OVERT

OBSERVABLE VERIFICATION EVIDENCE FOR RUNTIME TRUST

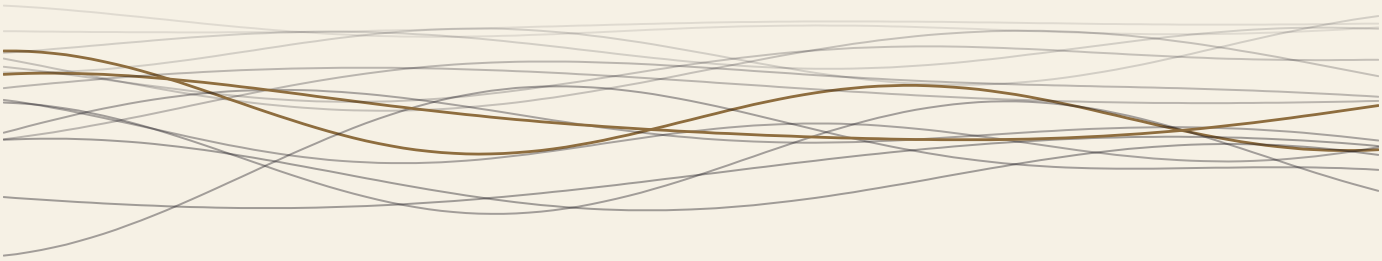

OVERT in the Context of Your Framework

NIST AI RMF · ISO/IEC 42001 · EU AI Act · AIUC-1/OWASP · SP 800-53/FedRAMP · OMB · DASF · IMDRF N93 · CHAI · RUAIH · DAGF

DATE	June 2026
PUBLISHED BY	GLACIS Technologies, Inc.
REPRODUCES	Crosswalk sections 23–31 (informative companion)
COMPLETE EDITION	overt.is
COMPANION NOTICE CONTACT	overt-review@glacis.io

This volume is the informative companion to OVERT Version 1.1. It contains the external-framework and regulatory crosswalks. Everything in it is informative: nothing here imposes requirements, modifies the standard's normative core, or determines compliance with any external framework. The Complete Edition is the sole authoritative text for conformance purposes.

Contents of this Volume

OVERT V1.1 CROSSWALKS — INFORMATIVE COMPANION

OVERT v1.1 Crosswalks — Informative Companion

This document is the informative companion to OVERT v1.1 — Observable Verification Evidence for Runtime Trust (`OVERT_v1.1_STANDARD.md`). It contains **all** external-framework and regulatory crosswalks, moved out of the standard's main text during the v1.1 finalization: NIST AI RMF, ISO/IEC 42001, the EU AI Act, AIUC-1/OWASP, NIST SP 800-53 Rev 5/FedRAMP, OMB M-25-21/M-25-22, the Databricks AI Security Framework (DASF) v3.0, IMDRF N93, the CHAI Governance Playbooks, the Joint Commission RUAIH guidance, and the Databricks AI Governance Framework (DAGF). The content below is preserved verbatim from the pre-extraction draft, including its original section numbering ("23" through "31"). One exception: Section 29.4 (Attestation Boundary Declaration) is normative and remains in the standard, renumbered as Section 22.10 — see the editor's note in place.

Everything in this document is informative. Nothing here imposes requirements, modifies the standard's normative core, or determines compliance with any external framework. Section and annex references in the text (e.g., "Section 17.5", "Annex C.10") refer to the standard. External-framework specifics below (structures, labels, figures) are reproduced as drafted and remain subject to the pre-publication fact-check register in the v1.1 review packet (`v1.1-review-packet/FLAGGED.md`); treat them as provisional until verified against primary sources.

23. Crosswalk: NIST AI RMF

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with the NIST AI Risk Management Framework.

This crosswalk maps OVERT controls to the NIST AI Risk Management Framework (AI RMF 1.0, January 2023) and the NIST AI RMF Generative AI Profile (July 2024).

NIST AI RMF Function / Category	OVERT Domain	OVERT Controls	Coverage	Notes
GOVERN 1.1–1.7 (Policies and procedures)	GOVERN, DRIFT	GOV-1, GOV-2, DRIFT-1	Partial	Machine-readable policy attestation and baseline intent declaration; OVERT attests policy existence and cryptographic binding but does not establish, implement, or evaluate policy adequacy
GOVERN 2.1–2.3 (Accountability structures)	GOVERN, DRIFT	GOV-2, DRIFT-5	Partial	Role and responsibility attestation; human oversight quality assessment. OVERT attests accountability structures but does not establish them
GOVERN 3.1–3.2 (Workforce diversity and training)	GOVERN	GOV-2	Adjacent	OVERT attests organizational role assignments but does not address workforce diversity, composition, or training programs
GOVERN 4.1–4.3 (Organizational commitments)	GOVERN	GOV-5	Partial	Disclosure and transparency attestation supports evidence of organizational commitments
GOVERN 5.1–5.2 (Engagement)	GOVERN	GOV-4	Adjacent	OVERT attests supply chain governance records; stakeholder engagement processes, community consultation, and feedback mechanisms are outside OVERT scope
GOVERN 6.1–6.2 (Policies and procedures review)	GOVERN	GOV-1	Partial	Policy review cycle attestation provides evidence that reviews occurred; review quality and adequacy are outside scope
MAP 1.1–1.6 (Context and use identification)	IDENTIFY	IDE-1	Partial	System inventory and classification attestation; context analysis and intended-use documentation are organizational responsibilities
MAP 2.1–2.3 (Context assessment)	IDENTIFY	IDE-1.2, IDE-2	Partial	Impact categorization and risk assessment attestation; contextual factors and deployment environment analysis are organizational responsibilities
MAP 3.1–3.5 (Benefits and costs)	IDENTIFY	IDE-1, IDE-2	Adjacent	Risk documentation supports evidence; benefit-cost analysis and societal impact assessment are outside OVERT scope
MAP 4.1–4.2 (Risk identification)	IDENTIFY	IDE-2, GOV-4	Partial	Risk registry with attestation; risk identification methodology and completeness are organizational responsibilities

NIST AI RMF Function / Category	OVERT Domain	OVERT Controls	Coverage	Notes
MAP 5.1–5.2 (Stakeholder impacts)	IDENTIFY	IDE-2	Adjacent	Impact assessment attestation provides supporting evidence; stakeholder identification and impact analysis are outside OVERT scope
MEASURE 1.1–1.3 (Metrics identification)	MEASURE	MEA-1, MEA-2	Direct	S3P sampling infrastructure produces attested metrics
MEASURE 2.1–2.13 (AI system evaluation)	MEASURE, DRIFT	MEA-2, MEA-3, MEA-4, DRIFT-2	Partial	Statistical safety signals with confidence intervals; behavioral drift detection per agent class. OVERT provides runtime measurement; comprehensive AI system evaluation including fairness, bias, and societal impact assessment requires additional evaluation methods
MEASURE 3.1–3.3 (Tracking and communication)	MEASURE, ATTEST	MEA-2, ATT-4	Direct	Transparency log integration provides attested tracking and communication records
MEASURE 4.1–4.3 (Measurement feedback)	MEASURE	MEA-3	Partial	TEVV process attestation supports evidence of feedback loops; measurement methodology adequacy is outside scope
MANAGE 1.1–1.4 (Risk response planning)	RESPOND	RES-1, RES-2	Partial	Attested response actions with cryptographic receipts; risk response planning, strategy development, and resource allocation are organizational responsibilities
MANAGE 2.1–2.4 (Risk treatment)	RESPOND, PROTECT, DRIFT	RES-1, PRO-1, DRIFT-3	Partial	Enforcement attestation; graph topology governance. OVERT attests that treatments executed but does not evaluate treatment effectiveness
MANAGE 3.1–3.2 (Risk monitoring)	MEASURE, DRIFT	MEA-1, MEA-2, DRIFT-2	Partial	Continuous monitoring via S3P and behavioral drift governance within attested scope; risk monitoring scope completeness and adequacy are organizational responsibilities
MANAGE 4.1–4.3 (Risk escalation)	RESPOND, DRIFT	RES-2, RES-3, DRIFT-4	Partial	Escalation and override attestation; causal drift attribution within attested scope. Escalation path design and organizational decision-making are outside scope
GOVERN 1.4 (Policies updated)	HITL	HITL-4	Direct	Policy approval attestation

NIST AI RMF Function / Category	OVERT Domain	OVERT Controls	Coverage	Notes
MANAGE 1.3 (Risk responses)	HITL	HITL-2, HITL-3	Direct	Human review and correction attestation
GOVERN 1.4 (Policies updated)	GOVERN	GOV-1	Partial	GOV-1 policy-existence/update attestation; control ownership of GOVERN 1.4 is shared with HITL-4 (policy approval)
GOVERN 6.1–6.2 (Policies and procedures review)	GOVERN, Agentic	GOV-4, MCP-1, MCP-2, MCP-3	Partial	Supply-chain/policy-review attestation also supports GOV-4 supplier governance and MCP server hosting/connection governance (MCP-1/2/3)
MAP 1.1, MAP 5.1 (Context and stakeholder identification)	GOVERN	GOV-3	Partial	GOV-3 risk-taxonomy attestation references context/stakeholder identification; primary mappings remain IDE-1 (MAP 1.1–1.6) and IDE-2 (MAP 5.1–5.2)
MAP 4.1 (Risk identification)	GOVERN	GOV-4	Partial	GOV-4 supply-chain risk-registry attestation; MAP 4.1–4.2 also maps to IDE-2
MEASURE 1.3 (Metrics identification)	MEASURE	MEA-3	Partial	MEA-3 TEVV metric attestation; MEASURE 1.1–1.3 also maps to MEA-1, MEA-2
MEASURE 2.1, 2.3, 2.5 (AI system evaluation)	MEASURE, DRIFT	MEA-4, DRIFT-4	Partial	MEA-4 evaluation and DRIFT-4 causal drift attribution; the MEASURE 2.1–2.13 range row also maps to MEA-2, MEA-3, DRIFT-2
GOVERN 1.1, GOVERN 1.3, MAP 2.1 (Policies, accountability, context)	Agentic	STATE-1, STATE-2	Partial	STATE-1/STATE-2 attest agent-state and prompt-registration governance referenced by these policy/context clauses; primary mappings remain GOV-1/GOV-2/DRIFT-1 and IDE-1.2/IDE-2
GOVERN 1.4, MANAGE 2.2, MANAGE 2.3 (Policy update and risk treatment)	Agentic	IDENT-1, STATE-1	Partial	IDENT-1 identity-delegation and STATE-1 agent-state attestation referenced by these clauses; primary mappings remain HITL-4 and the MANAGE 2.1–2.4 row (RES-1, PRO-1, DRIFT-3)
GOVERN 1.1, MANAGE 1.3 (Session lifecycle evidence)	Session Controls	SESS-1 through SESS-5	Partial	Session lifecycle attestation supports policy and risk-response clauses; primary mappings remain GOV-1/GOV-2/DRIFT-1 and HITL-2/HITL-3

24. Crosswalk: ISO/IEC 42001:2023

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with ISO/IEC 42001:2023.

This crosswalk maps OVERT controls to ISO/IEC 42001:2023 (Artificial Intelligence — Management System) clauses and Annex A controls.

ISO 42001 Clause / Annex	OVERT Domain	OVERT Controls	Coverage	Notes
4.1–4.4 Context of the organization	IDENTIFY	IDE-1	Adjacent	OVERT system inventory supports evidence of organizational context documentation; understanding of the organization, interested parties, and AIMS scope are management responsibilities
5.1–5.3 Leadership	GOVERN	GOV-1, GOV-2	Partial	Policy commitment and organizational role attestation provide evidence of leadership engagement; leadership commitment, resource allocation, and management review are outside OVERT scope
6.1.1–6.1.3 Risk assessment	IDENTIFY, GOVERN, DRIFT	IDE-2, GOV-3, DRIFT-1	Partial	AI risk classification with attestation; baseline intent declaration. OVERT supports evidence of risk assessment execution but does not prescribe risk assessment methodology or criteria
6.1.4 AI system impact assessment	IDENTIFY	IDE-2	Partial	Severity classification attestation supports evidence that impact assessments occurred; assessment completeness and methodology are organizational responsibilities
6.2 AI objectives	GOVERN	GOV-1.1	Adjacent	Objective-level attestation provides supporting records; establishing AI objectives and planning to achieve them is a management responsibility
7.5 Documented information	ATTEST	ATT-1, ATT-4	Partial	Tamper-evident records in transparency log address integrity and protection of documented information; OVERT does not address the full documented-information lifecycle including creation, updating, document

ISO 42001 Clause / Annex	OVERT Domain	OVERT Controls	Coverage	Notes
				control scope, and retention of all management system records
8.1 Operational planning and control	PROTECT	PRO-1 through PRO-5	Partial	Boundary enforcement attestation produces verifiable evidence of operational control execution within attested scope; operational planning, resource allocation, and control selection are organizational responsibilities
8.2–8.4 AI risk assessment and treatment	IDENTIFY, RESPOND	IDE-2, RES-1	Partial	Risk treatment attestation supports evidence of treatment execution; risk assessment completeness and treatment selection are organizational responsibilities
9.1 Monitoring, measurement, analysis	MEASURE, DRIFT	MEA-1, MEA-2, DRIFT-2	Direct	S3P statistical measurement; behavioral drift detection produces the monitoring artifacts
9.2 Internal audit	ATTEST	ATT-4	Adjacent	Transparency log provides audit-ready infrastructure and tamper-evident records; OVERT does not conduct internal audits, establish audit programs, define audit criteria, or ensure auditor independence — these are management system obligations
10.2 Corrective action	RESPOND, DRIFT	RES-2, DRIFT-4	Partial	Attested corrective actions; causal drift attribution. OVERT attests that corrective actions were taken but does not evaluate their adequacy or confirm nonconformity elimination
A.6.2.4 Verification and validation	MEASURE	MEA-3, MEA-4	Direct	TEVV attestation
A.6.2.6 Operation and monitoring	MEASURE, PROTECT, DRIFT	MEA-2, PRO-1, DRIFT-1, DRIFT-2	Direct	Runtime monitoring attestation; baseline intent and drift governance
A.6.2.8 Event log recording	ATTEST, TOOL	ATT-1, TOOL-5	Direct	Per-action attestation receipts
A.3.2 Roles and responsibilities	HITL, DRIFT	HITL-4.3, DRIFT-5	Partial	Separation of duties attestation; human oversight quality assessment. OVERT attests role assignments but does not define organizational roles or ensure competence

ISO 42001 Clause / Annex	OVERT Domain	OVERT Controls	Coverage	Notes
A.7.5 Data provenance	Agentic Controls	CAP-1	Direct	Capability-based data access attestation with provenance tracking
A.2.2, A.2.4 AI policy	GOVERN	GOV-1	Partial	OVERT attests existence and cryptographic binding of the AI policy and its alignment with organizational policies; policy adequacy is outside scope
A.3.2 Roles and responsibilities	GOVERN	GOV-2	Partial	Accountability-structure attestation; A.3.2 also maps to HITL-4.3, DRIFT-5 (see existing row)
A.5.2 AI system impact assessment process	IDENTIFY, GOVERN	IDE-2, GOV-3	Partial	Impact-assessment and risk-taxonomy attestation; assessment methodology is an organizational responsibility
A.5.3–A.5.5 AI system impact assessment	IDENTIFY	IDE-2	Partial	Severity/impact classification attestation supports evidence that documented impact assessments occurred; completeness is organizational
A.6.2.2 AI system requirements and specification	IDENTIFY	IDE-1	Partial	System inventory and classification attestation supports requirements/specification records; requirements definition is organizational
A.6.2.3 AI system design / verification configuration	Agentic	MCP-1, MCP-2, MCP-3, IDENT-1, STATE-1	Partial	OVERT attests runtime identity, MCP hosting/connection configuration, identity delegation, and agent-state design at mediation boundaries; full design documentation is organizational
A.8.2 Information for interested parties	GOVERN	GOV-5	Partial	Disclosure/transparency attestation supports evidence of information provided to interested parties; content adequacy is outside scope
A.8.4 Information for users / incident communication	RESPOND	RES-2	Partial	Attested corrective-action records support incident/user communication evidence; communication obligation is organizational
A.10.2, A.10.3 Responsible use / third-party and supplier	GOVERN	GOV-4	Partial	Supply-chain and third-party governance attestation; supplier-relationship management is organizational

25. Crosswalk: EU AI Act

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with Regulation (EU) 2024/1689.

This crosswalk maps OVERT controls to Regulation (EU) 2024/1689 (Artificial Intelligence Act). The Regulation generally applies from **2 August 2026**. Article 6(1) obligations and corresponding provisions apply from **2 August 2027**. Annex III systems are classified under Article 6(2) and follow the general application date. References in this crosswalk reflect the Regulation as adopted.

> ***Note on Digital Omnibus:*** *The Commission published the Digital Omnibus on AI on November 19, 2025. A provisional political agreement was reached on 6 to 7 May 2026 and confirmed by Member State representatives in the Council on 13 May 2026. While not yet formally adopted, it is expected to be published in the Official Journal before 2 August 2026. The agreement defers standalone Annex III high-risk obligations from 2 August 2026 to 2 December 2027; product-embedded Annex I high-risk obligations (which include medical devices) from 2 August 2027 to 2 August 2028; Article 50(2) AI-content watermarking to 2 December 2026; and the national regulatory sandbox obligation. It adds a new Article 5 prohibition on AI-generated non-consensual intimate imagery and CSAM. Until adoption and Official Journal publication, the original dates legally stand.*

OVERT attestation architecture is designed to support the requirements of the EU AI Act but does not determine conformity. Conformity assessment for high-risk AI systems under Article 43 is conducted by notified bodies or through internal control procedures as applicable.

EU AI Act Article	Subject	OVERT Domain	OVERT Controls	Coverage	Notes
Article 9	Risk management system	GOVERN, IDENTIFY, DRIFT	GOV-3, IDE-2, DRIFT-1, DRIFT-2	Partial	OVERT attests risk classification execution and provides behavioral drift detection, supporting evidence that risk monitoring occurred. Article 9 requires a complete risk management system lifecycle — including risk identification, estimation, evaluation, elimination/mitigation of risks, and testing — that extends well beyond runtime attestation

EU AI Act Article	Subject	OVERT Domain	OVERT Controls	Coverage	Notes
Article 10	Data and data governance	PROTECT	PRO-5, CAP-1	Partial	Data boundary enforcement attestation; capability-scoped data access. Article 10 data governance requirements (training data quality, representativeness, bias examination) are outside OVERT scope
Article 11	Technical documentation	ATTEST, GOVERN	ATT-1, ATT-4, GOV-1	Partial	Tamper-evident technical documentation in transparency log; machine-readable policy records. OVERT supports integrity and availability of documentation but does not generate the full technical documentation required by Annex IV
Article 12	Record-keeping	ATTEST, TOOL	ATT-1 through ATT-5, TOOL-5	Partial	Automatic logging with cryptographic attestation; tamper-evident records; retroactive reconstruction via transparency log. OVERT provides attested automatic logging records aligned with Article 12 record-keeping objectives, but Article 12 may require additional system-specific data elements outside OVERT scope
Article 13	Transparency and provision of information to deployers	GOVERN	GOV-5, DISC-1	Partial	AI system disclosure attestation; deployer-facing transparency records. OVERT supports evidence of transparency measures but does not generate all deployer-facing information required by Article 13
Article 14	Human oversight	HITL, DRIFT	HITL-1 through HITL-4, DRIFT-5	Partial	Human-in-the-loop attestation: consent, review, correction, and override with cryptographic receipts; human oversight quality assessment. OVERT attests that oversight occurred but does not ensure oversight effectiveness or design appropriate oversight measures
Article 15	Accuracy, robustness, and cybersecurity	PROTECT, MEASURE, DRIFT	PRO-1 through PRO-4, MEA-2, MEA-3, DRIFT-2	Partial	Boundary enforcement attestation; statistical safety measurement with confidence intervals; behavioral drift detection. OVERT provides runtime evidence but does not address the full accuracy, robustness, and cybersecurity lifecycle including design-time measures

EU AI Act Article	Subject	OVERT Domain	OVERT Controls	Coverage	Notes
Article 17	Quality management system	GOVERN, MEASURE	GOV-1, GOV-2, MEA-4	Adjacent	Governance policy attestation; quality management process records. Article 17 requires a comprehensive QMS covering design, development, testing, and post-market obligations; OVERT provides supporting attestation evidence for runtime aspects only
Article 26	Obligations of deployers	ATTEST, HITL	ATT-4, HITL-1, HITL-4	Partial	Deployer logging obligations; human oversight documentation. OVERT supports evidence of deployer obligations but does not address all deployer responsibilities under Article 26
Article 72	Reporting of serious incidents	RESPOND	RES-1, RES-2, RES-5	Partial	Attested incident detection and response; declared failure modes. OVERT provides incident detection evidence; the reporting obligation itself and causal investigation are organizational responsibilities
Article 13	Transparency and provision of information to deployers	Agentic	STATE-2	Partial	STATE-2 prompt-registration and prompt-to-action traceability supports deployer-facing transparency evidence; Article 13 also maps to GOV-5, DISC-1
Article 9	Risk management system	Agentic	IDENT-1	Partial	IDENT-1 identity-delegation attestation supports risk-management evidence at the agent boundary; Article 9 also maps to GOV-3, IDE-2, DRIFT-1, DRIFT-2

26. Crosswalk: AIUC-1 / OWASP

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with AIUC-1 or OWASP guidance.

This crosswalk maps OVERT controls to the AIUC-1 framework (January 2026 version) and the OWASP Top 10 for Agentic Applications, showing how OVERT attestation infrastructure supports evidence for these requirements by upgrading assurance levels over documentation-based approaches.

AIUC-1 / OWASP Reference	OVERT Domain	OVERT Controls	Coverage	Attestation Assurance Level Upgrade
A001–A007 (Data and Privacy)	PROTECT	PRO-5	Partial	Documentation to AAL-4 attestation receipts for data boundary enforcement; data privacy program design and DPIAs are outside scope
B001 (Adversarial Testing)	MEASURE	MEA-3.1	Partial	OVERT attests that testing occurred and publishes results to transparency log; OVERT does not conduct adversarial testing
B004 (Endpoint Protection)	PROTECT	PRO-3	Direct	Documentation to AAL-4 rate limit attestation receipts
B006 (Agent Actions)	TOOL, PROTECT	TOOL-1, PRO-1	Direct	Documentation to AAL-4 per-action attestation receipts
B008 (Deployment Security)	ATTEST	ATT-2	Direct	Documentation to AAL-4 co-epoch binding
C001 (Risk Taxonomy)	GOVERN	GOV-3	Partial	Document to AAL-4 machine-readable taxonomy; OVERT attests taxonomy existence and binding but does not evaluate taxonomy correctness or completeness
C003–C005 (Safety Controls)	PROTECT	PRO-4	Direct	Documentation to AAL-4 filter attestation receipts
C010–C012 (Third-party Testing)	MEASURE	MEA-3	Direct	Reports to AAL-2 + AAL-4 transparency log summary
D003 (Tool Call Restriction)	TOOL	TOOL-1 through TOOL-5	Direct	Documentation to AAL-4 per-call attestation
E015 (Logging)	ATTEST, TOOL	ATT-1 through ATT-4, TOOL-5	Direct	Logs to AAL-4 tamper-evident attested records
E016 (AI Disclosure)	GOVERN	GOV-5	Partial	Demonstrations to AAL-2 disclosure attestation; disclosure content adequacy is outside scope
E004 (Accountability)	HITL	HITL-4	Direct	Document to AAL-4 policy approval receipts

AIUC-1 / OWASP Reference	OVERT Domain	OVERT Controls	Coverage	Attestation Assurance Level Upgrade
D003.4 (Approval Gates)	HITL	HITL-1, HITL-2, HITL-3	Direct	Documentation to AAL-4 consent/review/correction receipts
OWASP Agentic #1 (Prompt Injection)	PROTECT, TOOL	PRO-1, TOOL-1	Partial	Boundary enforcement attestation at tool-call boundary; prompt injection prevention effectiveness depends on filter quality, which OVERT attests but does not determine
OWASP Agentic #2 (Tool Misuse)	TOOL, DRIFT	TOOL-1 through TOOL-5, DRIFT-3	Direct	Per-call policy evaluation with attestation receipt; graph topology governance
OWASP Agentic #3 (Privilege Escalation)	Agentic Controls, DRIFT	CAP-1, CAP-2, DRIFT-3	Partial	Capability-based access control with attestation; spawn authorization governance. CAP-2 architectural separation evidence is AAL-3 at Level 3 and AAL-4 at Level 4; privilege escalation prevention effectiveness depends on policy quality and scope completeness
E004 (Accountability)	GOVERN	GOV-2	Partial	Accountability-structure attestation; E004 also maps to HITL-4 (see existing row)
E006 (Organizational commitments)	GOVERN	GOV-4	Partial	Supply-chain/organizational governance attestation supports evidence of organizational commitments
E016 (AI Disclosure)	GOVERN	GOV-5	Partial	Disclosure attestation; this row also confirms GOV-5 ownership for the inline GOV-5 tag
E017 (Incident disclosure)	Disclosure	DISC-1	Partial	DISC-1 disclosure attestation supports incident/AI-disclosure evidence; disclosure content adequacy is outside scope
E001–E003 (Incident response and reporting)	RESPOND	RES-2	Partial	Attested incident response and corrective-action records; reporting obligation and timelines are organizational
B005 (Output safety / content controls)	PROTECT	PRO-4	Partial	Output-filter attestation; B005 effectiveness depends on filter quality
C006 (Safety controls)	PROTECT	PRO-4	Partial	Filter/safety-control attestation; C003–C006 range supports PRO-4

AIUC-1 / OWASP Reference	OVERT Domain	OVERT Controls	Coverage	Attestation Assurance Level Upgrade
A005, A006 (Data and Privacy)	PROTECT	PRO-5	Partial	Data-boundary enforcement attestation; A005/A006 fall within the A001–A007 → PRO-5 mapping (made explicit here)
C002 (Risk evaluation / measurement)	MEASURE	MEA-4	Partial	Evaluation attestation supports risk-measurement evidence; C002 also relates to MEASURE 2.3/2.5 runtime signals
D002, D004 (Tool / agent action controls)	MEASURE, TOOL	MEA-3, TOOL-1 through TOOL-5	Partial	TEVV testing attestation and per-call tool authorization; D002/D004 effectiveness depends on policy quality
D003.1, D003.2, D003.3 (Tool Call Restriction sub-controls)	TOOL	TOOL-1 through TOOL-5	Direct	Per-call attestation; D003.1/.2/.3 are sub-controls of D003, which maps to TOOL-1 through TOOL-5 (made explicit here)
OWASP Agentic #4 (Resource Overload / Unbounded Consumption)	PROTECT, TOOL	PRO-3, TOOL-3	Partial	Rate limiting and circuit breaking with attestation; referenced by inline MCP-3 tag

27. Crosswalk: NIST SP 800-53 Rev 5 / FedRAMP

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with NIST SP 800-53 Rev 5 or FedRAMP requirements.

This crosswalk maps OVERT controls to NIST SP 800-53 Revision 5 security and privacy control families. Organizations pursuing FedRAMP authorization at the Moderate or High baseline can use this mapping to identify how OVERT attestation architecture supports evidence for required control implementations.

OVERT does not replace 800-53 controls. It provides a cryptographic attestation layer that supports evidence for audit, integrity, and accountability controls by producing independently verifiable records of control execution.

800-53 Family	Representative Controls	OVERT Controls	Coverage	Alignment Notes
AU — Audit and Accountability	AU-2 (Event Logging), AU-3 (Content of Audit Records), AU-6 (Audit Record Review), AU-9 (Protection of Audit Information), AU-10 (Non-repudiation), AU-11 (Audit Record Retention), AU-12 (Audit Record Generation)	ATT-1 through ATT-5, TOOL-5, Section 21	Partial	Strong alignment for AI governance audit records within attested scope. OVERT attestation receipts produce tamper-evident, timestamped, independently verifiable audit records. Transparency log provides AU-9 (protection via append-only cryptographic structure). Section 21 addresses AU-11 (retention). ATT-2 (co-epoch binding) supports AU-10 (non-repudiation). OVERT covers AU controls for AI governance events; system-wide AU controls (non-AI event logging, centralized audit reduction, cross-domain audit) require separate implementation.
SI — System and Information Integrity	SI-4 (System Monitoring), SI-5 (Security Alerts), SI-6 (Security and Privacy Function Verification), SI-7 (Software, Firmware, and Information Integrity)	PRO-1 through PRO-5, MEA-1 through MEA-4, DRIFT-1 through DRIFT-3	Partial	OVERT boundary enforcement (PRO controls) provides attested system integrity monitoring for AI governance controls. S3P (MEA-2) produces statistical safety measurements with confidence intervals. Co-epoch binding (ATT-2) supports SI-7 by attesting binary and configuration integrity. DRIFT controls provide behavioral drift detection (SI-4), graph topology governance alerts (SI-5), and baseline intent verification (SI-6). OVERT addresses AI-layer integrity; host-level and network-level SI controls require separate implementation.
IA — Identification and Authentication	IA-2 (Identification and Authentication), IA-3 (Device Identification and Authentication), IA-8	HITL-1 through HITL-4, ATT-2	Partial	OVERT HITL controls provide attested identity binding for human actors in AI governance workflows. ATT-2 provides cryptographic identity binding for system components via notary-derived binary identity. OVERT does not implement authentication mechanisms, identity proofing, or

800-53 Family	Representative Controls	OVERT Controls	Coverage	Alignment Notes
	(Identification and Authentication — Non-Organizational Users), IA-12 (Identity Proofing)			credential management — these require separate IA control implementations.
CM — Configuration Management	CM-2 (Baseline Configuration), CM-3 (Configuration Change Control), CM-6 (Configuration Settings), CM-8 (System Component Inventory)	ATT-2 (co-epoch binding), Section 18	Partial	Co-epoch binding attests configuration state at each epoch. Configuration drift is cryptographically detectable (Section 18.6). Binary identity and network isolation state are independently verified by notaries. OVERT provides configuration integrity evidence for AI governance components; CM controls for the broader system environment require separate implementation.
IR — Incident Response	IR-4 (Incident Handling), IR-5 (Incident Monitoring), IR-6 (Incident Reporting), IR-8 (Incident Response Plan)	RES-1 through RES-5	Partial	OVERT RESPOND controls provide cryptographically gated incident response with attested escalation, override, and revocation procedures. RES-5 (declared failure modes) supports IR-8 planning requirements. OVERT addresses AI governance incidents; organization-wide IR program, staff training, and IR plan development are outside scope.
AC — Access Control	AC-3 (Access Enforcement), AC-4 (Information Flow Enforcement), AC-6 (Least Privilege), AC-25 (Reference Monitor)	CAP-1, CAP-2, TOOL-2	Partial	OVERT capability-based access control (CAP controls) implements least privilege with attestation at AI tool-call boundaries. TOOL-2 (schema enforcement) acts as a reference monitor for tool-call boundaries. Information flow enforcement is attested via non-egress architecture. OVERT provides access control attestation for AI system actions; system-wide AC controls (user account management, session controls, remote access) require separate implementation.
SC — System and Communi-	SC-7 (Boundary Protection), SC-8	PRO-2, ATT-1, Section 17	Partial	Non-egress architecture (Section 17) provides boundary protection with attestation for AI governance data flows. ATT-1 defines cryptographic

800-53 Family	Representative Controls	OVERT Controls	Coverage	Alignment Notes
cations Protection	(Transmission Confidentiality and Integrity), SC-12 (Cryptographic Key Establishment and Management), SC-13 (Cryptographic Protection)			commitment constructions. Split-knowledge key hierarchy supports SC-12. OVERT addresses SC controls for attestation infrastructure; network-level boundary protection and system-wide encryption require separate implementation.
SA — System and Services Acquisition	SA-4 (Acquisition Process), SA-9 (External System Services), SA-11 (Developer Testing and Evaluation)	GOV-1, MEA-3, ATT-5	Adjacent	OVERT governance policy attestation supports evidence for SA-4 acquisition documentation. MEA-3 (TEVV) provides attested testing and evaluation. ATT-5 (notary governance) documents external service dependencies. OVERT provides supporting evidence; SA controls address broader acquisition lifecycle and vendor management processes outside OVERT scope.
PM — Program Management	PM-9 (Risk Management Strategy), PM-28 (Risk Framing)	GOV-3, IDE-2	Adjacent	OVERT risk classification (GOV-3) and impact assessment (IDE-2) produce attested risk management records that support evidence for PM-level risk strategy documentation. Risk management strategy development and organizational risk framing are management responsibilities outside OVERT scope.

FedRAMP Considerations

For cloud service providers (CSPs) pursuing FedRAMP authorization:

- OVERT attestation architecture can serve as supporting evidence for AU and SI control implementations in the System Security Plan (SSP). OVERT does not satisfy these controls independently.
- Transparency log data supports evidence for continuous monitoring requirements under the FedRAMP Continuous Monitoring Strategy Guide.
- The non-egress architecture (Section 17) is designed to minimize transfer of customer content outside the operator boundary. Classification of attestation artifacts within a specific authorization boundary is determined by the system security plan and the authorizing official.

- S3P statistical safety signals provide quantitative metrics suitable for Plan of Action and Milestones (POA&M) tracking.

28. Crosswalk: OMB M-25-21 / M-25-22

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with OMB M-25-21, M-25-22, or any other federal requirement.

This crosswalk maps OVERT controls to federal AI procurement and governance requirements established by Office of Management and Budget memoranda M-25-21 ("Accelerating Federal Use of AI through Innovation, Governance, and Public Trust," April 2025) and M-25-22 ("Driving Efficient Acquisition of Artificial Intelligence in Government," effective for solicitations issued on or after **October 1, 2025** (180 days after issuance)).

Important: M-25-22 excludes National Security Systems (NSS) from its scope. Organizations operating NSS should consult applicable national security directives and their authorizing officials for AI governance requirements.

M-25-21: Agency AI Governance and Risk Management

M-25-21 requires federal agencies to inventory AI use cases, implement risk management practices, and report on AI governance (initial reporting on the schedule set by OMB implementation instructions). OVERT controls support evidence for these requirements as follows:

M-25-21 Requirement	OVERT Domain	OVERT Controls	Coverage	Alignment Notes
AI use case inventory and registration	GOVERN, IDENTIFY	GOV-1, GOV-5, IDE-1	Partial	Machine-readable governance policy (GOV-1) and system inventory (IDE-1) produce attested records that support evidence for AI use case registration. GOV-5 disclosure controls support transparency reporting. Inventory completeness and use case categorization are agency responsibilities.
Risk management practices for AI	GOVERN, IDENTIFY	GOV-3, IDE-2	Partial	Severity classification (GOV-3) and impact assessment (IDE-2) produce attested risk cat-

M-25-21 Requirement	OVERT Domain	OVERT Controls	Coverage	Alignment Notes
				egorization supporting evidence for NIST AI RMF alignment. Risk management practice design and implementation are agency responsibilities.
Rights-impacting and safety-impacting AI designation	IDENTIFY	IDE-2	Partial	Impact assessment attestation supports evidence for designation of rights-impacting and safety-impacting AI use cases with verifiable records. Designation decisions are made by agency officials.
Minimum risk management practices	GOVERN, MEASURE, RESPOND	GOV-1 through GOV-5, MEA-1 through MEA-4, RES-1 through RES-5	Partial	OVERT governance, measurement, and response domains collectively support evidence for minimum practices: impact assessment, monitoring, human oversight, and incident response — all with cryptographic attestation. Completeness and adequacy of minimum practices are agency responsibilities.
AI governance body oversight	GOVERN	GOV-2, HITL-4	Adjacent	Organizational role attestation (GOV-2) and policy approval attestation (HITL-4) document governance body decisions with tamper-evident records. Governance body establishment, composition, and oversight effectiveness are agency responsibilities.
Human oversight and appeal mechanisms	HITL	HITL-1 through HITL-4	Partial	OVERT human-in-the-loop controls provide attested consent (HITL-1), review (HITL-2), correction (HITL-3), and approval (HITL-4) workflows supporting evidence for human oversight requirements. Appeal mechanism design and due process obligations are agency responsibilities.
Ongoing monitoring and evaluation	MEASURE	MEA-1, MEA-2, MEA-4	Direct	S3P statistical safety signals provide continuous, quantitative, independently verifiable monitoring suitable for ongoing evaluation reporting.

M-25-22: AI in Federal Procurement

M-25-22 establishes requirements for the acquisition of AI capabilities by federal agencies, effective for solicitations issued on or after October 1, 2025 (180 days after issuance). OVERT controls support evidence for contractor compliance as follows:

M-25-22 Requirement	OVERT Domain	OVERT Controls	Coverage	Alignment Notes
AI performance and outcome tracking	MEASURE	MEA-1, MEA-2, MEA-4	Partial	S3P statistical safety signals provide quantitative performance metrics with confidence intervals within attested scope. OVERT tracks governance-control performance, not broad AI outcome tracking (accuracy, fairness, societal impact).
Data use restrictions and protections	PROTECT	PRO-1 through PRO-5	Partial	Boundary enforcement attestation supports evidence of data handling controls. Non-egress architecture (Section 17) supports evidence that declared boundary and non-egress controls executed within the attested scope. Data use restriction policy design is outside OVERT scope.
Transparency and explainability requirements	GOVERN	GOV-5, DISC-1	Partial	AI disclosure attestation provides verifiable transparency records. Machine-readable governance policy supports explainability documentation. Explainability of model decisions is outside OVERT scope.
Testing, evaluation, verification, and validation (TEVV)	MEASURE	MEA-3, MEA-4	Direct	OVERT TEVV controls produce attested test results at intervals defined in the operator's risk management policy.
Incident reporting and response	RESPOND	RES-1, RES-2, RES-5	Partial	Attested incident detection and response with cryptographic receipts. Declared failure modes (RES-5) support evidence for incident reporting requirements. Reporting obligations and timelines are agency responsibilities.
Continuous monitoring	MEASURE, ATTEST	MEA-2, ATT-4	Direct	S3P and transparency log provide continuous monitoring infrastructure with independently verifiable records.
Vendor risk assessment	ATTEST	ATT-5	Adjacent	Notary network governance model attestation provides documented independent verification of vendor AI governance claims. Ven-

M-25-22 Requirement	OVERT Domain	OVERT Controls	Coverage	Alignment Notes
				dor risk assessment methodology and vendor selection are agency responsibilities.

Federal AI Procurement Alignment Summary

Agencies MAY reference OVERT conformance levels in Statements of Work (SOWs) or evaluation criteria as one mechanism for supporting evidence of AI governance practices:

- **OVERT Level 1 (Foundation):** Supports evidence for AI use case inventory and risk documentation requirements under M-25-21.
- **OVERT Level 2 (Enforcement):** Supports evidence for data use restriction and boundary enforcement requirements under M-25-22.
- **OVERT Level 3 (Measurement):** Supports evidence for continuous monitoring, TEVV, and human oversight requirements under both M-25-21 and M-25-22.
- **OVERT Level 4 (Evidence-Grade):** Provides the highest OVERT evidence and preservation tier for high-risk AI procurements, with governance records that are independently verifiable within the declared scope and subject to the denominator-source and scope disclosures required by Section 22.4.

OVERT conformance does not determine compliance with M-25-21, M-25-22, or any other federal requirement. Compliance determinations are made by contracting officers, agency Chief AI Officers, and other designated officials applying the requirements of each memorandum to specific acquisitions and use cases.

29. Crosswalk: Databricks AI Security Framework (DASF) v3.0

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context). OVERT conformance does not determine compliance with DASF recommendations.

This crosswalk maps OVERT controls to the Databricks AI Security Framework (DASF) v3.0 compendium and companion agentic AI whitepaper. Per Databricks' published v3.0 release materials, the compendium inventories 97 technical security risks across 13 AI system components and 73 mitigation controls. Relative to the older pre-agentic OVERT draft baseline, DASF v3.0 adds Compo-

nent 13 (Agentic AI), 35 new agentic technical risks, and six new v3.0 controls (DASF 68-73); the compendium also includes controls DASF 65-67 introduced between the earlier v2.0 snapshot and the v3.0 agentic release. DASF is published by Databricks, Inc.; consult Databricks' published terms for licensing.

DASF describes what controls to configure across the AI system lifecycle. OVERT specifies how to produce cryptographic proof that runtime controls executed. The frameworks are complementary: DASF informs risk identification and control selection; OVERT provides the attestation layer that makes control execution independently verifiable. Where the companion PDF and earlier OVERT materials diverge on counts, this section follows the Databricks v3.0 compendium as the denominator source.

OVERT attestation is architecturally applicable to runtime enforcement, monitoring, governance, and agentic tool-use controls. DASF controls addressing training-time operations (data preparation, model training, experiment tracking), platform infrastructure (vulnerability management, SDLC, patching), and unmanaged client or secret-storage surfaces fall outside the OVERT attestation scope. These are noted as gaps below.

29.1 Risk-to-Control Crosswalk

The following table maps DASF risk categories to OVERT controls. Rows are grouped by DASF's 13 AI system components.

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
Data Operations (Components 1–4)				
1.1 Insufficient access controls	PROTECT, Agentic	CAP-1, PRO-5, TOOL-2	Direct	OVERT attests access enforcement at tool-call boundaries with per-action receipts; capability-based access control provides provenance-aware authorization
1.2 Missing data classification	IDENTIFY, GOVERN	IDE-1.2, GOV-3	Direct	System categorization and machine-readable risk taxonomy provide classification attestation artifacts
1.3 Poor data quality	MEASURE	MEA-2, MEA-4	Partial	S3P sampling detects quality degradation signals with confidence intervals; pre-deployment testing provides baseline. OVERT detects quality signals at runtime but does not address data quality management processes

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
1.4 Ineffective storage and encryption	PROTECT, ATTEST	PRO-2, ATT-1.4	Partial	OVERT attests network isolation state (NETATT) and TLS certificate pins; encryption implementation is outside OVERT scope
1.5 Lack of data versioning	—	—	—	No OVERT analog. Data versioning is an operational concern; OVERT attests policy and configuration versions
1.6 Insufficient data lineage	Agentic	CAP-1	Partial	Data provenance tracking covers lineage attestation for data flowing through tool calls; full-lifecycle data lineage is outside scope
1.7 Lack of data trustworthiness	ATTEST	ATT-1, ATT-4	Partial	Transparency log provides tamper-evident records of data access decisions with inclusion proofs; data quality and trustworthiness at source are outside scope
1.8 Legality of data	GOVERN, IDENTIFY	GOV-1, GOV-4, IDE-2	Adjacent	Governance policy and impact assessment support evidence of legal compliance documentation; supply chain governance covers third-party data. Legal compliance determination is outside OVERT scope
1.9 Stale data	—	—	—	No OVERT analog. Data freshness is an operational concern outside runtime attestation scope
1.10 Lack of data access logs	ATTEST, TOOL	ATT-1, ATT-4, TOOL-5	Direct	Tamper-evident, per-action audit trail with transparency log and notary attestation within attested scope
1.11 Compromised third-party datasets	GOVERN, PROTECT	GOV-4, PRO-1, PRO-4	Partial	Supply chain governance + boundary enforcement + input filtering with attestation receipts. OVERT attests enforcement at ingest boundaries; dataset integrity verification at source is outside scope
2.1 Preprocessing integrity	ATTEST	ATT-1, ATT-2	Partial	Co-epoch binding attests system configuration integrity; binary identity prevents unauthorized modification. Preprocessing logic correctness is outside scope
2.2 Feature manipulation	PROTECT	PRO-4, PRO-1	Partial	Input filtering + boundary enforcement attestation; runtime detection only
2.3 Raw data criteria	—	—	—	No OVERT analog. Data selection criteria are training-time decisions

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
2.4 Adversarial partitions	—	—	—	No OVERT analog. Train/test split manipulation is a training-time risk
3.1 Data poisoning	PROTECT, ATTEST	PRO-1, PRO-4, ATT-1	Partial	Boundary enforcement attests data ingestion policy compliance; transparency log records all decisions. Data poisoning prevention at the training level is outside scope
3.2 Ineffective storage and encryption (datasets)	PROTECT	PRO-2	Partial	Network isolation attestation (NETATT); storage encryption is infrastructure-level
3.3 Label flipping	—	—	—	No OVERT analog. Label manipulation is a training-time attack
4.1 Lack of traceability and transparency	GOVERN, ATTEST, Agentic	GOV-1, ATT-4, DISC-1.2	Direct	Machine-readable governance policy + transparency log + AI Bill of Materials
4.2 Lack of end-to-end ML lifecycle	GOVERN, HITL	GOV-1, GOV-2, HITL-4	Partial	Governance policy + accountability + configuration approval attestation; OVERT covers runtime governance within the ML lifecycle, not the full lifecycle
Model Operations (Components 5–8)				
5.1 Lack of experiment tracking	—	—	—	No OVERT analog. Experiment tracking is a development-time concern
5.2 Model drift	MEASURE, RESPOND, DRIFT	MEA-2, MEA-4, RES-1, DRIFT-2	Direct	S3P detects drift via violation rate changes with confidence intervals; adaptive control loop responds; DRIFT-2 detects behavioral drift within authorized bounds
5.2 Model drift (baseline intent)	DRIFT	DRIFT-1	Partial	Baseline intent declaration supports drift-detection evidence; DASF 5.2 also maps to DRIFT-2 (see existing row)
8.3 Model lifecycle without HITL (oversight quality)	DRIFT	DRIFT-5	Partial	Human-oversight-quality attestation supports lifecycle-without-HITL evidence; DASF 8.3 also maps to HITL-1 through HITL-4 (see existing row)
9.13 Excessive agency (baseline drift)	DRIFT	DRIFT-1	Partial	Baseline intent supports excessive-agency detection; DASF 9.13 also maps to TOOL-1–TOOL-5, CAP-1, CAP-2, DRIFT-3 (see existing row)

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
10.1 Lack of inference quality monitoring (drift)	DRIFT	DRIFT-2	Partial	Behavioral drift detection supports inference-quality monitoring; DASF 10.1 also maps to MEA-2, MEA-4, RES-1 (see existing row)
5.3 Hyperparameters stealing	PROTECT	PRO-2, PRO-5	Partial	Network isolation + data isolation attestation supports evidence of parameter exfiltration prevention
5.4 Malicious libraries	GOVERN, ATTEST	GOV-4, ATT-2.2	Partial	Supply chain governance + binary identity attestation detects unauthorized code at attestation boundaries; library vetting is outside scope
6.1 Evaluation data poisoning	MEASURE	MEA-3, MEA-4	Partial	Third-party testing + pre-deployment testing provide independent evaluation; evaluation data integrity is outside scope
6.2 Insufficient evaluation data	MEASURE	MEA-3, MEA-4	Partial	OVERT mandates testing scope documentation published to transparency log; evaluation data sufficiency determination is outside scope
6.3 Lack of interpretability	GOVERN, Agentic	GOV-5, DISC-1	Adjacent	AI disclosure + transparency documentation; OVERT attests disclosure existence, not model interpretability
7.1 Backdoor/Trojaned model	ATTEST, GOVERN	ATT-2.2, GOV-4	Partial	Binary identity attestation detects model artifact tampering; supply chain governance. OVERT detects modification but cannot detect backdoors embedded before initial attestation
7.2 Model assets leak	PROTECT, ATTEST	PRO-2, PRO-5, ATT-1	Direct	Non-egress architecture + data isolation + network isolation prevent model exfiltration with attestation proof
7.3 ML supply chain vulnerabilities	GOVERN	GOV-4	Partial	Supply chain and third-party governance attestation; vulnerability assessment and remediation are outside scope
7.4 Source code control attack	—	—	—	No OVERT analog. Source code management is development infrastructure security
8.1 Model attribution	ATTEST, Agentic	ATT-1, ATT-4, DISC-1	Direct	Per-interaction attestation receipts provide cryptographic attribution; AI Bill of Materials

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
8.2 Model theft	PROTECT, ATTEST	PRO-2, PRO-5, ATT-1	Direct	Non-egress architecture prevents content exfiltration; attestation proves containment
8.3 Model lifecycle without HITL	HITL	HITL-1 through HITL-4	Partial	Runtime HITL attestation: consent, review, correction, and approval; full-lifecycle human oversight is broader than runtime attestation
8.4 Model inversion	PROTECT, ATTEST	PRO-1, PRO-3, ATT-1.2	Partial	Boundary enforcement + rate limiting + keyed commitments ensure content never leaves operator boundary; model inversion prevention effectiveness depends on rate limit configuration
Model Serving — Inference (Components 9–10)				
9.1 Prompt injection	PROTECT, TOOL	PRO-1, PRO-4, TOOL-1	Partial	Boundary enforcement + input filtering + pre-execution policy enforcement with attestation receipts; prompt injection prevention effectiveness depends on filter quality
9.2 Model inversion (serving)	PROTECT	PRO-3, PRO-1	Partial	Rate limiting + boundary enforcement attestation reduce query volume for extraction attacks
9.3 Model breakout	PROTECT, ATTEST	PRO-2, ATT-2	Direct	Network isolation attestation (NETATT) + co-epoch binding proves sandbox containment
9.4 Looped input	TOOL	TOOL-3.4, TOOL-3.3	Direct	Loop detection + circuit breaker with attested termination
9.5 Infer training data membership	PROTECT	PRO-3, PRO-1	Partial	Rate limiting attestation reduces query budget for membership inference; does not prevent membership inference attacks
9.6 Discover ML model ontology	PROTECT	PRO-3, PRO-1	Partial	Rate limiting + boundary enforcement attestation reduce model fingerprinting; does not prevent ontology discovery
9.7 Denial of service	PROTECT, TOOL	PRO-3, TOOL-3	Direct	Rate limiting + circuit breaking with attestation receipts
9.8 LLM hallucinations	MEASURE, PROTECT	MEA-2, PRO-4	Partial	S3P measures hallucination rates with confidence intervals; output filtering attestation. OVERT detects and measures halluci-

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
				nation rates but does not prevent hallucinations
9.9 Input resource control	PROTECT, TOOL	PRO-3, TOOL-3.1	Direct	Rate limiting + per-tool rate limits with attestation
9.10 Accidental data exposure	PROTECT, Agentic	PRO-5, CAP-1	Direct	Data isolation + capability-based access control with provenance tracking
9.11 Model inference API access	PROTECT, TOOL	PRO-1, TOOL-2	Direct	Boundary enforcement + function authorization with attested policy evaluation
9.12 LLM jailbreak	PROTECT, MEASURE	PRO-4, PRO-1, MEA-2	Partial	Input/output filtering + boundary enforcement + S3P violation rate monitoring; jailbreak prevention effectiveness depends on filter quality
9.13 Excessive agency	TOOL, Agentic, DRIFT	TOOL-1 through TOOL-5, CAP-1, CAP-2, DRIFT-3	Direct	Per-action policy enforcement with attestation receipts and capability scoping within attested scope; graph topology governance constrains agent proliferation. Prevention effectiveness depends on policy quality and scope completeness
10.1 Lack of inference quality monitoring	MEASURE, RESPOND	MEA-2, MEA-4, RES-1	Direct	S3P continuous measurement with confidence intervals + adaptive control loop
10.2 Output manipulation	PROTECT, ATTEST	PRO-4, ATT-1	Partial	Output filtering attestation + receipt integrity proves output was not tampered with post-attestation; pre-attestation output manipulation is outside scope
10.3 Discover model ontology (output)	PROTECT	PRO-3, PRO-4	Partial	Rate limiting + output filtering attestation reduce fingerprinting surface
10.4 Discover model family	PROTECT	PRO-3	Partial	Rate limiting attestation reduces fingerprinting attack surface
10.5 Black box attacks	PROTECT, MEASURE	PRO-3, PRO-4, MEA-2	Partial	Rate limiting + filtering + statistical monitoring with S3P; OVERT detects and rate-limits but does not prevent all black box attacks
10.6 Sensitive data output	PROTECT	PRO-4, PRO-5.3	Direct	Output filtering + PII detection attestation
Operations and Platform (Components 11–12)				

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
11.1 Lack of MLOps standards	GOVERN, HITL	GOV-1, GOV-2, HITL-4	Partial	Governance policy + accountability + configuration approval attestation for operational processes; MLOps standard definition is outside scope
12.1 Lack of vulnerability management	—	—	—	No OVERT analog. Infrastructure vulnerability management is outside OVERT scope
12.2 Lack of pen testing / red teaming	MEASURE	MEA-3	Partial	OVERT mandates third-party AI-specific testing; infrastructure penetration testing is outside scope
12.3 Lack of incident response	RESPOND	RES-1 through RES-5	Partial	Cryptographically gated incident response with attested escalation, override, revocation, and failure mode declaration within attested scope; IR team composition, training, and organization-wide incident response are outside scope
12.4 Unauthorized privileged access	HITL, ATTEST	HITL-4, ATT-5.3	Partial	Separation of duties attestation + notary independence requirements; access management systems are outside scope
12.5 Poor SDLC	—	—	—	No OVERT analog. Software development lifecycle is outside OVERT scope
12.6 Lack of compliance	GOVERN, IDENTIFY	GOV-1, GOV-3, IDE-2	Partial	Machine-readable governance policy + risk taxonomy + impact assessment support evidence for compliance; compliance determination is outside OVERT scope
12.7 Initial access	PROTECT	PRO-1, PRO-2	Partial	Boundary enforcement + network isolation attestation (NETATT); initial access prevention is primarily an infrastructure security concern
Agentic AI (Component 13)				
13.1 Memory poisoning	ATTEST, Agentic, PROTECT	ATT-1, ATT-4, CAP-1, PRO-5	Partial	OVERT can attest provenance of tool-fed context and enforce scoped retrieval within the attested boundary; poisoning of external memory stores or unmanaged prompt state remains outside scope
13.2 Tool misuse	TOOL, Agentic	TOOL-1 through	Direct	Pre-execution policy checks, per-call authorization, capability scoping, rate limits, and attestation receipts directly address unsafe

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
		TOOL-5, CAP-1, CAP-2		or unauthorized tool invocation within scope
13.3 Privilege compromise	Agentic, TOOL, PROTECT	CAP-1, CAP-2, TOOL-2, PRO-5	Direct	Capability-scoped identities, function authorization, and provenance-aware data isolation constrain privilege escalation and prove access decisions
13.4 Resource overload	PROTECT, RESPOND, TOOL	PRO-3, TOOL-3, RES-1	Direct	Rate limiting, circuit breaking, and adaptive responses are directly attestable
13.5 Cascading hallucination attacks	MEASURE, RESPOND, DRIFT, TOOL	MEA-2, RES-1, DRIFT-3, TOOL-3.4	Partial	OVERT can measure violation-rate shifts, detect recursive loops, and trip circuit breakers, but it cannot guarantee prevention of upstream hallucinations
13.6 Intent breaking and goal manipulation	PROTECT, TOOL, Agentic	PRO-1, PRO-4, TOOL-1, CAP-2	Partial	Boundary enforcement, input/output policy checks, and scoped capabilities reduce manipulation risk, but effectiveness depends on policy quality and mediation completeness
13.7 Misaligned and deceptive behaviors	MEASURE, Agentic, HITL	MEA-2, MEA-4, CAP-2, HITL-4	Partial	OVERT can surface anomalous behavior, require approvals, and constrain roles, but latent model misalignment is not fully solved by attestation
13.8 Repudiation and untraceability	ATTEST, TOOL	ATT-1, ATT-4, TOOL-5	Direct	Receipts, transparency logs, and attested action traces provide non-repudiation within scope
13.9 Identity spoofing and impersonation	HITL, ATTEST, TOOL	HITL-1.2, HITL-4.3, ATT-1, TOOL-2	Partial	OVERT can attest authenticated identities and approval chains in receipts, but identity provider controls remain external
13.10 Overwhelming human in the loop	HITL, RESPOND	HITL-1 through HITL-4, RES-1	Partial	OVERT can force approvals, escalations, and threshold-triggered intervention, but it cannot by itself optimize human workload or review quality
13.11 Unexpected RCE and code attacks	PROTECT, ATTEST, TOOL	PRO-2, ATT-2.2, TOOL-2	Direct	Isolation, binary identity attestation, and per-call authorization directly address code execution risk within attested execution boundaries

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
13.12 Agent communication poisoning	ATTEST, Agentic, PROTECT	ATT-1, ATT-4, MULTI-1, MULTI-2, PRO-2	Partial	Inter-agent trust boundaries, message attribution, and network isolation provide strong evidence within the attested graph, but semantic correctness of upstream agent outputs remains partly external
13.13 Rogue agents in multi-agent systems	Agentic, DRIFT, RESPOND	CAP-2, MULTI-1, MULTI-2, DRIFT-3, RES-4	Direct	Graph-topology governance, bounded delegation, revocation, and circuit breakers directly constrain rogue-agent proliferation within scope
13.14 Human attacks on multi-agent systems	HITL, Agentic, TOOL	HITL-4, CAP-2, MULTI-1, TOOL-1	Partial	Approvals, bounded delegation, and policy gating help, but operator misuse and social engineering remain partly external
13.15 Human manipulation	HITL, PROTECT, TOOL	HITL-1 through HITL-4, PRO-4, TOOL-1, CAP-2	Partial	Consent, review, correction, and policy gating reduce manipulation, but human susceptibility cannot be eliminated by attestation alone
13.16 Prompt injection (MCP server)	PROTECT, TOOL	PRO-1, PRO-4, TOOL-1, TOOL-2	Partial	Guardrails and pre-execution policy checks reduce unsafe tool use, but prevention depends on policy and filter quality
13.17 Confused deputy (MCP server)	TOOL, Agentic	TOOL-2, CAP-1, CAP-2	Direct	Function authorization and capability scoping directly address deputy misuse within scope
13.18 Tool poisoning (MCP server)	GOVERN, ATTEST, TOOL	GOV-4, ATT-2.2, TOOL-2, ATT-5	Partial	Supply-chain governance and binary identity attestation help detect poisoned tools, but upstream dependency vetting remains external
13.19 Credential and token exposure (MCP server)	PROTECT, ATTEST	PRO-2, PRO-5, ATT-1.4	Partial	Non-egress isolation, data containment, and transport-state evidence help, but secret storage and rotation controls are largely outside OVERT scope
13.20 Insecure server configuration (MCP server)	PROTECT, ATTEST, GOVERN	PRO-2, ATT-2.2, GOV-1	Partial	OVERT can attest configuration state at mediation boundaries and runtime identity, but full server hardening remains external

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
13.21 Supply chain attacks (MCP server)	GOVERN, ATTEST	GOV-4, ATT-2.2, ATT-5	Partial	Third-party governance and attested binary identity support evidence, but supply-chain risk management is broader than runtime attestation
13.22 Excessive permissions and scope creep (MCP server)	Agentic, TOOL	CAP-1, CAP-2, TOOL-2	Direct	Least-privilege capabilities and per-call authorization directly address scope creep within the attested boundary
13.23 Data exfiltration (MCP server)	PROTECT, ATTEST, TOOL	PRO-2, PRO-5, ATT-1, TOOL-2	Direct	Non-egress enforcement, data isolation, and per-call receipts directly support exfiltration prevention within scope
13.24 Context spoofing and manipulation (MCP server)	ATTEST, PROTECT, TOOL	ATT-1, ATT-4, PRO-4, TOOL-1	Partial	Provenance and policy checks help identify spoofed context, but upstream context trustworthiness is only partly attestable
13.25 Insecure communication (MCP server)	PROTECT, ATTEST	PRO-2, ATT-1.4	Partial	Network isolation and transport-state attestation support evidence of secure channels; protocol implementation details remain external
13.26 Malicious server connection (MCP client)	GOVERN, TOOL, ATTEST	GOV-4, TOOL-2, ATT-5	Partial	Third-party governance and attested connection policy support evidence, but complete vendor validation remains external
13.27 Insecure credential storage (MCP client)	PROTECT	PRO-2, PRO-5	Partial	OVERT constrains egress and data exposure within scope, but client credential storage itself is outside the attested runtime boundary
13.28 UI/UX deception (MCP client)	HITL, DISC, TOOL	HITL-1, HITL-2, HITL-3, DISC-1, TOOL-1	Partial	OVERT can attest disclosure, consent, and review checkpoints, not full interface-design integrity
13.29 Insufficient server validation (MCP client)	GOVERN, ATTEST, TOOL	GOV-4, ATT-5, TOOL-2	Partial	Attested trust-chain and external-service governance help, but complete server-validation logic is external
13.30 Client-side data leakage	PROTECT, Agentic	PRO-4, PRO-5, CAP-1	Partial	Output filtering, data isolation, and capability scoping reduce leakage within scope; client endpoint controls remain external
13.31 Excessive permission granting	Agentic, TOOL	CAP-1, CAP-2, TOOL-2	Direct	Capability scoping and per-call authorization directly address overbroad grants

DASF Component / Risk	OVERT Domain	OVERT Controls	Coverage	Notes
13.32 Client-side code execution	PROTECT, ATTEST	PRO-2, ATT-2.2, TOOL-2	Partial	Isolation and binary identity attestation help where client execution occurs inside the attested boundary; unmanaged client code remains external
13.33 Insecure communication handling	PROTECT, ATTEST	PRO-2, ATT-1.4	Partial	Network isolation and transport-state evidence help, but client transport-implementation details remain external
13.34 Session and state management failures	ATTEST, Agentic, HITL	ATT-1, ATT-4, CAP-1, HITL-4	Partial	Receipts, provenance, scoped capabilities, and approval chains help govern session state within scope; full session lifecycle controls are broader
13.35 Update and patch management	GOVERN, ATTEST	GOV-4, ATT-2.2	Partial	Runtime identity attestation and supply-chain governance support version-control evidence, but patch management is broader than OVERT

29.2 Control-to-Control Crosswalk

The following table maps DASF mitigation controls to OVERT controls, grouped by functional category.

DASF Control	Description	OVERT Controls	Coverage	Notes
Identity and Access				
DASF 1	SSO with IdP and MFA	HITL-1.2, HITL-4.3	Adjacent	OVERT attests authenticated identity in governance receipts; does not prescribe or implement authentication mechanism
DASF 2	Sync users and groups	GOV-2	Adjacent	Organizational role attestation; OVERT does not manage identity provisioning
DASF 5	Object-level access control	CAP-1, TOOL-2	Direct	Capability-based access control + function authorization with per-action attestation receipts
DASF 43	Access control lists	CAP-1, TOOL-2	Partial	Capability-based access + provenance-aware authorization; ACL management is outside scope

DASF Control	Description	OVERT Controls	Coverage	Notes
DASF 57	Attribute-based access control	CAP-1, TOOL-2	Partial	Provenance tracking supports evidence for attribute-based policy enforcement
DASF 64	Limit AI agent access	TOOL-1 through TOOL-5, CAP-1, CAP-2	Direct	Per-action policy enforcement + capability scoping + architectural separation within attested scope
DASF 67	Federate authentication	HITL-1.2, HITL-4.3, ATT-1, IDENT-1	Direct	IDENT-1 attests the full identity delegation chain including originating principal, each token exchange, scope narrowing verification, and token lifecycle. Federation protocol implementation remains external
Network and Isolation				
DASF 3	IP access lists	PRO-2	Partial	OVERT attests network isolation state (NETATT); IP list enforcement is infrastructure-level
DASF 4	Private link	PRO-2, ATT-1.4	Partial	Non-egress architecture + TLS certificate pinning attestation; private link configuration is infrastructure-level
DASF 34	Model isolation	PRO-2, PRO-5	Direct	Network isolation + data isolation attestation with co-epoch binding
DASF 56	Restrict outbound connections	PRO-2	Direct	NETATT attests egress policy, network controller identity, and eBPF state at each epoch
DASF 62	Network segmentation	PRO-2	Partial	Network isolation attestation covers segmentation enforcement at the AI boundary; broader network segmentation is infrastructure-level
MCP and Agent Hosting				
DASF 68	Use securely hosted managed MCP servers	GOV-4, TOOL-2, ATT-5, PRO-2, MCP-1	Direct	MCP-1 attests managed server identity, transport security, governance metadata, and per-call routing. Vendor internal operations remain external

DASF Control	Description	OVERT Controls	Coverage	Notes
DASF 69	Securely host custom MCP servers	GOV-4, ATT-2.2, PRO-2, TOOL-2, MCP-2	Direct	MCP-2 attests binary identity, network isolation, per-call authorization, and configuration change detection for operator-hosted MCP servers. Server deployment lifecycle remains external
DASF 70	Securely connect to external MCP servers	TOOL-2, PRO-2, ATT-5, CAP-1, MCP-3	Direct	MCP-3 attests connection governance, allowlist enforcement, capability scoping, output filtering, and connection lifecycle. External server internal posture remains external
DASF 72	Securely store and reuse agent state	PRO-5, ATT-1, ATT-4, CAP-1, STATE-1	Direct	STATE-1 attests state sealing, integrity verification, hash-chained lineage, mutation provenance, and access scoping. Storage infrastructure security remains external
DASF 73	Register prompts	GOV-1, DISC-1.2, ATT-4, TOOL-1, STATE-2	Direct	STATE-2 attests prompt registration, session binding, change detection, prompt-to-action traceability, and change approval governance. Prompt content quality and engineering methodology remain external
Data Security				
DASF 6	Classify data	IDE-1.2, GOV-3	Direct	System categorization + machine-readable risk taxonomy
DASF 8	Encrypt data at rest	—	—	No OVERT analog; encryption at rest is infrastructure-level
DASF 9	Encrypt data in transit	PRO-2	Partial	OVERT attests TLS certificate pins in NETATT; does not implement encryption
DASF 16	Secure model features	PRO-5, CAP-1	Partial	Data isolation + provenance-aware capability enforcement; feature store security is outside scope
DASF 51	Secure data sharing	PRO-5, ATT-1	Partial	Data isolation + non-egress attestation; data sharing governance is outside scope
DASF 46	Store and retrieve embeddings securely	PRO-5, PRO-2	Partial	Data isolation + network isolation attestation; embedding storage security is infrastructure-level

DASF Control	Description	OVERT Controls	Coverage	Notes
DASF 58	Data filters and masking	PRO-4, PRO-5.3	Direct	Output filtering + PII detection attestation
Audit and Monitoring				
DASF 14	Audit data actions	ATT-1, ATT-4, TOOL-5	Direct	Cryptographic upgrade: OVERT produces tamper-evident, notary-signed audit records vs. conventional logs
DASF 37	Inference tables	ATT-1, ATT-4, TOOL-5	Direct	Per-action attestation receipts in transparency log vs. mutable inference tables
DASF 55	Monitor audit logs	ATT-1, ATT-4, TOOL-5	Direct	OVERT audit trail is append-only with inclusion proofs and split-view detection
DASF 65	Implement end-to-end AI traceability	ATT-1, ATT-4, TOOL-5, DISC-1.2	Direct	Per-action receipts, transparency logs, and AI Bills of Materials provide end-to-end traceability within the attested scope
DASF 71	Log and register AI agents	ATT-1, ATT-4, DISC-1.2, GOV-1	Direct	Receipts, transparency logs, governance records, and AI Bills of Materials directly support agent inventory and activity logging within scope
DASF 21	Monitoring dashboard	MEA-2, RES-1	Adjacent	S3P provides quantitative monitoring signals; OVERT does not prescribe dashboard implementation
DASF 35	Track model performance	MEA-2, MEA-4	Partial	S3P statistical safety signals + pre-deployment testing requirements provide runtime performance measurement. DASF 35 encompasses broader model performance tracking including accuracy benchmarks, feature drift, and retraining triggers that extend beyond OVERT runtime attestation scope
DASF 36	Monitoring alerts	RES-1, RES-2	Direct	Adaptive control loop + incident response attestation
Guardrails and Enforcement				

DASF Control	Description	OVERT Controls	Coverage	Notes
DASF 31	Secure model serving endpoints	PRO-1, PRO-2, PRO-3	Partial	Boundary enforcement + network isolation + rate limiting attestation for serving endpoints; endpoint hardening is infrastructure-level
DASF 54	Implement AI guardrails	PRO-1, PRO-4	Direct	OVERT proves guardrails executed (not just configured): per-action permit/deny receipts with policy reference
DASF 60	Rate limiting	PRO-3, TOOL-3	Direct	Rate limiting with attested enforcement receipts and circuit breaking
Model Lifecycle				
DASF 18	Govern model assets	GOV-1, DISC-1.2	Partial	Governance policy attestation + AI Bill of Materials; model asset governance beyond runtime is outside scope
DASF 19	ML lifecycle management	GOV-1, GOV-2, HITL-4	Partial	Governance + accountability + approval attestation; covers runtime governance, not full ML lifecycle
DASF 23	Register/version/deploy model	HITL-4, GOV-1	Partial	OVERT attests deployment approvals; model registry is outside scope
DASF 24	Model access control	CAP-1, TOOL-2, PRO-5	Direct	Capability-based access + authorization + data isolation
DASF 29	MLOps workflows	GOV-1, HITL-4	Partial	Policy and approval attestation within operational workflows; MLOps workflow design is outside scope
DASF 42	MLOps/LLMOps	GOV-1, GOV-2	Adjacent	Governance attestation for ML operations processes; MLOps/LLMOps platform design is outside scope
Evaluation and Testing				
DASF 38	Pen testing and red teaming	MEA-3	Partial	OVERT mandates third-party AI testing; infrastructure pen testing is outside scope
DASF 45	Evaluate models	MEA-3, MEA-4	Direct	Third-party and pre-deployment testing requirements with transparency log publication

DASF Control	Description	OVERT Controls	Coverage	Notes
DASF 49	Automated LLM evaluation	MEA-2	Partial	S3P provides automated, statistically rigorous quality measurement for runtime behavior. DASF 49 encompasses broader automated evaluation including pre-deployment benchmarks, evaluation dataset curation, and model comparison that extend beyond OVERT runtime attestation scope
Supply Chain and Compliance				
DASF 32	LLM provider management	GOV-4, ATT-5	Partial	Supply chain governance + notary governance for third-party verification; vendor management processes are outside scope
DASF 50	Platform compliance	GOV-1, GOV-3	Partial	Governance policy + risk taxonomy support evidence for compliance; compliance determination is outside OVERT scope
DASF 53	Third-party library control	GOV-4	Partial	Supply chain and third-party governance attestation; library vetting and vulnerability scanning are outside scope
Incident Response				
DASF 39	Incident response team	RES-1 through RES-5	Partial	Cryptographically gated incident response with attested escalation, override, revocation, and failure modes; IR team composition and training are outside scope
DASF 40	Internal access controls	HITL-4, ATT-5.3	Partial	Separation of duties + notary independence requirements; access control system implementation is outside scope
Human Oversight				
DASF 66	Use human-in-the-loop feedback	HITL-1 through HITL-4, RES-1	Partial	OVERT can attest consent, review, correction, approval, and escalation gates; broader feedback pipelines and product UX remain external
DASF 44	Event-triggered actions	RES-1	Direct	OVERT control loop triggers attested responses to threshold exceedances
DASF 48	Hardened ML runtime	ATT-2.2	Partial	Binary identity attestation proves runtime integrity via hardware-rooted measure-

DASF Control	Description	OVERT Controls	Coverage	Notes
				ment; runtime hardening implementation is outside scope
No OVERT Analog				
DASF 11	Capture and view data lineage	CAP-1	Partial	OVERT data provenance tracking covers tool-call data flows; full-lifecycle data lineage is outside scope
DASF 7, 10, 12, 13, 15, 17, 20, 22, 25, 26, 27, 28, 30, 33, 41, 47, 52, 59, 61, 63	Data quality checks, versioning, deletion, real-time data, EDA, training tracking, experiment tracking, representative data, RAG, fine-tuning, pre-training, model tags, encryption, secrets, secure SDLC, LLM comparison, source control, clean rooms, security training, software updates	—	—	These controls address training-time operations, data management, infrastructure security, or platform features outside the scope of runtime attestation. OVERT complements but does not replace these controls

29.3 Coverage Summary

DASF risks addressable by OVERT attestation:

- Section 29.1 intentionally distinguishes **Direct**, **Partial**, and **Adjacent** mappings. Those row-level qualifiers govern; aggregate percentages should not be read as equivalent to direct coverage or comprehensive DASF alignment.
- This revised crosswalk maps all **97** DASF v3.0 risk rows. **88 of 97** have at least a partial OVERT analog; **9 of 97** remain outside OVERT scope.
- All **35 Component 13 agentic risks** have at least a partial OVERT analog. The strongest alignment appears in runtime enforcement, tool-use governance, exfiltration controls, non-repudiation, circuit breaking, and multi-agent capability scoping.

DASF controls with OVERT analogs:

- Section 29.2 likewise distinguishes **Direct**, **Partial**, and **Adjacent** mappings. Combined mapping counts should not be interpreted as direct control equivalence.
- This revised crosswalk maps all **73** DASF controls. **58 of 73** have at least an adjacent or stronger OVERT analog; **15 of 73** remain outside OVERT scope.
- Six controls (DASF 67-70, 72-73) upgrade from Partial or Adjacent to Direct through the addition of MCP, STATE, and IDENT control families.

Key gaps — DASF risks not addressed by OVERT:

1. **Training-time attacks** (3.3 label flipping, 2.4 adversarial partitions, 5.1 experiment tracking): OVERT attests runtime behavior; training pipeline integrity is outside scope.
2. **Data lifecycle management** (1.5 versioning, 1.9 stale data, 2.3 raw data criteria): Operational data management concerns, not runtime governance.
3. **Platform infrastructure** (12.1 vulnerability management, 12.5 SDLC, 7.4 source code control): Traditional InfoSec controls outside AI attestation scope.
4. **Unmanaged secret and client surfaces** (13.19, 13.27, 13.30, 13.32-13.34): OVERT can attest bounded behavior within the mediated runtime, but credential storage, unmanaged clients, and full session lifecycle controls remain partly external.

Key gaps — OVERT controls not addressed by DASF:

1. **Cryptographic attestation** (ATT-1 through ATT-5): DASF has no equivalent to non-egress attestation, co-epoch binding, three-phase attestation, transparency logs, or notary governance. This is OVERT's primary contribution.
2. **Statistical safety measurement** (MEA-1, MEA-2): DASF recommends monitoring (DASF 21, 35, 36) but does not specify cryptographically verifiable, auditor-reproducible measurement with exact confidence intervals.
3. **Risk signals** (OVERT risk signals (see Section 4.6 and Annex D)): DASF has no equivalent verifier-usable signal architecture. OVERT provides quantitative, independently verifiable

runtime signals within the declared mediation scope for monitoring, audit, and external risk analysis.

4. **Evidence-grade agentic runtime control** (TOOL-1 through TOOL-5, CAP-1/CAP-2, MULTI-1/MULTI-2): DASF v3.0 adds meaningful agentic controls (64-73), especially around MCP, agent registration, state, and prompt governance; OVERT still provides per-action attestation, capability-based access control, inter-agent trust boundaries, and loop detection with independent verification of execution.
5. **Failure mode and override** (RES-3, RES-4, RES-5): DASF 39 covers incident response broadly; OVERT specifies attested emergency override, scoped revocation, and explicit fail-open/fail-closed declarations.

Differentiation summary: DASF v3.0 materially narrows the conceptual gap by explicitly modeling agents, MCP servers, MCP clients, agent state, and prompt governance. OVERT remains differentiated because it upgrades those runtime controls within the attested scope from configuration guidance to tamper-evident, independently assessable execution evidence. Organizations deploying in regulated environments can use DASF for risk identification and control selection, then use OVERT to prove that the most consequential runtime controls actually executed.

[Editor's note: Section 29.4 (Attestation Boundary Declaration) is normative and remains in the standard, renumbered as Section 22.10; it is not part of this companion.]

30. Crosswalk: IMDRF N93

OVERT attestation artifacts support evidence for the mapped requirements below. Coverage qualifiers indicate the degree of alignment: **Direct** (OVERT directly produces the required artifact), **Partial** (OVERT supports some aspects but not all), **Adjacent** (OVERT evidence is relevant context).

This crosswalk maps OVERT controls to the IMDRF N93 draft, Technical Framework for AI Life Cycle Management, scoped to the post-market lifecycle. N93 states of itself that it is neither regulation nor guidance and that it does not replace existing IMDRF SaMD or cybersecurity publications. OVERT does not make a device N93-compliant; OVERT conformance does not determine compliance with N93. Manufacturers cite Glacis as the implementation and cite the recognized standards OVERT builds on (RFC 6962, NIST AI RMF, ISO/IEC 42001); OVERT mechanisms enter a document like N93 only through a recognized body, with ISO/IEC 23053 the natural on-ramp. The non-egress architecture does not remove BAA obligations: consistent with Section 17.5 and Annex C.10, this standard does not assert that receiving cryptographic commitments derived from PHI falls outside

the Business Associate definition. N93 builds on GMLP N88:2025; its own traceability appendix ties every post-market lifecycle step to GMLP Guiding Principle 10, the single principle OVERT maps to most cleanly.

The controls-ran-not-clinical-accuracy boundary holds in every row: OVERT attests that controls executed and that the record is genuine. It does not attest clinical accuracy, model quality, or bias acceptability.

IMDRF N93 Section	Subject	OVERT Controls	Coverage	Notes
4.1 QMS	Post-market monitoring mechanisms; configuration management and traceability of model and data versions	GOV-1.3, ATT-2, ATT-4	Partial	OVERT evidences that monitoring ran and that policy and configuration versions are tamper-evident. It is not a quality management system and does not satisfy ISO 13485; the QMS remains the manufacturer's.
4.2.4 Deployment and post-market risks	Version control when the model or device is modified; performance degradation from changes in third-party general-purpose models; configuration drift	ATT-2, Section 18.6, STATE-2, MCP-1	Partial	OVERT attests model and configuration identity per epoch and detects change at the boundary. Detecting a provider-side model substitution behind an opaque API depends on what the provider exposes.
4.3 Human oversight	Ability to override or reverse automated output; post-market human monitoring for degradation; guarding against automation bias	HITL-2, HITL-3, DRIFT-5	Partial	OVERT attests that review, correction, and override events occurred, with identity binding, and flags when human-engagement quality degrades. It does not evaluate whether an override was clinically correct.
4.4 Cybersecurity	Logging, auditing, anomaly detection; post-market validation after updates; distinguishing poisoning from natural drift	PRO-1, PRO-2 (NETATT), ATT-1, ATT-2.2, DRIFT-2	Partial	Covers the AI-layer monitoring and integrity controls N93 lists: tamper-evident telemetry, notary-derived binary identity (detecting unauthorized code or model substitution), and drift signals. Host, network, and platform security are out of scope. OVERT evidences that drift occurred; it does not by itself distinguish malicious poisoning from natural drift.

IMDRF N93 Section	Subject	OVERT Controls	Coverage	Notes
5.5 Deployment: traceability and version control	Failure-event identification and root cause; rollback; UDIs or alternative auditable identifiers	ATT-1, ATT-4, RES-4, Section 4.8	Partial	Per-action receipts in the transparency log are a candidate for N93's "alternative auditable identifier." OVERT does not issue or manage UDIs and does not handle regulatory model-version registration.
5.6 Operations and monitoring: infrastructure, data, and logging	Detailed logs and audit trails of predictions, inputs, outputs, and metadata for traceability, accountability, and compliance; local and central records	ATT-1 through ATT-5, TOOL-5	Direct	The strongest fit. N93 requires the log to exist and be detailed; it does not require the log to be tamper-evident or independently verifiable. OVERT closes that gap, turning a self-maintained log into a record a third party can verify without protected-content egress. OVERT attests the metadata and integrity layer; the clinical content of the log remains the device's own output.
5.6 Performance degradation and drift detection	Distinguishing expected variation from drift in non-deterministic and generative models; prompt-response histories; prompt classification with risk-based escalation	PRO-4 (scanner and local classifier), MEA-2, DRIFT-2	Partial	OVERT measures policy-violation rate (with exact confidence intervals) and behavioral or distributional drift. It does not measure clinical-accuracy drift against a reference standard (e.g., AUC, sensitivity); N93's drift concept includes both.
5.6 Alerting and monitoring of advanced models	Thresholds on critical metrics; automated alerts; usage-aware monitoring; risk-based escalation	RES-1, PRO-3, TOOL-3, DRIFT-3	Partial	OVERT attests that alerts and bounded response actions executed. Alert-threshold design and clinical interpretation remain with the manufacturer.
5.7 Real-world performance evaluation	Investigating divergence between override rates and monitoring signals; performance indicators including human-AI team measures; subgroup performance	HITL-3, HITL-3.3, DRIFT-5	Partial	N93's override-rate-versus-monitoring-signal scenario is directly evidenced by attested HITL-3 override records and DRIFT-5 oversight-quality signals with tamper-evident provenance. OVERT does not compute clinical accuracy or subgroup performance; those require ground-truth labels outside its scope.
6 Transparency and labelling	Telling the user what risk controls and monitoring are in	GOV-5.6, DISC-1.3	Adjacent	OVERT supports a verifiable per-interaction receipt reference and a published attestation summary (coverage, drift, override fre-

IMDRF N93 Section	Subject	OVERT Controls	Coverage	Notes
	place; tailored transparency at each update, including PCCP changes			quency). It does not generate the clinical labelling content enumerated in N93 Appendix C.

Out of scope. The following N93 areas are pre-deployment, data-centric, or model-quality concerns to which OVERT evidence does not apply: 5.1 planning and design and 5.3 model building and tuning (architecture choice, explainability, feature engineering); 5.2 data collection and management (training-data representativeness, bias mitigation, training-data lineage); 5.4 verification and validation, including clinical evaluation; and Appendix B model-performance metrics (accuracy, precision, recall, AUC, perplexity, coherence). OVERT operates at the evidence-integrity layer, not the model-performance layer.

HIPAA record retention (migrated reference). The session-lifecycle controls (SESS-1 through SESS-5) produce tamper-evident, time-stamped session and state records relevant to the HIPAA documentation-retention requirement at 45 CFR §164.530(j) (retain required documentation for six years from creation or last-effective date). OVERT evidences that such records exist and are tamper-evident; it does not by itself establish a covered entity's or business associate's retention program, and — consistent with the §30 note above and Section 17.5 — the non-egress architecture does not remove BAA obligations. (This reference was migrated from an inline tag in the standard body during the v1.1 crosswalk consolidation; HIPAA is otherwise outside the companion's enumerated framework set.)

31. Major Framework Crosswalks

OVERT helps an organization produce the tamper-evident evidence these frameworks require. Each framework governs the organization, not the tool. OVERT is not itself certified or compliant with these frameworks.

CHAI Governance Playbooks

This crosswalk maps OVERT controls to the published CHAI Governance Playbooks. OVERT covers the evidentiary portions of the controls, not the organizational ones.

- **Monitoring and performance:** MEASURE, DRIFT (**Direct**)

- **Responsible data management and data boundary:** PROTECT (**Direct/Partial**)
- **Lifecycle and third-party management:** ATT cross-boundary attestation and GOVERN (**Partial**)
- **Risk and impact assessment records:** GOVERN (**Partial** - OVERT evidences that the assessed controls operated, not that the assessment was adequate)
- **Organizational structure and policy domains: Adjacent** (OVERT does not stand up a committee)

Joint Commission RUAIH

OVERT helps a health system produce the evidence each standard asks for, so the organization can pursue certification. OVERT is never "RUAIH-certified," because RUAIH certifies organizations, not tools.

- **Governance:** GOVERN records and disclosure (**Partial**)
- **Effective data management:** PROTECT (**Direct/Partial**)
- **Risk and bias reduction:** MEASURE, DRIFT (**Partial** - OVERT proves the monitoring ran, not that bias is acceptable)
- **Monitoring, evaluating and validating:** MEASURE, DRIFT (**Direct** - this is the strongest fit, OVERT is the tamper-evident monitoring record and the safety-event provenance)
- **Transparency, education and training: Adjacent** (training is organizational)

Databricks AI Governance Framework (DAGF)

- **AI Organization: Adjacent**
- **Legal and Regulatory Compliance:** GOVERN plus regulatory crosswalks (**Partial** - OVERT evidences controls, not legal conclusions)
- **Ethics, Transparency and Interpretability:** GOVERN disclosure and receipts (**Partial** - not model interpretability)
- **Data, AIOps and Infrastructure:** PROTECT, ATT, runtime evidence (**Direct** - strong fit)
- **AI Security / DASF:** PROTECT, scanner and arbiter (**Direct** - strong fit, maps to DASF control areas)