THOUGHT LEADERSHIP

OPERATIONALIZING AI COMPLIANCE

LEVERAGING NIST AI RMF TO SUPPORT ISO 42001 & GLOBAL COMPLIANCE

CO-AUTHORED BY

schellman + solidcore.ai + carahsoft

EXECUTIVE SUMMARY

Enterprise adoption of generative artificial intelligence has outpaced organizations' ability to demonstrate that it is governed, and regulators have now defined what demonstration means.

The EU AI Act's high-risk obligations, including Article 12's requirement for continuous, producible logging of system operation following the Digital Omnibus amendments adopted in June 2026. In the United States, NIST AI RMF alignment is an emerging expectation in federal procurement and a reference point for state law and sector regulators.

The question facing security and compliance leaders is no longer whether AI governance policies exist, but whether their operation can be evidenced.

Two frameworks anchor that demonstration. ISO/IEC 42001 , the first certifiable AI management system standard, defines what an organization must establish; the NIST AI Risk Management Framework (AI RMF 1.0, together with the NIST AI 600-1 Generative AI Profile) as well as other technical frameworks provide the operational model for how.

APPROACH AND FINDINGS

SolidCore is a containerized application deployed within the customer's own AWS environment , not SaaS. It integrates read-only with cloud service APIs, so data never leaves the environment boundary and existing inference pipelines are unmodified. Its scope is first-party generative AI systems operating in AWS. Schellman examined SolidCore.ai's product capabilities against the AI RMF's Govern, Map, Measure, and Manage functions. Appendix A presents the resulting mapping capability by capability, with each rating validated as full, partial, or not evaluated. Appendix B provides a crosswalk to ISO 42001 Annex A controls to support audit planning.

INTENDED USE

This white paper is a representative example of the assessment pattern Schellman sees often: anchor governance in ISO 42001 certification and leverage more technical implementation frameworks like NIST AI RMF for the operationalization of AI compliance. The paper is intended for CISOs, security leaders, AI governance leads, and internal audit teams building evidence-based AI compliance programs and for the external assessors who will evaluate them. It goes without saying that compliance with regulations and associated industry standards like these soley reside with the company implementing them. The goal is this analysis is to show how SolidCore can support those requirements.

Schellman performed an independent review of SolidCore.ai's practices as reflected in this white paper. Schellman holds no financial or other interest in SolidCore.ai and receives no benefit contingent on the conclusions presented herein. The information provided is "AS IS" without warranties of any kind. Schellman and SolidCore.ai expressly disclaim any warranties or representations, including implied warranties of fitness for a particular purpose, and express no opinion as to SolidCore.ai's suitability for any specific use case.

THE REGULATORY MOMENT WE'RE IN

A converging set of global frameworks has established what responsible AI governance looks like - on paper and in process. The challenge now is demonstrating compliance operationally, across a regulatory landscape that is both broader and more specific than most organizations anticipated.

EU AI ACT - RISK-TIERED OBLIGATIONS WITH ACTIVE ENFORCEMENT

The EU AI Act establishes a four-tier risk classification - unacceptable, high, limited, and minimal - with obligations scaled to the potential for harm. Prohibited practices (social scoring, manipulative AI, certain real-time biometric identification) have been banned since February 2025. Obligations for providers of general-purpose AI models, including LLMs, took effect on August 2, 2025, requiring technical documentation, transparency disclosures, copyright compliance measures, and - for models with systemic risk - adversarial testing and incident reporting.

In June 2026, the EU formally adopted the Digital Omnibus on AI, the first set of amendments to the Act, which reset the compliance timeline for high-risk AI systems.

Obligations for stand-alone high-risk systems under Annex III - covering AI used in employment decisions, credit scoring, education assessments, essential services, and law enforcement - now apply from December 2, 2027, and obligations for AI embedded in regulated products under Annex I (medical devices, machinery, and similar) apply from August 2, 2028. The amendments enter into force upon publication in the Official Journal, expected in July 2026. Not everything moved: Article 50 transparency obligations still apply from August 2, 2026, a new Article 5 prohibition on AI-generated non-consensual intimate imagery takes effect December 2, 2026 (alongside the watermarking deadline for legacy generative systems), and penalties for non-compliance still reach €35 million or 7 percent of global annual turnover.

Article 12 is the provision most directly relevant to this paper's thesis. It requires providers of high-risk AI systems to implement automatic logging capabilities that record system operation throughout the AI lifecycle - including input data, system outputs, and the operation of human oversight mechanisms. These logs must be retained, accessible, and producible for regulatory authorities. In practice, Article 12 creates a statutory requirement for exactly the kind of continuous, tamper-evident audit trail that most organizations have not yet built. The deferral moves the deadline, not the obligation. The EU postponed enforcement precisely because most organizations lacked the supporting infrastructure to comply - and the logging, inventory, and evidence architecture Article 12 demands is the kind that takes quarters to build, not weeks. Organizations that treat December 2027 as distant will find themselves in 2026's position all over again.

ISO/IEC 42001 - THE GLOBAL AI MANAGEMENT SYSTEM STANDARD

ISO/IEC 42001 is the first internationally certifiable management system standard for artificial intelligence, published by ISO in December 2023. Structured around 10 clauses and 38 Annex A controls, it requires organizations to establish an AI Management System (AIMS) that covers the full lifecycle - from data sourcing and model development through deployment, monitoring, and decommissioning.

What matters beyond the certification itself is what the standard requires operationally. ISO 42001 mandates a systems-level risk assessment (Clause 6) that extends beyond model performance to encompass data pipelines, downstream business processes, and human oversight mechanisms. It requires documented accountability structures with assigned roles and escalation paths (Clause 5). And it requires organizations to produce evidence that controls are operating as designed - not merely that they exist.

Schellman was the first US (ANAB) accredited ISO 42001 certification body and has issued AIMS certificates for scores of companies including AWS's initial certification along with most of the major model providers.

Across our assessments, the most consequential controls are the ones organizations most frequently underestimate: AI system impact assessments that account for affected stakeholders beyond the immediate user, human oversight procedures that are operational rather than aspirational, and monitoring mechanisms that produce auditable records of system behavior in production.

For the CISO, the standard integrates naturally with ISO 27001 - both share the Annex SL management system structure - but it adds AI-specific dimensions around data governance, model transparency, and algorithmic fairness that security controls alone do not address.

ISO 42001

NIST AI RMF – A CORE OPERATIONAL MODEL

The NIST AI Risk Management Framework (AI RMF 1.0), published in January 2023, organizes AI governance into four functions - Govern, Map, Measure, and Manage - spanning 19 categories and 72 subcategories. Unlike ISO 42001, which is a certifiable management system standard, the AI RMF is a voluntary, outcome-based framework designed to translate risk management principles into operational practice. In effect, ISO 42001 defines the 'what' of an AI management system; the NIST AI RMF provides a detailed operational model for 'how.'

The framework's significance extends well beyond its voluntary designation. OMB policy memoranda have made NIST AI RMF alignment functionally required for federal agencies and increasingly for the contractors that serve them. Federal procurement solicitations now routinely reference AI RMF compliance, and agencies are flowing governance requirements down through contract modifications and new acquisitions. For government contractors - particularly those operating under CMMC 2.0 and DFARS - the AI RMF's Govern and Manage functions map directly to existing access control, audit trail, and monitoring obligations, making an AI RMF-aligned program an efficient path to satisfying multiple federal requirements simultaneously.

Beyond the federal context, the framework is gaining traction as a de facto standard of care. Like the NIST Cybersecurity Framework, while developed by the government, the standard was designed to be utilized across different industries.

The Colorado AI Act explicitly references recognized risk management frameworks for safe harbor protection. Sector regulators - including the OCC, FDIC, FDA, and FTC - are increasingly citing NIST AI RMF principles in their expectations for responsible AI deployment. For organizations operating across multiple regulatory regimes, the framework provides a common operational layer beneath jurisdiction-specific compliance requirements.

NIST released the Generative AI Profile (NIST AI 600-1) in July 2024, adding 12 risk categories specific to LLM and generative AI deployments, and continues to expand its guidance - including a new Critical Infrastructure AI profile released in April 2026. AI RMF 1.0 itself is also under revision, per the 2025 White House AI Action Plan's direction to NIST. Organizations building AI governance programs today should treat the AI RMF as an evolving operational baseline, not a one-time implementation.

ADDITIONAL FRAMEWORKS

It is noteworthy that there are additional frameworks for consideration when evaluating AI governance and implementation models. These include but are not limited to:

AGENTIC / AI-AGENT-SPECIFIC

- **AIUC-1** — https://www.aiuc-1.com/
- **OWASP Top 10 for Agentic Applications** — https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/
- **OWASP AI Vulnerability Scoring System (AIVSS)** — https://aivss.owasp.org/
- **CSA AI Controls Matrix (AICM)** — https://cloudsecurityalliance.org/artifacts/ai-controls-matrix-v1-1

LLM / MODEL-SECURITY SPECIFIC

- **OWASP Top 10 for LLM Applications** — https://genai.owasp.org/llm-top-10/
- **MITRE ATLAS** — https://atlas.mitre.org/

VENDOR/INDUSTRY TECHNICAL FRAMEWORKS

- **Google Secure AI Framework (SAIF)** — https://saif.google/
- **IBM AI Risk Atlas** — https://www.ibm.com/docs/en/watsonx/saas?topic=ai-risk-atlas

SCOPE AND APPROACH

The Artificial Intelligence Risk Management Framework (AI RMF 1.0 & NIST AI 600-1) is is organized around four core functions that work together as a continuous cycle.SolidCore designed its solution with NIST AI RMF as technological backdrop, as well as best practices across security and compliance standards. Within NIST AI RMF, individual product policies (e.g. rules, risks, and configurations) include any relevant mappings to the four functions of the NIST AI RMF:

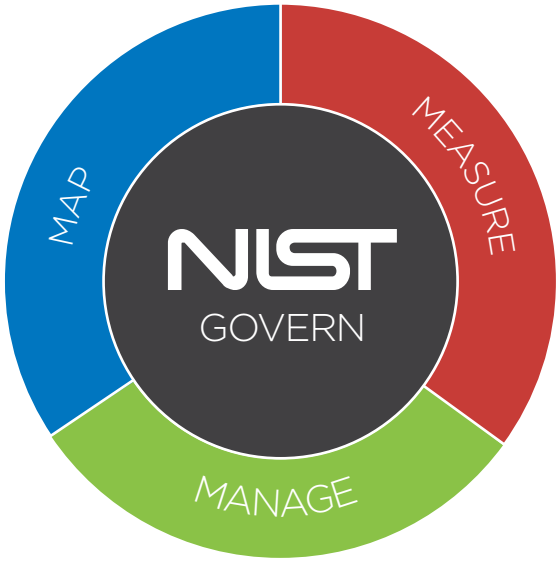

SolidCore's goal to operationalize AI compliance resulted in a number of policy objectives built into the solution through specific capabilities. These may use existing cloud-native capabilities such as AWS's Bedrock or CloudTrail. Together, they provide on-demand, legally defensible, NIST-aligned, actionable compliance reporting and audit trails.

Each policy objective is mapped to a severity level to reflect the potential impact of non-compliance. Data is collected via read-only integration with AWS cloud services APIs, requiring no modification to your existing resources or inference pipelines. These policy objectives and solution architecture provide an indispensable tool to demonstrate that the GAI systems are complying with NIST guidelines.

SolidCore's solution can be implemented in less than half an hour. Since it resides in the cloud tenant where it is deployed, it does not need proxies and data never leaves the environment's boundaries. Scanning model invocation logs, SolidCore sees the data after the fact, which does not prevent or stop an AI workflow. It performs like an observability layer and derives its data by using prorietary service and log integrations.

For each policy objective, the mapping to NIST AI RMF 1.0 was examined by Schellman. The categories of People, Process, and Technology are subjective labels developed by SolidCore. The breakdown below compares the requirements that NIST AI RMF requires and the product capabilities that the SolidCore solution offers. In most cases, this comes through as a product objective that can be enabled. For others, the capability may be more architectural or provide an input into a future workflow (especially those with a Process label).

The key concepts to SolidCore's solution include:

- Projects
- System of Record
- Knowledge Base
- Test, Evaluation, Validation, and Verification (TEVV)
- Metadata from API calls
- Logging

NIST GOVERN

AI RMF's Govern function is the core profile that establishes governance. This cross-cutting function cannot be isolated. It offers a structure by which the other risk management functions can be aligned with an organization's principles, policies, and strategic priorities.

By successfully implementing and continually executing the governance function, organizations can more effectively drive and enhance internal practices and norms throughout the lifespan of an AI system. This includes adapting to the ever-evolving needs, cultures, expectations and knowledge of a risk management framework.

SOLIDCORE PRODUCT CAPABILITIES

SolidCore gives internal audit and AI governance teams the ability to evaluate AI compliance without involving engineering. The solution has the ability out-of-the-box to support the Govern category and it will not interrupt existing AI workflows (or changes to those workflows).

The platform creates a centralized inventory of AI projects (think GAI systems such as agents, applications or LLM gateways), models (that operate in the cloud environment), and associated resources (such as knowledge bases, guardrails/content filters). This supports governance activities such as AI asset management, lifecycle tracking, and ongoing monitoring. With in-product compliance mappings, technical controls, and risk dashboards, SolidCore can support risk-informed decision making and improve AI ecosystem transparency.

Using the concept of Projects to create units to track GAI usage, SolidCore can embed technical controls, monitor policy compliance, and track AI system status throughout its life. This includes features to configure log retention, LLM usage monitoring, and continuous assessment of AI risks.

While several capabilities partially support NIST AI RMF Govern outcomes because organizational policies, processes, and human oversight remain critical.

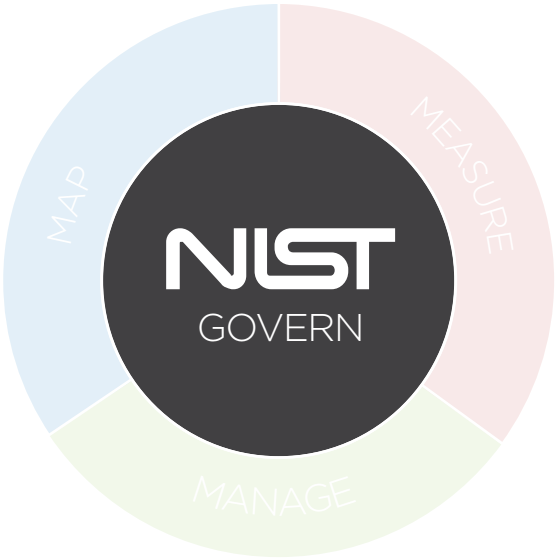

schellman + solidcore.ai + carahsoft

NIST MAP

Risk management is difficult to perform when lacking contextual knowledge and awareness of the risks within the identified contexts. The purpose of the AI RMF's Map function is to enhance an organization's ability to identify not only risks but also their contributing factors.

Throughout the lifecycle of AI systems, their context, capabilities, risks, benefits, and potential impacts will continually evolve. Therefore, it is incumbent on Framework users to establish Mapping as an ongoing process.

SOLIDCORE PRODUCT CAPABILITIES

SolidCore enables organizations to understand how AI systems are built, what data they rely upon, and how those systems operate throughout their lifecycle. Without requiring engineering teams to modify existing AI workflows, the platform automatically creates a System of Record that inventories AI models, inference activity, data sources, retrieval mechanisms, and supporting resources such as Knowledge Bases used in Retrieval-Augmented Generation (RAG). This visibility helps governance and internal audit teams map AI systems, understand their dependencies, and establish context for AI risk assessments.

The platform captures technical metadata about AI systems, including the models invoked during inference, associated Knowledge Bases, data vectors, and connected data sources. These capabilities provide organizations with the information needed to understand AI system behavior, identify knowledge boundaries, document third-party dependencies, and support technology evaluations such as testing, evaluation, verification, and validation (TEVV). By maintaining detailed operational records and system relationships, SolidCore helps organizations establish a comprehensive understanding of their AI environment and the factors that influence model outputs.

SolidCore also provides the operational insight needed to support AI risk identification through detailed logging, policy-based issue tracking, and assignment of governance responsibilities. Users can monitor AI interactions, examine privacy and third-party data risks, investigate issues associated with technical controls, and use captured operational data to refine ongoing monitoring activities. While organizational judgment remains necessary to assess business impacts, determine risk tolerance, and document AI risks, SolidCore provides the technical evidence and system visibility needed to support the NIST AI RMF Map function.

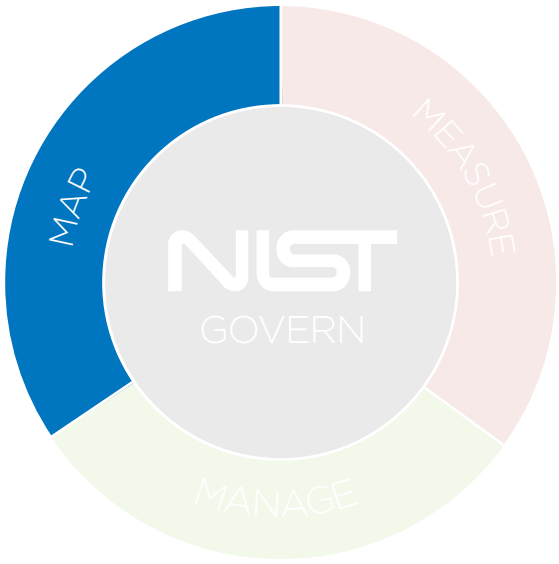

NIST MEASURE

Organizations are expected to analyze, assess, benchmark, and monitor AI risk and related impacts. The Measure function necessitates rigorous software testing and performance assessment methodologies to determine uncertainties, compare performance benchmarks and document results through formalized reporting.

Metrics and measurement methodologies should lead to objective, repeatable, or scalable Test, Evaluation, Validation, and Verification (TEVV) processes that then can be followed and documented.

SOLIDCORE PRODUCT CAPABILITIES

SolidCore enables organizations to measure AI system performance by providing the operational data and technical evidence needed to evaluate AI systems throughout their lifecycle. Through its System of Record, the platform captures detailed logging, inference activity, applied policies, and runtime metrics such as inference latency and token consumption. These capabilities provide governance, risk, and internal audit teams with objective information to monitor AI system behavior, document testing and evaluation activities (TEVV), and assess whether AI systems continue to perform as intended in production.

The platform also supports ongoing measurement by providing project-level reporting, policy compliance information, and visibility into AI interactions that may introduce operational or privacy risks. Organizations can use this information to review AI system performance with relevant stakeholders, identify performance trends, document privacy-related observations, and continuously refine monitoring activities as AI systems evolve.

Measurement is a key component in the detection and response to AI risks, both in real time and over the lifecycle of your organization's program. SolidCore offers custom views that prioritize behaviors an organization wants to monitor (as defined by the outcomes from the previous Map function). In parallel, SolidCore institutes filters, guardrails, and compensating controls at a granular level to adapt an organization's privacy policies to desired projects.

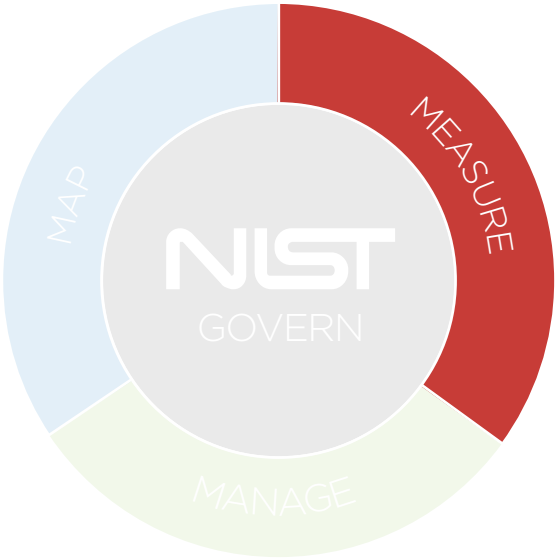

schellman + solidcore.ai + carahsoft

NIST MANAGE

To comply with the Manage function, organizations should prioritize allocating risk resources to mapped and measured risks on a regular basis and as defined by the Govern function. Risk treatment comprises plans to respond to, recover from, and communicate about incidents or events.

Relevant AI actors within an organization should offer inputs to generate, plan, prepare, implement, and document strategies that will maximize AI benefits and minimize the magnitude or likelihood of potential negative impacts.

SOLIDCORE PRODUCT CAPABILITIES

SolidCore tracks stages of Projects to log system details and capture them in the System of Record for monitoring AI systems post-deployment. In tandem, SolidCore evaluates the cloud architecture and isolation guardrails to identify content from third parties that need to be reviewed before delivering it to an LLM.

Without disrupting existing AI workflows, the platform tracks AI Projects, monitors pre-trained model versions, captures production activity within its System of Record, and maintains visibility into AI system configuration and usage. Integrated compliance mappings, technical controls, and configurable policy evaluations help organizations align operational activities with internal governance requirements while supporting ongoing monitoring and maintenance of AI systems.

The platform also provides the operational insight needed to prioritize and respond to AI risks as they emerge. Risk trends dashboards, project-level reporting, and policy-based issue tracking help governance and internal audit teams identify and prioritize risks based on observed system activity. SolidCore tracks unauthorized access to AI systems, captures production telemetry, and leverages guardrails provided by supported AI platforms to evaluate interactions with third-party models and resources.

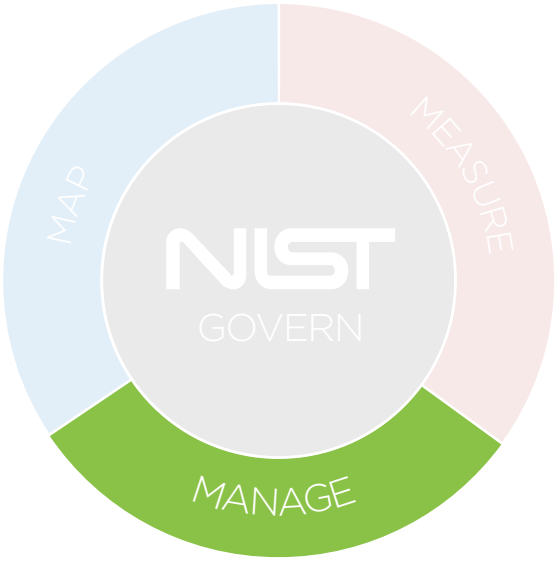

APPENDIX A: SOLIDCORE CAPABILITIES — NIST

The fields can be interpreted as follows:

- **CAPABILITY NUMBER:** Unique Identifier for each SolidCore capability
- **CAPABILITY STATEMENT:** Text describing the feature/capability/architecture
- **NIST AI RMF MAPPING:** Compliance mapping to NIST AI RMF 1.0
- **MAPPING Y/N/PARTIAL:** Mapping between NIST AI RMF and SolidCore's Capability (values can be Yes, None, or Partial)
- **TYPE (PPT):** Designation of People, Process, and/or Technology objective of the requirement.

NIST AI RMF	MAPPING	TYPE
GOVERN 1.1	Partial	Process
GOVERN 1.2	Partial	Technology
GOVERN 1.3	Partial	Process, People
GOVERN 1.4	Partial	Technology
GOVERN 1.5	Partial	Technology, Process
GOVERN 1.6	Yes	Technology
GOVERN 1.7	Partial	Process
GOVERN 2.1	Partial	People
GOVERN 2.2	None	Not Evaluated
GOVERN 2.3	None	Not Evaluated
GOVERN 3.1	None	Not Evaluated
GOVERN 3.2	None	Not Evaluated
GOVERN 4.1	Partial	Process
GOVERN 4.2	Partial	Process
GOVERN 4.3	Partial	Technology
GOVERN 5.1	None	Not Evaluated
GOVERN 5.2	None	Not Evaluated
GOVERN 6.1	None	Not Evaluated
GOVERN 6.2	None	Not Evaluated

NIST AI RMF	MAPPING	TYPE
MAP 1.1	None	Not Evaluated
MAP 1.2	None	Not Evaluated
MAP 1.3	None	Not Evaluated
MAP 1.4	None	Not Evaluated
MAP 1.5	None	Not Evaluated
MAP 1.6	None	Not Evaluated
MAP 2.1	Partial	Technology
MAP 2.2	Partial	Technology, Process
MAP 2.3	Partial	Technology
MAP 3.1	None	Not Evaluated
MAP 3.2	None	Not Evaluated
MAP 3.3	None	Not Evaluated
MAP 3.4	None	Not Evaluated
MAP 3.5	None	Not Evaluated
MAP 4.1	Partial	Process
MAP 4.2	None	Not Evaluated
MAP 5.1	Partial	Process
MAP 5.2	None	Not Evaluated
MEASURE 1.1	None	Not Evaluated
MEASURE 1.2	None	Not Evaluated
MEASURE 1.3	None	Not Evaluated
MEASURE 2.1	Partial	People, Process
MEASURE 2.2	None	Not Evaluated
MEASURE 2.3	None	Not Evaluated
MEASURE 2.4	Partial	People, Process
MEASURE 2.5	None	Not Evaluated
MEASURE 2.6	None	Not Evaluated
MEASURE 2.7	None	Not Evaluated
MEASURE 2.8	None	Not Evaluated

NIST AI RMF	MAPPING	TYPE
MEASURE 2.9	None	Not Evaluated
MEASURE 2.10	Partial	People, Process
MEASURE 2.11	None	Not Evaluated
MEASURE 2.12	None	Not Evaluated
MEASURE 2.13	None	Not Evaluated
MEASURE 3.1	None	Not Evaluated
MEASURE 3.2	None	Not Evaluated
MEASURE 3.3	None	Not Evaluated
MEASURE 4.1	None	Not Evaluated
MEASURE 4.2	Partial	People, Process
MEASURE 4.3	Partial	Technology, Process
MANAGE 1.1	None	Not Evaluated
MANAGE 1.2	Partial	People, Process
MANAGE 1.3	None	Not Evaluated
MANAGE 1.4	None	Not Evaluated
MANAGE 2.1	None	Not Evaluated
MANAGE 2.2	None	Not Evaluated
MANAGE 2.3	None	Not Evaluated
MANAGE 2.4	Partial	Technology
MANAGE 3.1	Yes	Technology
MANAGE 3.2	Yes	Technology
MANAGE 4.1	Partial	Technology
MANAGE 4.2	None	Not Evaluated
MANAGE 4.3	None	Not Evaluated

CAPABILITY #	CAPABILITY STATEMENT	NIST AI RMF	MAPPING	TYPE
1	Installing SolidCore into the AWS VPC provides an observability platform for the organization to establish a Source of Record (SoR). The SoR lists all the cloud resources involved in AI usage. AI is monitored through system policies that state the manner to operate AI for the organization.	**GOVERN 1.6:** Mechanisms are in place to inventory AI systems and are resourced according to organizational risk priorities.	Yes	Technology
2	View a mapping to NIST AI RMF via the Compliance page. This contains mappings to Policies, Issues found and controls in-place.	**GOVERN 1.1:** Legal and regulatory requirements involving AI are understood, managed, and documented.	Partial	Process
3	An AI inventory and usage of AI can be tracked as Projects. The system will capture several technical controls.	**GOVERN 1.5:** Ongoing monitoring and periodic review of the risk management process and its outcomes are planned and organizational roles and responsibilities clearly defined, including determining the frequency of periodic review.	Partial	Technology
4	The Risk Trends dashboard aggregates risks based on their designation (Total, Critical, High, Medium, Low) based on each Project. This information is reported over a 30-day period. It can be sorted by Projects, Risk Categories, and Time period.	**GOVERN 1.3:** Processes, procedures, and practices are in place to determine the needed level of risk management activities based on the organization's risk tolerance.	Partial	People, Process
5	Technical controls such as reviewing the retention period for LLM logs. This setting can be evaluated against the organization's AI activity data retention policies via the Log Retention Policy.	**GOVERN 1.4:** The risk management process and its outcomes are established through transparent policies, procedures, and other controls based on organizational risk priorities	Partial	Technology

CAPABILITY #	CAPABILITY STATEMENT	NIST AI RMF	MAPPING	TYPE
6	Technical controls are represented in SolidCore as Policy. These can be enabled on a Project and any Issues identified are associated with the respective Policy. These can assist with the practice of defining characteristics of trustworthy AI and monitoring AI applications based on this baseline.	**GOVERN 1.2:** The characteristics of trustworthy AI are integrated into organizational policies, processes, procedures, and practices.	Partial	Technology
7	Monitor projects and their AI model usage to evaluate the status of AI systems that are being phased out (release stage status drop downs include: development, preview, and production).	**GOVERN 1.7:** Processes and procedures are in place for decommissioning and phasing out AI systems safely and in a manner that does not increase risks or decrease the organization's trustworthiness.	Partial	Process
8	Roles and responsibilities can be assigned to SolidCore that align to the governance team to map, measure, and manage AI risks using the technical controls as feedback.	**GOVERN 2.1:** Roles and responsibilities and lines of communication related to mapping, measuring, and managing AI risks are documented and are clear to individuals and teams throughout the organization.	Partial	People
9	Projects follow a lifecycle process and after each change, the date of that review for the corresponding stage is captured.	**GOVERN 4.1:** Organizational policies and practices are in place to foster a critical thinking and safety-first mindset in the design, development, deployment, and uses of AI systems to minimize potential negative impacts. **GOVERN 4.2:** Organizational teams document the risks and potential impacts of the AI technology they design, develop, deploy, evaluate, and use, and they communicate about the impacts more broadly.	Partial Partial	Process

CAPABILITY #	CAPABILITY STATEMENT	NIST AI RMF	MAPPING	TYPE
10	SolidCore provides an inventory of LLM based on usage and monitor it for safe, secure, and reasonable usage guidelines through the lens of Projects.	**GOVERN 4.3:** Organizational practices are in place to enable AI testing, identification of incidents, and information sharing.	Partial	Process
11	SolidCore information around Project risks can be used to treat them in a prioritized fashion through a manual process where the system reports serve as an input.	**MANAGE 1.2:** Treatment of documented AI risks is prioritized based on impact, likelihood, and available resources or methods.	Partial	People, Process
12	SolidCore is monitoring for inconsistent user, primarily driven by unauthorized access to the AI systems represented as Projects.	**MANAGE 2.4:** Mechanisms are in place and applied, and responsibilities are assigned and understood, to supersede, disengage, or deactivate AI systems that demonstrate performance or outcomes inconsistent with intended use.	Partial	Technology
13	SolidCore uses the third-party guardrails established by the AI platforms. These are used to evaluate content from third-party resources and can be monitored by the user.	**MANAGE 3.1:** AI risks and benefits from third-party resources are regularly monitored, and risk controls are applied and documented. **GOVERN 6.1:** Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights.	Yes Partial	Technology Technology
14	SolidCore monitors the pre-trained AI models capturing the version of the model and logging other information visible throught the System of Record interface.	**MANAGE 3.2:** Pre-trained models which are used for development are monitored as part of AI system regular monitoring and maintenance.	Yes	Technology

CAPABILITY #	CAPABILITY STATEMENT	NIST AI RMF	MAPPING	TYPE
15	SolidCore tracks stages of Projects and "Production" includes logging system details and capturing them in the System of Record for monitoring AI systems post-deployment.	**MANAGE 4.1:** Post-deployment AI system monitoring plans are implemented, including mechanisms for capturing and evaluating input from users and other relevant AI actors, appeal and override, decommissioning, incident response, recovery, and change management.	Partial	Technology
16	SolidCore interfaces with the Knowledge Base resource that can specify tasks and methods that are implemented to support the AI system	**MAP 2.1:** The specific tasks and methods used to implement the tasks that the AI system will support are defined (e.g., classifiers, generative models, recommenders).	Partial	Technology
17	SolidCore interfaces with the Knowledge Base and this is captured in the System of Record, which can be used to understand the system's knowledge limits.	**MAP 2.2:** Information about the AI system's knowledge limits and how system output may be utilized and overseen by humans is documented. Documentation provides sufficient information to assist relevant AI actors when making decisions and taking subsequent actions.	Partial	Technology, Process
18	SolidCore captures the type of model used in the inference call. This runtime inference of the model can assist in TEVV considerations such as the data points used to enrich the analysis. Insights captured in the Data Vector and RAG inputs can also assist the user monitoring the system for TEVV considerations.	**MAP 2.3:** Scientific integrity and TEVV considerations are identified and documented, including those related to experimental design, data collection and selection (e.g., availability, representativeness, suitability), system trustworthiness, and construct validation.	Partial	Technology

CAPABILITY #	CAPABILITY STATEMENT	NIST AI RMF	MAPPING	TYPE
19	The System of Record captures the data sources and this can be used to detect when inputs/outputs are violating a third-party risk controls.	**MAP 4.1:** Approaches for mapping AI technology and legal risks of its components – including the use of third-party data or software – are in place, followed, and documented, as are risks of infringement of a third party's intellectual property or other rights.	Partial	Process
20	The total number of issues are tallied around each control. The user can select the Issues to drill-down and see the Policy items that are associated with the control.	**MAP 5.1:** Likelihood and magnitude of each identified impact (both potentially beneficial and harmful) based on expected use, past uses of AI systems in similar contexts, public incident reports, feedback from those external to the team that developed or deployed the AI system, or other data are identified and documented.	Partial	Process
21	A list of applied Policies to each Project can be used to document the TEVV.	**MEASURE 2.1:** Test sets, metrics, and details about the tools used during TEVV are documented.	Partial	People, Process
22	SolidCore has a variety of data sources that it can log based on the AI system and inference call. These can be examined by the user to monitor privacy risks of the AI system.	**MEASURE 2.10:** Privacy risk of the AI system – as identified in the MAP function – is examined and documented.	Partial	People, Process
23	Supporting the System of Record is raw logging/data that can be used to refine what is needed to support monitoring AI system behavior.	**MEASURE 2.4:** The functionality and behavior of the AI system and its components – as identified in the MAP function – are monitored when in production.	Partial	People, Process

CAPABILITY #	CAPABILITY STATEMENT	NIST AI RMF	MAPPING	TYPE
24	SolidCore reports on a Project basis can be used as input into the review and measurement of the system. This information could be used in compiling system performance and be reviewed by relevant AI actors to validate the system is performing consistently as intended.	**MEASURE 4.2:** Measurement results regarding AI system trustworthiness in deployment context(s) and across the AI lifecycle are informed by input from domain experts and relevant AI actors to validate whether the system is performing consistently as intended. Results are documented.	Partial	People, Process
25	SolidCore captures some performance values such as the inference latency and the token volume consumed with the inference.	**MEASURE 4.3:** Measurable performance improvements or declines based on consultations with relevant AI actors, including affected communities, and field data about context relevant risks and trustworthiness characteristics are identified and documented.	Partial	Technology, Process

APPENDIX B: SOLIDCORE CAPABILITIES TO ISO

This content is provided as a crosswalk. This white paper did not delve into the details around ISO 42001. However, due to the common request for ISO 42001 certification, this mapping may may assist in situations where an internal or external auditor is looking to better understand how SolidCore product capabilities might crosswalk to ISO 42001.

CAPABILITY #	CAPABILITY STATEMENT	ISO 42001 Annex A
1	Installing SolidCore into the AWS VPC provides an observability platform for the organization to retain an inventory of AI systems (Projects) and the resources they access (Policies).	A.2.2 Roles & responsibilities A.6.1 AI system impact assessment
2	View a mapping to NIST AI RMF via the Compliance screen. This contains mappings between Issues and controls.	A.2.6 Objectives A.9.1 Monitoring & measurement
3	An AI inventory and usage of AI can be tracked as Projects. The system will capture several technical controls.	A.6.1 AI system impact assessment A.9.1 Monitoring & measurement
4	The Risk Trends dashboard aggregates risks based on their designation (Total, Critical, High, Medium, Low) based on each Project. This information is reported over a 30 day period. It can be sorted by Projects, Risk Categories, and Period.	A.9.1 Monitoring & measurement A.2.6 Objectives
5	Technical controls such as reviewing the retention period for LLM logs. This setting can be evaluated against the organization's data retention policies via the Log Retention Policy.	A.7.4 Documentation A.9.1 Monitoring & measurement
6	Technical controls are represented in SolidCore as Policy. These can be enabled on a Project and any Issues identified are associated with the respective Policy. These can assist with the practice of defining characteristics of trustworthy AI and monitoring AI applications based on this baseline.	A.2.2 Roles & responsibilities A.4.1 Responsible AI policy
7	Monitor projects and their AI model usage to evaluate the status of AI systems that are being phased out (release stage status drop downs include: development, preview, and production).	A.6.2 AI system lifecycle

CAPABILITY #	CAPABILITY STATEMENT	ISO 42001 Annex A
8	Roles and responsibilities can be assigned to SolidCore that align to the governance team to map, measure, and manage AI risks using the technical controls as feedback.	A.2.2 Roles & responsibilities A.2.1 Leadership & commitment
9	Projects follow a lifecycle process and after each change, the date of that review for the corresponding stage is captured.	A.2.2 Roles & responsibilities A.7.4 Documentation
10	SolidCore provides an inventory of LLM based on usage and monitor it for safe, secure, and reasonable usage guidelines through the lens of Projects.	A.9.1 Monitoring & measurement A.10.1 Incident response
11	SolidCore information around Project risks can be used to treat them in a prioritized fashion through a manual process where the system reports serve as an input.	A.6.1 AI system impact assessment A.9.1 Monitoring & measurement
12	SolidCore is monitoring for inconsistent use, primarily driven by unauthorized access to the AI systems represented as Projects.	A.8.5 Human oversight A.9.1 Monitoring & measurement
13	SolidCore uses the third party guardrails established by the AI platforms. These are used to evaluate content from third-party resources and can be monitored by the user.	A.5.3 Supplier relationships A.6.1 AI system impact assessment
14	SolidCore monitors the pre-trained AI models capturing the version of the model and logging other information visible throught the System of Record interface.	A.6.2 AI system lifecycle A.9.1 Monitoring & measurement
15	SolidCore tracks stages of Projects and “Production” includes logging system details and capturing them in the System of Record for monitoring AI systems post-deployment.	A.9.1 Monitoring & measurement A.10.1 Incident response
16	SolidCore interfaces with the Knowledge Base resource that can specify tasks and methods that are implemented to support the AI system.	A.6.1 AI system impact assessment A.7.4 Documentation

CAPABILITY #	CAPABILITY STATEMENT	ISO 42001 Annex A
17	SolidCore interfaces with the Knowledge Base and this is captured in the System of Record, which can be used to understand the system's knowledge limits.	A.6.1 AI system impact assessment
18	SolidCore captures the type of model used in the inference call. This runtime inference of the model can assist in TEVV considerations such as the data points used to enrich the analysis. Insights captured in the Data Vector and RAG inputs can also assist the user monitoring the system for TEVV considerations.	A.6.1 AI system impact assessment
19	The System of Record captures the data sources and this can be used to detect when inputs/outputs are violating a third-party risk controls.	A.5.3 Supplier relationships A.7.4 Documentation
20	The total number of issues are tallied around each control. The user can select the Issues to drill-down and see the Policy items that are associated with the control.	A.6.1 AI system impact assessment A.9.1 Monitoring & measurement
21	A list of applied Policies to each Project can be used to document the TEVV.	A.9.1 Monitoring & measurement A.7.4 Documentation
22	SolidCore has a variety of data sources that it can log based on the AI system and inference call. These can be examined by the user to monitor privacy risks of the AI system.	A.6.1 AI system impact assessment A.8.3 Data governance
23	Supporting the System of Record is raw logging/data that can be used to refine what is needed to support monitoring AI system behavior.	A.9.1 Monitoring & measurement
24	SolidCore reports on a Project basis can be used as input into the review and measurement of the system. This information could be used in compiling system performance and be reviewed by relevant AI actors to validate the system is performing consistently as intended.	A.9.1 Monitoring & measurement A.9.3 Management review
25	SolidCore captures some performance values such as the inference latency and the token volume consumed with the inference.	A.9.1 Monitoring & measurement

ABOUT THE AUTHORS

schellman

Schellman is a leading provider of attestation, compliance, and certification services, and the only company in the world that is simultaneously a CPA firm (SOC 2), an ISO Certification Body, a globally licensed PCI Qualified Security Assessor, a HITRUST CSF Assessor, CMMC 3rd Party Assessment Organization (C3PAO) and the leading FedRAMP 3PAO.

Schellman was the first ANAB-accredited certification body for ISO/IEC 42001, the world's first certifiable AI management system standard, and is the first authorized auditor for AIUC-1, the technical assurance standard for AI agents. Schellman's AI practice combines this certification authority with supplemental attestations (SOC 2 + AI) as well as hands-on assessment experience across NIST AI RMF assessments, ISO 42001 audits, and AI red teaming, giving the firm a practitioner's view of how AI governance frameworks perform against real production

SCHELLMAN.COM
4010 W Boy Scout Blvd / Ste 600 / Tampa, FL 33607
1.866.254.0000
Outside of the United States, please dial:
1.813.288.8833

Follow @Schellman on Social Media:

solidcore.ai

Every company should be able to innovate with AI as confidently as they build in the cloud. Generative AI adoption can be safe, accountable, and scalable — when organizations have visibility and control, they need to manage risk, meet compliance standards, and build trust. Our founders are security and startup veterans whose products are used to protect the world's largest enterprises. Our team of experts in security and AI are dedicated to tackling hard problems.

SOLIDCORE.AI
887 Oak Grove Ave / Ste 203 / Menlo Park, CA 94025
hello@solidcore.ai

Follow @SolidCore.ai on Social Media:

carahsoft.

Carahsoft is The Trusted IT Solutions Provider®, supporting Federal, State, Local Government, Education, Healthcare, and Private Sector industries. As the Master Aggregator™ for our vendor, reseller and integrator partners, we deliver solutions for Cybersecurity, FinOps, MultiCloud, DevSecOps, Artificial Intelligence, Human Capital, Legal & Courtroom Technology, Customer Experience & Engagement, and more. Carahsoft is consistently recognized by its partners as a top revenue producer and is listed annually among the industry's fastest growing firms by CRN, Inc., Forbes, Inc. 5000, and Washington Business Journal. Founded in 2004, Carahsoft is headquartered in Reston, VA.

CARAHSOFT.COM
11493 Sunset Hills Rd / Ste 100 / Reston, Virginia 20190
1.888.662.2724
sales@carahsoft.com

Follow @Carahsoft on Social Media: