organizational Profile | A description of an organization's current and/or target cybersecurity posture in terms of CSF Core outcomes. |
| Partial (Tier 1) | The lowest CSF Tier. Risk strategy is managed ad hoc; cybersecurity risk awareness is limited. |
| Profile | Short form of Organizational Profile. |
| Repeatable (Tier 3) | The third CSF Tier. Risk management practices are formally approved and expressed as policy. Risk-informed policies, processes, and procedures are defined and reviewed. |
| Risk Informed (Tier 2) | The second CSF Tier. Risk management practices are approved by management but may not be organization-wide policy. |
| Subcategory | The most specific level of the CSF Core. Each Subcategory describes a detailed outcome that supports a Category. CSF v2.0 has 106 Subcategories. |
| Target Profile | An Organizational Profile specifying the desired outcomes the organization has selected and prioritized for achieving its risk management objectives. |
| Tier | A characterization of the rigor of an organization's cybersecurity risk governance and management practices. CSF v2.0 retains the four Tiers from v1.1: Partial, Risk Informed, Repeatable, Adaptive. |
| v1.1 | The previous version of the CSF, published in 2018. Superseded by v2.0 but not formally invalidated; legacy programs may continue to reference v1.1. |
| v2.0 | The current version of the CSF, published February 26, 2024. |