

Defensible 10

Annex E (Normative): D05-Data Security Architecture

Technical Standard

Standards Committee
6-1-2026



© 2026 **ISAUnited.org**. Non-commercial use permitted under CC BY-NC. Commercial integration requires ISAUnited licensing.

Obsolete and withdrawn documents should not be used; please use replacements.

Copyright 2026. The Institute of Security Architecture United. All rights reserved

About ISAUnited

The Institute of Security Architecture United is the first dedicated Standards Development Organization (SDO) focused exclusively on cybersecurity architecture and engineering through a security-by-design approach. As an international support institute, ISAUnited helps individuals and enterprises unlock the full potential of technology by promoting best practices and fostering innovation in security.

Technology drives progress; security enables it. ISAUnited equips practitioners and organizations across cybersecurity, IT operations, cloud/platform engineering, software development, data/AI, and product/operations with vendor-agnostic standards, education, credentials, and a peer community—turning good practice into engineered, testable outcomes in real environments.

Headquartered in the United States, ISAUnited is committed to promoting a global presence and delivering programs that emphasize collaboration, clarity, and actionable solutions to today's and tomorrow's security challenges. With a focus on security by design, the institute champions the integration of security into every stage of architectural and engineering practice, ensuring robust, resilient, and defensible systems for organizations worldwide.

Obsolete and withdrawn documents should not be used; please use replacements.

Disclaimer

ISAUnited publishes the ISAUnited Defensible 10 Standards Technical Guide to provide information and education on security architecture and engineering practices. While efforts have been made to ensure accuracy and reliability, the content is provided “as is,” without any express or implied warranties. This guide is for informational purposes only and does not constitute legal, regulatory, compliance, or professional advice. Consult qualified professionals before making decisions.

Limitation of Liability

ISAUnited and its authors, contributors, and affiliates shall not be liable for any direct, indirect, incidental, consequential, special, exemplary, or punitive damages arising from the use of, inability to use, or reliance on this guide, including any errors or omissions.

Operational Safety Notice

Implementing security controls can affect system behavior and availability. First, validate changes in non-production, use change control, and ensure rollback plans are in place.

Third-Party References

This guide may reference third-party frameworks, websites, or resources. ISAUnited does not endorse and is not responsible for the content, products, or services of third parties. Access is at the reader’s own risk.

Use of Normative Terms (“Shall,” “Should,” “Must”)

- Must / Shall: A mandatory requirement for conformance to the standard.
- Must Not / Shall Not: A prohibition; implementations claiming conformance shall not perform the stated action.
- Should: A strong recommendation; valid reasons may exist to deviate in particular circumstances, but the full implications must be understood and documented.

Acceptance of Terms

By using this guide, readers acknowledge and agree to the terms in this disclaimer. If you disagree, refrain from using the information provided.

For more information, please visit our [Terms and Conditions](#) page.

Obsolete and withdrawn documents should not be used; please use replacements.

License & Use Permissions

The Defensible 10 Standards (D10S) are owned, governed, and maintained by the Institute of Security Architecture United (ISAUnited.org).

This publication is released under a Creative Commons Attribution–NonCommercial License (CC BY-NC).

Practitioner & Internal Use (Allowed):

- You are free to download, share, and apply this standard for non-commercial use within your organization, departments, or for individual professional, academic, or research purposes.
- Attribution to ISAUnited.org must be maintained.
- You may not modify the document outside of Sub-Standard authorship workflows governed by ISAUnited, excluding the provided Defensible 10 Standards templates and matrices.

Commercial Use (Prohibited Without Permission):

- Commercial entities seeking to embed, integrate, redistribute, automate, or incorporate this standard in software, tooling, managed services, audit products, or commercial training must obtain a Commercial Integration License from ISAUnited.

To request permissions or licensing:
info@isaunited.org

Standards Development & Governance Notice

This standard is one of the ten Parent Standards in the Defensible 10 Standards (D10S) series. Each Parent Standard is governed by ISAUnited's Standards Committee, peer-reviewed by the ISAUnited Technical Fellow Society, and maintained in the Defensible 10 Standards GitHub repository for transparency and version control.

Contributions & Collaboration

ISAUnited maintains a public GitHub repository for standards development.

Practitioners may view and clone materials, but contributions require:

- ISAUnited registration and vetting
- Approved Contributor ID
- Valid GitHub username

All Sub-Standard contributions must follow the Defensible Standards Submission Schema (D-SSF) and are peer-reviewed by the Technical Fellow Society during the annual Open Season.

Obsolete and withdrawn documents should not be used; please use replacements.

Abstract

The ISAUnited Defensible 10 Standards provide a structured, engineering-grade framework for implementing robust and measurable cybersecurity architecture and engineering practices. The guide outlines the frameworks, principles, methods, and technical specifications required to design, build, verify, and operate reliable systems.

Developed under the ISAUnited methodology, the standards align with modern enterprise realities and integrate Security by Design, continuous technical validation, and resilience-based engineering to address emerging threats. The guide is written for security architects and engineers, IT and platform practitioners, software and product teams, governance and risk professionals, and technical decision-makers seeking a defensible approach that is testable, auditable, and scalable.

This document includes a series of Practitioner Guidance, Cybersecurity Students & Early-Career Guidance, and Quick Win Playbook callouts.



Practitioner Guidance- Actionable steps and patterns to apply the technical standards in real environments.



Cybersecurity Student & Early-Career Guidance- Compact, hands-on activities that turn each section's ideas into a small, verifiable artifact.



Quick Win Playbook- Immediate, evidence-driven actions that improve posture now while reinforcing good engineering discipline.

Together, these elements help organizations translate intent into engineered outcomes and sustain long-term protection and operational integrity.

Obsolete and withdrawn documents should not be used; please use replacements.

Foreword

Message from ISAUnited Leadership

Cybersecurity is at a turning point. As digital systems scale, reactive and checklist-driven practices do not keep pace with adversaries. The ISAUnited position is clear: security must be practiced as engineered design, grounded in scientific principles, structured methods, and defensible evidence. Our mission is to professionalize cybersecurity architecture and engineering with standards that are actionable, testable, and auditable.

ISAUnited Defensible 10 Standards: First Edition is a practical framework for that shift. The standards in this book are not theoretical. They translate intent into measurable specifications, controls, and verification, and enable teams to design and operate resilient systems at enterprise scale.

About This First Edition

This edition publishes 10 Parent Standards, one for each core domain of security architecture and engineering. Sub-standards will follow in subsequent editions, contributed by ISAUnited members and reviewed by our Technical Fellow Society, to provide focused, technology-aligned detail. Adopting the Parent Standards now positions organizations for seamless integration of Sub Standards as they are released on the ISAUnited annual update cycle.

Why “Defensible Standards”

Defensible means the work can withstand technical, operational, and adversarial scrutiny. These standards are designed to be demonstrated with evidence, featuring clear architecture, measurable specifications, and verification, so that practitioners can confidently stand behind their designs.

Obsolete and withdrawn documents should not be used; please use replacements.

Contents

Section 1. Standard Introduction.....	10
Section 2. Definitions	11
Section 3. Scope.....	14
Section 4. Use Case	16
Section 5. Requirements (Inputs)	18
Section 6. Technical Specifications (Outputs)	20
Section 7. Cybersecurity Core Principles.....	25
Section 8. Foundational Standards Alignment.....	27
Section 9. Security Controls	29
Section 10. Engineering Discipline	32
Section 11. Associate Sub-Standards Mapping.....	37
Section 12. Verification and Validation (Tests)	41
Section 13. Implementation Guidelines	47
Appendices.....	53
Appendix A: EP-05 Engineering Traceability Matrix (ETM).....	53
Appendix B: EP-05 Evidence Pack Matrix	56

Obsolete and withdrawn documents should not be used; please use replacements.

Annex E (Normative): D05-Data Security Architecture

Obsolete and withdrawn documents should not be used; please use replacements.

ISAUnited's Defensible 10 Standards**Parent Standard:** D05-Data Security Architecture**Document:** ISAU-DS-DS-1000_v2**Last Revision Date:** June 2026**Peer-Reviewed By:** ISAUnited Technical Fellow Society**Approved By:** ISAUnited Standards Committee

Obsolete and withdrawn documents should not be used; please use replacements.

Section 1. Standard Introduction

The Data Security Architecture Parent Standard (ISAU-DS-DS-1000) establishes the engineering baseline for safeguarding structured, semi-structured, and unstructured data across its lifecycle—at rest, in transit, and in use—spanning on-premises, cloud, hybrid, and multi-cloud environments. As a Parent Standard, it defines common terminology, scope, Requirements (inputs), Technical Specifications (outputs), and Verification and Validation expectations that subordinate sub-standards inherit. It is vendor-neutral and implementation-agnostic, aligning with recognized NIST and ISO/IEC foundations while extending them with normative, testable specifications. The intent is to provide a defensible, measurable, and auditable approach to data security that links protections to classification and ownership and demonstrates integrity and recoverability.

Objective

This standard defines the foundational principles for Data Security Architecture (ISAU-DS-DS-1000), engineered to keep protections bound to the data itself—wherever it resides or flows. It provides data security architects, platform and storage engineers, data stewards and owners, SOC and incident response teams, and backup and disaster-recovery engineers with a structured, defensible methodology for governing access, protecting data, monitoring activity, and recovering with evidence.

Emphasis is placed on:

- Enforce data classification and ownership as the first step; automated discovery and tagging drive downstream controls and auditability.
- Implementing Zero-Trust Data Access (ABAC) with purpose-bound decisions; requiring MFA/JIT for privileged data actions and deny-by-default for sensitive classes.
- Requiring encryption by policy for data at rest and in transit, governed centrally through KMS/HSM; cryptographic parameters and agility are defined in the CEK Parent Standard.
- Preventing loss and misuse with tag-driven DLP across endpoint, network, cloud, and SaaS channels; correlating events in SIEM for end-to-end traceability.
- Proving integrity and resilience through immutable (WORM) backups, cryptographic verification, multi-region redundancy, and tested restores to RTO/RPO.
- Emitting structured, tamper-evident data access and modify telemetry with correlation identifiers to support detection, investigations, and incident response.
- Producing evidence—catalog and tag exports, ABAC policies and decision logs, encryption posture and KMS logs, DLP incidents and tests, immutability settings and restore drills—that makes data security measurable and auditable.

By integrating these engineering-focused capabilities, the standard provides an actionable, measurable, and defensible framework for securing data across databases,

Obsolete and withdrawn documents should not be used; please use replacements.

warehouses, and lakehouses; object and file stores; SaaS platforms; analytics and AI pipelines; messaging and streaming; edge; and archives.

Justification

Modern enterprises operate distributed data ecosystems—object stores, lakehouses, SaaS analytics, and edge pipelines—that change rapidly. Breaches frequently stem not from missing encryption, but from inconsistent classification, weak or unscoped entitlements, unmanaged shadow data, permissive sharing links, incomplete telemetry, and backups that can be altered before recovery.

Policy-only governance and product-first deployments are insufficient. Protections must be engineered as data-centric behaviors, enforced as policies-as-code, validated against explicit acceptance thresholds, and evidenced with durable artifacts. This standard addresses these realities by unifying discovery and classification, purpose-bound access, encryption-by-policy with centralized key governance, tag-driven DLP, tamper-evident logging with anomaly detection, and immutable, testable recovery. Detailed cryptography, application-layer, and delivery mechanics are deferred to their respective parents (CEK, Application Security, and DevSecOps). Through structured requirements, measurable outputs, and rigorous validation, the standard enables teams to proactively protect data, reduce the risk of unauthorized access and exfiltration, accelerate investigations, and ensure reliable recovery across hybrid and multi-cloud deployments.

Section 2. Definitions

ABAC (Attribute-Based Access Control) – Authorization model using subject, object, action, and environmental attributes (including sensitivity tags) to grant or deny access.

ADR (Architecture Decision Record) – A structured decision artifact that records the problem, options, assumptions, trade-offs, decision, and related evidence plan for a data security architecture choice.

Anonymization – an irreversible transformation that prevents re-identification by reasonable means; distinct from pseudonymization.

CEK (Architectural Reference) – ISAUnited’s Cryptography, Encryption & Key Management Parent Standard governing cryptographic profiles, algorithms/modules, and key lifecycles consumed by this annex.

Data Access Event – A normalized, tamper-evident record containing at a minimum: timestamp, subject, source, object, action, result, purpose, trace_id, and policy_decision (allow/deny).

Obsolete and withdrawn documents should not be used; please use replacements.

Data Catalog – The authoritative registry of datasets, metadata, owners/stewards, locations, lineage, and tags used for discovery, access decisions, and audit.

Data Classification – The engineering process of assigning sensitivity and handling categories to datasets (e.g., Public, Internal, Confidential, Restricted) that drive controls for access, retention, protection, and monitoring.

Data Egress Control – Allowlist-based, auditable controls governing movement of data from trusted zones to external destinations, tenants, or sharing mechanisms; out-of-policy egress is denied and logged.

Data Erasure / Cryptographic Erasure – Verified destruction of data through secure deletion processes or by destroying/invalidating keys protecting encrypted data.

Data Integrity – Assurance that data is complete, accurate, and unaltered except by authorized, logged actions; commonly verified via cryptographic hashes or authenticated checks.

Data Lineage / Provenance – End-to-end trace of dataset origin, transformations, copies, and consumers; proves where data came from and how it changed across systems.

Data Masking – Obfuscation of sensitive fields (irreversible or reversible) for non-production use, analytics, or support while preserving utility.

Data Minimization – Collect, process, and retain only what is necessary for declared purposes; remove, mask, or tokenize excess fields.

Data Owner / Data Steward – Roles accountable for classification, retention, and policy enforcement (owner) and day-to-day data quality/metadata lifecycle (steward).

Data Residency / Sovereignty – Physical or jurisdictional constraints on where data is stored/processed and which legal regimes apply; enforced technically in this annex.

Differential Privacy (DP) – Technique that adds calibrated noise to outputs to limit re-identification risk while preserving statistical utility.

DLP (Data Loss Prevention) – Policy-driven inspection and control of sensitive data across channels (endpoint, network, storage, SaaS) to prevent unauthorized egress or misuse.

Encryption (Architectural Reference) – Protection of data in transit/at rest using approved cryptography with centrally governed keys; algorithm and module specifics are defined in CEK.

Obsolete and withdrawn documents should not be used; please use replacements.

Evidence Pack (EP) / Evidence Pack ID – The signed, versioned artifact bundle—and its identifier—that proves conformance for a dataset/service and release/change (this annex uses EP-05, EP-05.0, EP-05.1, EP-05.2, EP-05.3).

High-Value Dataset (HVD) – Designation for collections whose compromise would materially impact safety, financials, or mission; subject to elevated controls and monitoring.

ICD (Interface Control Document) – A technical artifact that documents a data interface (access pattern, identity type, ABAC context, sensitivity tag use, telemetry fields, and invariants) and how it is enforced and tested.

JIT (Just-in-Time) Data Access – Time-bounded elevation for privileged data actions requiring explicit approval, MFA, and automatic revocation with full audit.

KMS / HSM (Architectural Reference) – Central key services and hardware modules providing generation, storage, rotation, and policy enforcement for keys used by data systems (parameters in CEK).

Legal Hold – Suspension of routine retention/disposal for datasets subject to investigation; overrides standard deletion schedules with auditable control.

MFA (Multi-Factor Authentication) – Authentication requiring two or more independent factors; mandated for privileged data actions in this annex.

MTTD (Mean Time To Detect) – The average elapsed time to detect anomalous or suspicious activity; a detection SLO used in §6.6 and §12.

Policy-as-Code (Data) – Declarative policies for access, retention, masking/tokenization, and DLP are compiled and enforced by engines across platforms.

Pseudonymization – Replacement of identifiers with reversible tokens held under separate controls; reduces exposure while preserving linkage under strict conditions.

Retention Schedule – Time-bound policies (by class/purpose) dictating how long data is kept and when/who may dispose of it.

RTO / RPO (Data) – Recovery Time Objective / Recovery Point Objective for datasets and services; quantitative targets proven via drills.

SDD (Sensitive Data Discovery) – Automated scanning of repositories/streams to locate and tag sensitive elements (PII, PHI, PCI, secrets) with confidence scores and owners.

Sensitivity Tag – A machine-readable label attached to data objects/streams indicating classification, residency, owner, and retention—consumed by policies (ABAC, DLP, retention).

Obsolete and withdrawn documents should not be used; please use replacements.

SIEM (Security Information and Event Management) – Aggregation, correlation, and analysis of security events and telemetry (e.g., DLP incidents and data access logs) to support detection and investigation.

SLO (Service Level Objective) – A measurable target (e.g., discovery coverage, denial rate, DLP efficacy, MTTD, restore attainment) used to verify outcomes in §6 and §12.

Shadow Data – Unmanaged copies (exports, sandboxes, sync caches) outside approved controls or catalogs.

Write Once, Read Many (WORM) – Write Once, Read Many storage, aka Immutable Storage, enforces non-alterable retention windows for backups/archives and evidentiary records; modify/delete attempts during retention are denied and logged.

Tamper-Evident Logging – Logging architecture that prevents undetected alteration (e.g., append-only/WORM plus hash chaining and external time-stamping).

Tokenization – The substitution of a sensitive value with a token that lacks intrinsic meaning; the original values are retrieved only through controlled, audited detokenization.

Zero-Trust Data Access (ZTDA) – Continuous, context-aware evaluation of requests to data—no implicit trust from network location; decisions bound to classification, identity, device posture, risk, and purpose.

Section 3. Scope

Modern enterprise data estates span databases, warehouses/lakehouses, object and file stores, SaaS platforms, analytics and AI pipelines, streaming systems, edge locations, and archives. The scope of this standard covers data security architecture across on-premises, cloud, multi-cloud, and hybrid environments, including partner-managed exchanges and shared services.

This standard defines the architectural expectations and technical guardrails required to maintain protections that are bound to the data itself. It is designed to help practitioners classify and tag data, enforce purpose-bound access, require policy-based encryption, prevent out-of-policy egress, generate investigation-ready telemetry, and maintain verifiable recoverability—without duplicating responsibilities already defined by adjacent parent standards.

Obsolete and withdrawn documents should not be used; please use replacements.

Applicability

- Enterprise data platforms in production contexts across industries and sizes.
- Hybrid and multi-environment estates: databases, lakehouses, object stores, SaaS data planes, pipelines/streaming, edge, and archives.
- Greenfield and brownfield deployments: new builds, modernization, and third-party integrations where data is ingested, transformed, shared, retained, or erased.
- Practitioner roles focused on technical execution: data security architects, data stewards and owners, platform and storage engineers, SOC and incident response analysts, backup and disaster-recovery engineers.

Key Focus Areas

- **Discovery, Classification, and Governance:** Automated discovery with defined coverage targets; sensitivity tagging with owners and stewards; lineage; retention flags that drive enforcement.
- **Zero-Trust Data Access (ZTDA):** Attribute-based access control tied to classification and purpose; MFA/JIT for privileged actions; deny-by-default on sensitive classes; decision logging with purpose.
- **Encryption and Key Governance (architectural scope):** Encryption by policy for data at rest and in transit with centralized KMS/HSM governance; cryptographic profiles and parameters are defined in the CEK parent standard.
- **Data Loss Prevention and Monitoring:** Tag-driven DLP across endpoint, network, and cloud/SaaS channels; SIEM correlation and anomaly analytics for misuse, exfiltration, and insider risk.
- **Integrity and Resilience:** Immutable storage (WORM) for backups and archives, cryptographic integrity verification, multi-region redundancy, and restore drills aligned to RTO/RPO.
- **Logging, Audit, and Anomaly Detection:** Tamper-evident access and modification of logs with the standard schema (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision); analytics to surface suspicious patterns and support end-to-end reconstruction.
- **Data Lifecycle Management:** Minimization, retention, archival, and secure deletion or erasure tied to classification and residency; demonstrable enforcement in platforms and policies.
- **Privacy-preserving Techniques (technical scope):** Masking, tokenization, and differential privacy are embedded in data paths where design requires; residency and sovereignty constraints are enforced technically.

Obsolete and withdrawn documents should not be used; please use replacements.

Outcomes

By defining this scope, the standard ensures that the data security architecture is:

- **Defensible:** Controls are bound to classification and ownership, documented, auditable, and centrally governed, with CEK alignment for cryptographic posture.
- **Measurable:** Performance is validated through discovery and tagging coverage, out-of-policy denial rates, DLP efficacy with stated FP/FN bounds, integrity checks, and restore-drill attainment.
- **Adaptive:** Protections follow data across platforms, copies, jurisdictions, and processing styles (analytics, AI, edge) without redesign.
- **Aligned:** Interfaces cleanly with adjacent domains (IAM, CEK, MDR, DevSecOps) through clear technical boundaries and shared evidence identifiers.

This scope provides the foundation for resilient, engineering-grade data protection that preserves confidentiality, integrity, and availability while enabling analytics, innovation, and operational agility.

Section 4. Use Case

Achieving resilient data security requires engineered practice across data stores and flows—not just policy. The following consolidated use case reflects a complex enterprise with distributed data platforms. It exposes common data-layer weaknesses, maps them to data-centric controls, and defines measurable outcomes. This links architectural decisions directly to defensible, auditable results.

Table E-1. Use Case:

Use Case Name	Preventing sensitive data exfiltration and ensuring recoverability at scale.
Objective	Classify/tag sensitive datasets; enforce Zero-Trust Data Access (ABAC + MFA/JIT); prevent out-of-policy movement (DLP); verify integrity and rapid recovery using immutable backups.
Scenario	A global enterprise maintains customer PII and financial data across data lakes, warehouses, object stores, and SaaS analytics. Findings: inconsistent classification, permissive sharing links, unmanaged sandbox copies, in-place edits on backups, limited visibility into high-risk access.
Actors	Data Security Architect; Data Steward; Platform/Storage Engineer; SOC/IR Analyst; Backup/DR Engineer.

Obsolete and withdrawn documents should not be used; please use replacements.

Challenges Identified	Incomplete discovery/classification coverage; ABAC gaps on sensitive tables/objects; ad-hoc exports via SaaS sync and email; backups alterable during the retention window; sparse or non-tamper-evident access/modify logs.
Technical Solution	Discovery & Classification. Automated scans across in-scope repositories; apply sensitivity tags with owners; map retention flags to classes. Zero-Trust Data Access (ZTDA). ABAC with purpose context; MFA/JIT for privileged data actions; deny-by-default for sensitive classes; log policy_decision and trace_id on each access. Encryption & Key Governance (architectural). Encrypt by policy for data at rest and in transit per CEK cryptographic profiles; enforce centralized KMS/HSM governance with rotation/use logging. DLP & Monitoring. Tag-driven DLP across endpoint, network, and cloud/SaaS; SIEM correlation and anomaly analytics. Integrity & Resilience. Immutable backups (WORM) with multi-region copies; cryptographic integrity verification; periodic encrypted restore drills to RTO/RPO. Logging & Auditability. Tamper-evident centralized logs using the standard schema and append-only/WORM storage.
Expected Outcomes (SLOs)	Discovery coverage ≥ 98 % and tagging latency ≤ 24 hours; out-of-policy access denial rate ≥ 99 % on sensitive classes with MFA/JIT for privileged actions; simulated exfiltration blocked ≥ 95 % with documented FP/FN bounds; immutable backups deny alteration and quarterly restores meet RTO/RPO; 100 % sensitive-class access/modify events logged and correlated to identities.
Evidence	Classification policy and discovery reports; tag exports and owner maps; ABAC policy definitions and decision logs; MFA/JIT elevation trails; storage/transport posture scans; KMS rotation and key-use logs; DLP incidents/tests with FP/FN; SIEM dashboards/queries; immutable backup configs and denied-alter events; restore drill reports—filed under the Evidence Pack ID.

Key Takeaways

- Tags before access. Discovery and sensitivity tagging precede any control; downstream ABAC, DLP, retention, and monitoring consume tags as truth (§5.1–§5.2, §6.1).
- Zero Trust is enforceable. Deny-by-default on sensitive classes with purpose-bound ABAC, MFA/JIT for privileged actions, and per-request policy_decision and trace_id in logs (§6.2, §6.6).
- Encryption is policy. Data at rest and in transit is encrypted per CEK cryptographic profiles with centralized KMS/HSM governance and auditable key-use/rotation logs (§6.3).
- DLP must be tag-driven and correlated. Endpoint, network, and cloud/SaaS DLP policies key off sensitivity tags and are correlated in SIEM (§6.4, §6.6).
- Recovery is proven. WORM-enforced backups, integrity verification jobs, multi-region copies, and routine encrypted restore drills demonstrate RTO/RPO attainment (§6.5).

Obsolete and withdrawn documents should not be used; please use replacements.

- Telemetry is investigation-ready. Standardized events (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision) are stored in append-only/WORM and forwarded to SIEM (§6.6).
- Outcomes are measured. Coverage $\geq 98\%$, tagging latency ≤ 24 hours, out-of-policy denials $\geq 99\%$, DLP block rate $\geq 95\%$ (with FP/FN bounds), restore drills meet RTO/RPO (§6, §12).
- Evidence binds it all. Each change references an Evidence Pack ID with catalog/tag exports, ABAC decisions, DLP incidents/tests, KMS logs, WORM configs, and restore drill reports (§12).

Section 5. Requirements (Inputs)

These inputs are baseline preconditions. They Must exist before teams design, implement, or validate controls in this Parent Standard or any aligned sub-standard. Each input should have an owner, a status, and a proof link (Evidence Pack ID).

5.1 Enterprise Data Catalog & Classification Framework

A centralized, enforced data catalog Must exist with authoritative classification/tags for datasets (structured, semi-structured, unstructured), recorded owners/stewards, lineage pointers, and retention attributes. Catalog scope Must include cloud-native data stores and SaaS repositories within the defined estate, so discovery and tagging coverage targets apply uniformly across on-premises and cloud environments.

5.2 Sensitivity Tags Bound to Controls

Sensitivity tags Must be machine-readable and drive policy: ABAC rules, DLP policies, retention/erasure, masking/tokenization, and logging. Tag to control bindings are documented and testable.

5.3 Zero-Trust Data Access Baseline

An ABAC model Must be defined for sensitive classes (with purpose context); MFA/JIT elevation Must be required for privileged data actions; deny-by-default applies outside declared purpose.

5.4 Encryption Policy & KMS Integration (architectural)

Enterprise encryption policies Must require encryption at rest and in transit and integrate with centralized KMS/HSM for key governance. *(The CEK Parent Standard governs cryptographic profiles, algorithms/modules, and key lifecycles.)*

5.5 DLP Channel Coverage & Policy Registry

DLP capabilities Must be deployed across endpoint, network, and cloud/SaaS channels with a registry of policies keyed by sensitivity tags and data types; incident routing and tuning procedures are defined.

Obsolete and withdrawn documents should not be used; please use replacements.

5.6 Centralized Data Access Logging & Schema

All access/modify events on sensitive classes Must be logged with the standard schema—timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision—and aggregated into tamper-evident storage with SIEM onboarding. *Ingest schema conformance for required fields is 100 %.*

5.7 Immutable Backup & Recovery Objectives

Backups of critical datasets Must be immutable (WORM) for defined retention, have multi-region copies, and have documented RTO/RPO targets with a scheduled encrypted restore-drill cadence.

5.8 Residency & Processing Constraints (technical)

Residency/sovereignty constraints Must be enforced technically in platforms and data paths; retention and hold flags Must be mapped to classifications and systems and verified in tests. *(No policy/regulatory narrative here; technical enforcement only.)*

5.9 Discovery & Lineage Coverage Targets

Automated discovery and lineage tooling Must be in place with coverage targets (for example, $\geq 98\%$ of in-scope stores scanned/tagged; new/changed data tagged within ≤ 24 hours) and named owners for gaps.

5.10 Evidence & Metrics Readiness

Dashboards/queries Must exist for recurring checks (classification coverage/latency, ABAC denials, DLP incidents, KMS rotation, restore drills). Evidence sources (catalog exports, policy bindings, logs, drill reports) are pre-staged under an Evidence Pack ID.



Practitioner Guidance:


- Readiness gate (one page). List 5.1–5.10 with owners, status, and links to the Evidence Pack ID; do not proceed to design until all are green.
- Blockers policy. If 5.1/5.2 (catalog/tags to controls) or 5.3/5.4 (ABAC/MFA/JIT, encryption + KMS) is missing, halt solutioning and open a tracked risk—downstream controls are not defensible without them.
- Traceability first. For each sensitivity tag, document its bound controls (ABAC, DLP, retention/erasure, masking/tokenization, logging) and the named test that proves the binding; store artifacts in the Evidence Pack.
- Measure continuously. Establish baseline metrics now (coverage $\geq 98\%$, tagging latency ≤ 24 hours, out-of-policy denial rate $\geq 99\%$, DLP block/quarantine rate $\geq 95\%$ with FP/FN bounds, restore attainment to RTO/RPO) to compare against §6 SLOs during V&V.

Obsolete and withdrawn documents should not be used; please use replacements.

Section 6. Technical Specifications (Outputs)

These specifications translate the data security policy into measurable, testable data-layer behavior. Outputs are enforced in data platforms, catalogs/policy engines, gateways proximate to data access, and operational controls under the data/security teams' control. (Annex J governs pipeline gates, provenance/SBOM, and promotion controls; Annex I (CEK) governs cryptographic profiles and key lifecycles.)

Figure 6-1. This illustrates the Defensible Data Security output model for Section 6. The model shows how data governance, classification, access policy, privileged data actions, protected repositories, encryption, key management, DLP, integrity controls, telemetry, automated response, continuous assessment, and validation activities work together to establish measurable, enforceable, and auditable data security outcomes.



Obsolete and withdrawn documents should not be used; please use replacements.

Outputs must be:

- **Measurable:** validated by scans, logs, audits, or tests
- **Actionable:** implementation-ready, not policy slogans
- **Aligned:** traceable to §5 Requirements and sub-standards

6.1 Data Classification & Governance

- Automated discovery Must scan in-scope repositories (structured, semi-structured, unstructured), including cloud-native and SaaS repositories, and assign sensitivity tags with owners/stewards.
- Coverage SLO: discovery/tagging coverage $\geq 98\%$ of in-scope stores; new/changed data tagged within ≤ 24 hours.
- Shadow Data discovery Must identify unmanaged copies (exports, sandboxes, sync caches) of sensitive classes and route them to remediation or containment workflows within a defined operational window.
- Target: newly detected shadow copies are tagged or quarantined within ≤ 24 hours.
- Retention enforcement Must be bound to classification (including legal hold overrides) and be auditable per dataset.
- Evidence: discovery reports, tag exports, owner maps, retention-enforcement proofs (Evidence Pack ID).

6.2 Secure Data Access Controls (ZTDA)

- ABAC Must enforce deny-by-default for sensitive classes; decisions include purpose context.
- MFA/JIT Must be required for privileged data actions; elevation is time-bounded with automatic revocation and a complete audit trail.
- Denial SLO: out-of-policy access requests to sensitive classes are denied $\geq 99\%$ in validation scenarios.
- Evidence: ABAC policy definitions, MFA/JIT settings, denial/approval logs mapped to tags (Evidence Pack ID).

6.3 Encryption & Data Protection (Architectural Scope)

- Encryption by policy: sensitive data at rest and in transit Must be encrypted per CEK cryptographic profiles.
- Central KMS/HSM Must govern key lifecycle; rotation/usage logging Must meet CEK policy thresholds. (*Algorithms/modules and profiles are defined in the CEK Parent Standard.*)
- Resilience SLO: encrypted restore drills meet RTO/RPO; no plaintext export paths for protected classes.
- Evidence: transport/storage posture scans, KMS rotation/usage logs, restore-drill reports (Evidence Pack ID).

6.4 Data Loss Prevention & Monitoring

- DLP controls Must operate across endpoint, network, and cloud/SaaS channels with policies driven by sensitivity tags and data types.
- Data Egress Control Must enforce allowlisted export and sharing paths for sensitive classes; out-of-policy egress attempts Must be denied and logged; approved exceptions Must be correlated in SIEM with data-access events.

Obsolete and withdrawn documents should not be used; please use replacements.

- Efficacy SLO: simulated exfiltration via email/web/cloud sync is blocked ≥ 95 %, with documented FP/FN bounds.
- Integration: DLP telemetry Must feed SIEM for correlation with data-access events and anomalies.
- Evidence: DLP policy registry, incident/test logs, SIEM correlation dashboards (Evidence Pack ID).

6.5 Data Integrity & Resilience

- Immutable storage (WORM) Must protect backups/archives of critical datasets for defined retention; attempts to alter/delete Must be denied and logged.
- Integrity verification Must use cryptographic hashing for critical files/objects with periodic verification jobs.
- Recovery SLO: quarterly restore exercises meet RTO/RPO; multi-region redundancy is active and tested.
- Evidence: WORM configurations and denied-alter logs, hash-verification reports, restore-drill outputs (Evidence Pack ID).

6.6 Logging, Auditing & Anomaly Detection

- Standard event schema: access/modify events for sensitive data Must include timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision and be aggregated to a tamper-evident store.
- Detection SLO: anomalies on sensitive classes (after-hours, unusual volume/source) are detected within ≤ 15 minutes (MTTD), with investigations able to reconstruct events end-to-end.
- Acceptance: ingest schema conformance for required fields is 100 %.
- Evidence: logging-schema validators, immutability settings, SIEM queries, and alert timelines (Evidence Pack ID).



Practitioner Guidance:

- Implement as code: express classification jobs, ABAC policies, DLP rules, logging schemas, and backup/WORM policies declaratively and with version control.
- Phase with proof: classify/tag first; then enforce ABAC (with MFA/JIT); then enable DLP block actions; then tune anomaly detection—capturing artifacts in an Evidence Pack ID.
- Profile encryption consistently: treat storage and transport protections as reusable profiles; validate with automated scans and KMS logs.
- Make logs useful: ensure events carry purpose and trace_id and include policy_decision; store in append-only/tamper-evident systems.
- Manage exceptions: any temporary waivers require scope, compensating control, owner, sunset date, and weekly review.

Obsolete and withdrawn documents should not be used; please use replacements.



Quick Win Playbook:

Title: Tag-Driven DLP + Deny-by-Default ABAC for a High-Value Dataset

Objective: Prove that sensitivity tags drive enforceable controls by
 (1) denying out-of-policy access via ABAC with a purpose,
 (2) blocking email/cloud-sync exfiltration via DLP, and
 (3) emitting investigation-ready events to immutable storage—using measurable SLOs and Evidence Packs EP-05.0/05.1/05.2 (and EP-05.3 if immutability is in scope).

Target: Enforce tag-driven DLP and deny-by-default ABAC for one high-value dataset to block email and cloud-sync exfiltration (§6.2, §6.4, §6.6).

Component/System: Data catalog + policy engine (ABAC); endpoint and SaaS DLP; email and SaaS egress control points; SIEM; append-only/WORM evidence store.

Protects: Sensitive data from out-of-policy egress across common channels (email, web upload, cloud sync) and from unscoped access.

Stops/Detects: Out-of-policy transfers of Restricted/Confidential classes; access without required attributes/purpose; missing tags on new data.

Action:

1. Classify and tag the target dataset; bind tags to ABAC (deny-by-default) and to DLP rules. Ensure the access request includes a purpose attribute that is evaluated by ABAC and recorded in the access event.
2. Enable DLP block/quarantine for email and cloud-sync channels; forward DLP telemetry to SIEM.
3. Validate Data Egress Control for the dataset: allowlisted export path succeeds; non-allowlisted destination is denied and logged.
4. Run a smoke test: (a) in-policy access (allow with policy_decision=allow), (b) out-of-policy access (deny with policy_decision=deny), (c) email/cloud-sync of tagged sample (block/quarantine).
5. Verify that all events contain trace_id and policy_decision and are stored in an append-only/WORM location.

Proof (attach artifacts to Evidence Packs):

- DLP simulation results and SIEM correlation views are filed under EP-05.2 (decision/correlation logs) and referenced from EP-05.0.
- EP-05.1 — Catalog/tag exports, class and owner mappings (classification and tagging).
- EP-05.2 — ABAC policy bundle, deny/allow decision logs, MFA/JIT elevation trails (if invoked).
- EP-05.3 — WORM immutability settings/logs for the dataset's evidence location (if immutable storage is used for the test artifacts or protected backups).
- EP-05.0 — Readiness summary referencing the items above.

Obsolete and withdrawn documents should not be used; please use replacements.

	Metric: Out-of-policy denials $\geq 99\%$; simulated exfiltration block $\geq 95\%$ with documented FP/FN bounds; 100 % of relevant events include the standard schema fields; artifacts present under EP-05.1 and EP-05.2 (with any immutability proof in EP-05.3) and summarized in EP-05.0. Rollback: Disable the DLP block rules and revert the ABAC policy bundle to the previous commit; record a time-bounded exception with owner and expiry in EP-05.0 (and link affected artifacts in EP-05.1/05.2).
--	--

Section 7. Cybersecurity Core Principles

The following ISAUnited Cybersecurity Core Principles guide the intent, design, and implementation of the Data Security Architecture Parent Standard. Each principle Must translate into concrete, testable behaviors—traceable to §6 outputs and verified in §12.

Table E-2. Applicable ISAUnited Recommended Principles for Data Security Architecture:

Principle name	Code	Applicability to Data Security Architecture
Least privilege	ISAU-RP-01	Enforce deny-by-default for sensitive classes; use ABAC decisions bound to tags, purpose, and context; require MFA/JIT for privileged data actions.
Zero Trust	ISAU-RP-02	No implicit trust from network or location; continuously evaluate identity, device posture, risk, and purpose for every data access.
Complete mediation	ISAU-RP-03	Every read/write/modify on sensitive data is evaluated and logged after classification/tag checks; no cached grants for sensitive operations without re-check.
Defense in depth	ISAU-RP-04	Layer controls driven by tags: ABAC to encryption by policy to DLP to tamper-evident logging to anomaly detection to immutable recovery; failure of one control does not compromise data.
Secure by design	ISAU-RP-05	Require classification and ownership before access; bind controls to tags via policies-as-code; design for residency constraints and retention from inception.

Obsolete and withdrawn documents should not be used; please use replacements.

Principle name	Code	Applicability to Data Security Architecture
Data minimization & purpose limitation	ISAU-RP-06-D	Collect and retain only necessary fields; enforce purpose-bound processing and masking/tokenization per class; prevent excessive exposure in API responses/exports.
Secure defaults	ISAU-RP-10	Encryption is enabled by policy; sensitive classes use deny-by-default access; logging is enabled with tamper-evident storage; WORM is required for critical backups. Explicit, approved action is needed to relax these protections.
Evidence production	ISAU-RP-15	Emit tamper-evident access/modify logs with timestamp, subject, source, object, action, result, purpose, trace_id, and policy_decision; maintain an Evidence Pack ID per validation cycle.
Protect confidentiality	ISAU-RP-18	Encrypt sensitive data at rest and in transit per CEK profiles; restrict access via ABAC/MFA/JIT; prevent out-of-policy movement with tag-driven DLP.
Protect integrity	ISAU-RP-19	Verify integrity of critical datasets and files (hashing, authenticated checks); protect backups/archives with WORM and verification jobs.
Protect availability	ISAU-RP-20	Design for recoverability with multi-region redundancy and tested restore to RTO/RPO; throttle abusive or anomalous access patterns to preserve service.

Implementation note (normative). Each selected principle Must map to at least one §6 output and one §12 verification/validation activity, with artifacts stored under the annex's Evidence Pack ID.

	Practitioner Guidance: - Map cybersecurity principle to output to test. For every entry in Table E-2, record the corresponding §6 output(s) and the Test-ID(s) in §12 that prove it, then add the artifact path to the Evidence Pack ID. - Make it measurable. Convert principle intent into an acceptance threshold (for example, RP-01: out-of-policy denial rate $\geq 99\%$ on sensitive classes; RP-20: quarterly restore meets RTO/RPO).
---	--

Obsolete and withdrawn documents should not be used; please use replacements.

	Re-validate on change. When classifications, retention rules, or processing purposes change, update policies and tests in the same change set and re-run the affected validations.
--	--

Section 8. Foundational Standards Alignment

Internationally recognized frameworks from NIST and ISO/IEC establish baseline expectations for engineering practice. Data Security Architecture builds upon these foundations, integrating them into a defensible, testable model for modern data estates.

Purpose and Function

- Demonstrate alignment with globally accepted NIST/ISO practices.
- Bridge baseline expectations to ISAUnited's engineering methodology.
- Provide a consistent baseline for clause-level mapping in sub-standards.
- Enhance credibility and traceability for adoption and audit readiness.

Table E-3. Applicable Foundational Standards:

Framework	Standard ID	Reference focus (data-security relevance)
NIST	SP 800-53 Rev. 5	Security and privacy controls for information systems (data classification, access, audit, protection, and recovery families).
NIST	SP 800-160 Vol. 1	Systems security engineering (engineered, testable protections across the data lifecycle).
NIST	SP 800-207	Zero Trust Architecture (principles for purpose-bound, attribute-driven data access).
NIST	SP 800-111	Guide to Storage Encryption (architectural considerations for data-at-rest encryption; crypto specifics deferred to CEK).
NIST	SP 800-209	Security Guidelines for Storage Infrastructure (integrity, isolation, and resilience for storage systems).
NIST	SP 800-88 Rev. 1	Media Sanitization (technical erasure/cryptographic erasure of data at end of life).

Obsolete and withdrawn documents should not be used; please use replacements.

Framework	Standard ID	Reference focus (data-security relevance)
ISO/IEC	27001:2022	ISMS requirements (control families that underpin data security governance and measurement).
ISO/IEC	27002:2022	Code of practice for information security controls (data classification, access, logging, and protection controls).
ISO/IEC	27040	Storage security (architectures, controls, and operations for storage platforms).
ISO/IEC	27017	Code of practice for information security controls for cloud services (provider/consumer data-security responsibilities).
ISO/IEC	27018	Protection of PII in public cloud acting as PII processors (data-handling protections in cloud data planes).

NOTE: *ISAUnited Charter Adoption of Foundational Standards.*

Per the ISAUnited Charter, the institute formally adopts the International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) and the National Institute of Standards and Technology (NIST) as its foundational standards bodies, consistent with their public encouragement of organizational adoption. Parent Standards align with ISO/IEC and NIST for architectural grounding and auditability, and this alignment flows down to Sub-Standards as invariants and minimum requirements that may be tightened but not weakened. ISAUnited does not restate or speak on behalf of ISO/IEC or NIST; practitioners shall consult the official publications and terminology of these organizations, verify scope and version currency against the latest materials, and implement controls in a manner consistent with ISAUnited security invariants and the requirements of this standard.

As sub-standards are developed and published under this parent standard, more specific references to NIST and ISO foundational standards will be included to provide detailed, control-level alignment and facilitate practical implementation.

Obsolete and withdrawn documents should not be used; please use replacements.

Sub-Standard Expectations

Sub-standards under ISAU-DS-DS-1000 Must:

- Cite specific clauses from Table E-3 (for example, NIST SP 800-53 AC-6, AU-, CP-; ISO/IEC 27002:2022 5/8/10/12/18; ISO/IEC 27040 sections) for each normative output they extend.
- Convert those clauses into testable engineering behaviors (policies-as-code / data-controls-as-code) with verification/validation in §12.
- Document any divergence with compensating controls, an acceptance rationale, and a sunset date; store passing artifacts under the Evidence Pack ID.
- Include a concise mapping table: §6 Output to Framework to Clause to Test-ID(s) to Evidence Pack ID.

	Practitioner Guidance: • Map at the clause level only. For each §6 output (e.g., 6.1 Classification & Governance, 6.2 ABAC/MFA/JIT, 6.5 Integrity & Resilience), add a row: <i>Spec to NIST/ISO clause, how enforced (policy/code), to Evidence Pack ID.</i> • Keep mappings current. When a control or policy changes, update the citation in the same change and store the diff in the Evidence Pack. • Scope discipline. Reserve CSA CCM, CIS Controls, and OWASP for Section 9 (Security Controls); do not list them as foundational standards here.
---	--

Section 9. Security Controls

This section specifies the technical control families and control references enforced by the Data Security Architecture Parent Standard. These mappings ensure traceability between data-layer requirements and recognized industry frameworks—providing explicit, actionable guidance for engineers, reviewers, and auditors.

Purpose and Function

Security controls bridge architectural objectives and actionable safeguards at the data layer—protecting confidentiality, integrity, availability, authentication, authorization, and auditability across datasets and data paths. By mapping to CSA CCM, CIS Controls v8, and OWASP ASVS/API Top 10, ISAUnited ensures:

- Alignment with widely adopted best practices,
- Interoperability across platforms and processing styles,
- Audit-ready traceability into sub-standards and project implementations.

Obsolete and withdrawn documents should not be used; please use replacements.

Implementation Guidance

Authors and practitioners Must:

- Reference controls from CSA CCM, CIS Controls v8, and OWASP (ASVS/API Top 10) that are directly enforced in the data layer.
- Provide framework acronym, control family/ID, and a concise, implementation-oriented description.
- Map each control to one or more §6 outputs and to named tests in §12 (Verification & Validation).
- Favor enforceable controls (tag-driven ABAC rules, DLP policies, retention/WORM configs, logging schemas) over policy-only statements.

Table E-4. Control Mappings for Data Security Architecture:

Framework	Control ID	Control name/description (data-layer)	Primary linkage to §6 outputs
CIS v8	3.4	Configure entitlements (ACLs/ABAC) based on business need-to-know for sensitive classes.	6.2 Secure Data Access Controls
CIS v8	3.11	Encrypt sensitive data in transit (architectural reference to CEK profiles).	6.3 Encryption & Data Protection
CIS v8	3.12	Encrypt sensitive data at rest (architectural reference to CEK profiles).	6.3 Encryption & Data Protection
CIS v8	8.x	Audit log management: generate, aggregate, retain, and protect security-relevant logs.	6.6 Logging, Auditing & Anomaly Detection
CIS v8	11.3	Protect recovery data (isolation/immutability, access controls, integrity checks).	6.5 Data Integrity & Resilience
CIS v8	13.7	Deploy DLP to monitor and protect sensitive data in transit across channels.	6.4 Data Loss Prevention & Monitoring
CSA CCM	DSI-01	Maintain the current data inventory/flows, including classification and sensitivity levels.	6.1 Data Classification & Governance

Obsolete and withdrawn documents should not be used; please use replacements.

Framework	Control ID	Control name/description (data-layer)	Primary linkage to §6 outputs
CSA CCM	DSI-03	Enforce retention, archival, and secure disposal bound to classification and hold flags.	6.1 Data Classification & Governance; 6.5 Data Integrity & Resilience
CSA CCM	EKM-02	Encryption key management lifecycle (architectural reference; CEK defines specifics).	6.3 Encryption & Data Protection
CSA CCM	IAM-05	Least-privilege access to data assets; restrict to need-to-know with purpose context.	6.2 Secure Data Access Controls
OWASP ASVS v4.0	V2.1	Authentication/authorization architecture for access to sensitive data systems.	6.2 Secure Data Access Controls
OWASP ASVS v4.0	V9.1	Data protection verification for sensitive data at rest and in transit (ref. CEK).	6.3 Encryption & Data Protection
OWASP API Top 10 (2023)	API3	Prevent excessive data exposure by filtering and minimizing API responses/exports.	6.1 Data Classification & Governance; 6.2 Secure Data Access Controls; 6.4 DLP & Monitoring

NOTE. Use exact clause/requirement references in sub-standards (for example, “CIS 8.x” to cite the specific sub-control) when mapping named tests in §12. Cryptographic parameters and key lifecycles are governed in Annex I (CEK).

NOTE: *Use of External Control Frameworks.*

Per the ISAUnited Charter, the institute formally adopts and maps to external control frameworks to provide alignment and traceability, but does not speak on behalf of those organizations. Practitioners shall consult and follow the official practices, recommendations, and implementation guidance of the Center for Internet Security (CIS), the Cloud Security Alliance (CSA), and the Open Worldwide Application Security Project (OWASP) when applying controls. Always verify control identifiers, scope, and version currency against the publishers’ latest materials. Where wording differs, use the framework’s official documentation while maintaining consistency with ISAUnited security invariants and this standard's requirements.

Obsolete and withdrawn documents should not be used; please use replacements.

Additional References

As data-layer threats and practices evolve, sub-standards may incorporate additional OWASP controls (beyond those listed) in ASVS/API sections that are directly enforced at the data layer. Foundational NIST/ISO references remain limited to §8.

Sub-Standard Expectations

Sub-standards under this Parent Standard Must:

- Select and enforce explicit data-layer controls relevant to their scope (for example, catalog/tags, ABAC, DLP, retention/WORM, logging schema).
- Provide detailed mappings from each control to §6 outputs, §12 tests, and an Evidence Pack ID.
- Document any deviation from control families with compensating controls and a sunset date; include passing verification artifacts.

Evidence Pack ID.

Document any deviation from control families with compensating controls and a sunset date; include passing verification artifacts.

	Practitioner Guidance: • Build and maintain a Controls to Outputs to Tests sheet per data domain: each control from Table E-4 points to §6 items, Test-IDs in §12, and the Evidence Pack ID. • Keep the sheet current in the same change that modifies ABAC/DLP/retention/logging; attach proofs (policy diffs, posture scans, test results). • Favor controls that can be expressed as code or declared in policy engines, catalogs, and DLP registries, and verified automatically by tests.
---	--

Section 10. Engineering Discipline

This section defines the architectural thinking, rigorous engineering processes, and disciplined operational behaviors required to implement Data Security Architecture (ISAU-DS-DS-1000). ISAUnited's Defensible Standards are not compliance checklists; they are engineered systems, grounded in systems thinking, critical reasoning, and

Obsolete and withdrawn documents should not be used; please use replacements.

Verification & Validation (V&V), that produce measurable, auditable, defensible outcomes across data stores, data paths, and data-driven services.

10.1 Purpose & Function

Purpose. Establish a repeatable, auditable way of working that integrates systems thinking, lifecycle controls, adversary-aware design, and measurable outcomes for data security.

Function in D10S. Parent Standards set expectations and invariants. Sub-Standards convert them into policies-as-code/data-controls-as-code, test specifications, and evidence artifacts embedded in delivery and operations.

10.2 Systems Thinking

Goal: Make the data system end-to-end legible—repositories, flows, interfaces, and dependencies—so controls bind to where risk manifests.

10.2.1 System Definition & Boundaries

- Declare system purpose, scope, stakeholders, and in-/out-of-scope assets (databases, warehouses/lakehouses, object/file stores, SaaS data planes, analytics/AI pipelines, messaging/streaming, edge, backup/archive).
- Model trust zones and boundary crossings (user/service to data gateway, service to data store, pipeline to store, inter-tenant/share, export/egress).
- Write Once, Read Many (WORM) is platform-enforced immutability for stored objects. After write, an object cannot be modified or deleted until its retention period expires or a legal hold is released. WORM is a technical control, not a policy alone.

10.2.2 Interfaces & Data Contracts

- Maintain Interface Control Documents (ICDs) for data access patterns (SQL/NoSQL, object APIs, streaming topics, export jobs, SDKs/ODBC/JDBC, service accounts).
- For each interface, specify: identity type (human/service), ABAC model and purpose context, sensitivity class and Sensitivity Tag, rate/flow limits, Data Egress Control rules, error/deny semantics, telemetry fields (trace_id, policy_decision), retention/residency flags, and invariants (for example, “deny-by-default for Restricted,” “no plaintext export for Protected classes,” “schema change requires re-classification”).

10.2.3 Dependencies & Emergent Behavior

- Map shared services (catalog/lineage, KMS/HSM, SIEM, secrets, scheduler, backup/WORM, data gateway).

Obsolete and withdrawn documents should not be used; please use replacements.

- Identify emergent risks from composition (for example, permissive sharing links + shadow sandboxes to uncontrolled copies; broad “reader” role + missing purpose tag to data spill; export job + no egress allowlist to exfiltration; restore without encryption proof to integrity failure).

10.2.4 Failure Modes & Safeguards

- For critical paths, document failure modes (classification drift, orphaned datasets, ABAC bypass, untagged exports, mutable backups, missing audit fields) and safeguards (automated discovery with SLO, deny-by-default, tag-driven DLP, append-only logging, WORM, encrypted restore drills).

Required Artifacts (minimum): Context diagram with trust boundaries; system data-flow map and lineage overlays; ICD set; invariants register.

10.3 Critical Thinking

Goal: Replace assumptions with explicit reasoning that survives review, attack, and audit.

10.3.1 Decision Discipline

- Use Architecture Decision Records (ADRs): problem to options to constraints/assumptions to trade-offs to decision to invariants to test/evidence plan (who/when/how measured).

10.3.2 Engineering Prompts

- **Boundaries:** What is the data system? Where are the trust boundaries and why?
- **Interfaces:** What must always be true at each data interface (invariants)? How is it tested?
- **Adversary:** Which techniques are credible here (insider misuse, excessive data exposure, out-of-policy egress, ransomware on backups)? What is the shortest attack path?
- **Evidence:** Which objective signals prove this control works today and after the change?
- **Failure:** When this fails, does it fail safe (deny, quarantine, immutable log)? What is the operator’s following action?

Required Artifacts (minimum): ADRs; assumptions and constraints log; evidence plan per decision.

10.4 Domain-Wide Engineering Expectations

Secure System Design

- Define data boundaries (stores, pipelines, gateways, catalogs, KMS/HSM, SIEM, backup/WORM).
- Validate boundaries and trust relationships with §10.2 artifacts; ensure protections bind to classification and purpose at each hop.

Obsolete and withdrawn documents should not be used; please use replacements.

Implementation Philosophy — “Built-in, not bolted-on.”

- Integrate classification, ABAC, encryption by policy, DLP, logging, integrity, and recovery at design time; avoid retrofits that create Shadow Data.
- Express controls as policies-as-code/data-controls-as-code bound to invariants in §10.2.4 (for example, deny-by-default for Restricted, tag-driven DLP, no plaintext export, WORM retention).

Lifecycle Integration

- Embed data controls into design reviews, backlog, build/test, deployment, and operations; keep delivery mechanics in Annex J and crypto specifics in CEK.
- Enforce version-controlled reviews with required ADRs and Evidence Pack ID updates on every change.

Verification Rigor (V&V)

- Combine automated checks (discovery coverage/latency, ABAC decision tests, DLP simulations, event-schema conformance, immutability checks, encrypted-restore drills) with targeted manual probes (purpose-misuse, unusual joins/exports, residency edge cases).
- Require continuous validation in pipelines and scheduled runtime checks tied to invariants (for example, classification SLO, deny out-of-policy, block egress, WORM deny-alter).

Operational Discipline

- Monitor for drift and unauthorized change (classification gaps, permissive shares, orphaned datasets, stale ABAC/DLP/retention, residency violations); auto-remediate where safe with time-bounded exceptions.
- Maintain runbooks/SOPs for access-policy violations, suspected exfiltration, integrity anomalies, and recovery events; record outcomes in the Evidence Pack.

10.5 Engineering Implementation Expectations

- **Policies/Controls as Code:** Manage catalog schemas, tag bindings, ABAC rules, DLP registries, retention/WORM profiles, and logging schemas as code with peer review and provenance.
 - **Structured Enforcement Path:** Build to discovery/tag jobs to ABAC tests to DLP simulations to encrypted-restore drill to canary to promote/rollback (execution in Annex J; semantics here).
 - **Explicit Security Boundaries:** Maintain diagrams and ICDs; continuously validate posture (deny-by-default, no plaintext export, tag to control bindings, WORM) with targeted audits and smoke tests.
 - **Automated Security Testing:** Integrate classification coverage/latency checks, ABAC decision suites, DLP exfil simulations, event-schema validators (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision), and immutability proofs before production.
 - **Traceable Architecture Decisions:** Link ADRs to controls, tests, and evidence; update ADRs and evidence on each change request.
- Required Artifacts (minimum):** Policies-as-code repository;

Obsolete and withdrawn documents should not be used; please use replacements.

enforcement/test gates; boundary/ICD set; automated test results; evidence ledger (see §10.7 and §12).

10.6 Sub-Standard Alignment (inheritance rules)

Sub-Standards must operationalize this discipline with data-specific detail:

- Catalog, Tags, Owners, Lineage, Retention (for example, ISAU-DS-DS-5010). Automated discovery targets, tag schema, and owners; retention/hold bindings. Tests: coverage, latency, retention enforcement.
- Storage/Transport Encryption & KMS Integration (for example, ISAU-DS-DS-5020). CEK-aligned profiles; posture scans; KMS rotation/use logs; encrypted-restore proof.
- DLP Policy Engineering & Enforcement (for example, ISAU-DS-DS-5030). Tag-driven DLP across endpoints, networks cloud/SaaS; FP/FN bounds; SIEM correlation.
- Immutable Backup & Recovery Architecture (for example, ISAU-DS-DS-5040). WORM retention; deny-alter logs; multi-region drills to RTO/RPO.
- Zero-Trust Data Access (ABAC/MFA/JIT) (for example, ISAU-DS-DS-5050). Deny-by-default; purpose-bound decisions; privileged elevation with MFA/JIT.
- Logging, Audit & Anomaly Analytics (for example, ISAU-DS-DS-5060). Standard event schema; tamper-evident storage; MTTD targets; end-to-end reconstruction.

10.7 Evidence & V&V (what proves it works)

Establish a Data Evidence Pack per system containing:

- **Design Evidence:** trust-boundary diagrams, data-flow/lineage maps with ICDs, invariants register, ADRs.
- **Build Evidence:** policies-as-code (catalog, tags, ABAC, DLP, retention/WORM, logging schema), discovery/latency reports, ABAC decision tests, DLP simulations with FP/FN, encrypted-restore drill results, KMS/HSM logs.
- **Operate Evidence:** runtime deny/allow logs with trace_id/policy_decision, SIEM correlation dashboards/queries, immutability proofs (WORM deny-alter), residency/retention mappings, incident and rollback records.
- **Challenge Evidence:** red-team/exfil tests, adversary emulation on data paths, remediation closure with re-test.

Each control requires objective pass/fail criteria, a specified test frequency, a designated responsible owner, and a defined retention policy. Map Evidence Pack IDs into §12 traceability. Use EP-05.1 for catalog/tags/lineage/retention items; EP-05.2 for ABAC/MFA/JIT; EP-05.3 for encryption/KMS and encrypted-restore; and EP-05.0 for readiness summaries, referencing the above.

Obsolete and withdrawn documents should not be used; please use replacements.

10.8 Example: Sub-Standard Discipline Alignment (Catalog, Tags, and ABAC)

Scope: ISAU-DS-DS-5010 / 5050 (Classification & Governance; ZTDA).

Design: Define trust boundaries and invariants (for example, “classification precedes access,” “deny-by-default for Restricted,” “purpose required on sensitive reads”)—document decision points and enforcement at data gateways and stores.

Implement: Express tag bindings, ABAC rules, DLP policies, and logging schema as code; require MFA/JIT for privileged data actions; block plaintext export paths for protected classes.

V&V: Run discovery coverage/latency checks; ABAC deny/allow suites with purpose; DLP exfil simulations; verify event schema (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision) and WORM immutability.

Operate: The Evidence Pack includes policy repository history, coverage/latency reports, ABAC decision logs, DLP incidents/tests, SIEM correlations, encrypted-restore results, and closed-loop remediation (reference EP-05.1/05.2/05.3, summarized in EP-05.0).



Practitioner Guidance:

- Maintain the Controls to Outputs to Tests to SLO to Evidence Pack sheet per data domain; update it in the same change that modifies catalog/tags, ABAC, DLP, retention, or logging; attach proofs (policy diffs, posture scans, test results), and reference EP-05.1/05.2/05.3 (summary in EP-05.0).
- Prefer controls as code with automatic §12 validation; reserve exceptions for time-bounded, owner-approved waivers with compensating controls, and record them in the sheet with a sunset date and Evidence Pack link.

Section 11. Associate Sub-Standards Mapping

Purpose of Sub-Standards

ISAUnited Defensible Sub-Standards under Data Security Architecture are tightly scoped, engineering-driven extensions that:

- Define granular, data-layer requirements (DSR-IDs) for specialized domains.
- Translate architectural intent into enforceable behaviors in platforms and policies (catalog/tags, ABAC, DLP, retention/WORM).
- Specify verification/validation methods that yield test artifacts (discovery coverage, deny/allow logs, DLP simulations, encrypted restore drills) referenced in §12.
- Align directly to the Parent Standard’s §6 outputs and §7 principles, with traceable Evidence Pack artifacts.

Obsolete and withdrawn documents should not be used; please use replacements.

Interface notes (non-normative):

- Annex E produces data-layer requirements (DSR-IDs), control bindings to tags, and tests.
- Annex J ensures those tests run in CI/CD and at promotion; provenance/SBOM and gates live there.
- Annex I (CEK) governs cryptographic profiles and key lifecycles; Annex E governs correct data-layer application (encrypted posture, encrypted restore).
- Annex F (IAM) and Annex H (MDR) provide identity foundations and detection/IR workflows that consume data-layer events.

Scope and Focus of Data Security Sub-Standards**Data Classification & Governance Implementation**

Example Sub-Standard: ISAU-DS-DS-1010 – Catalog, Tags, Owners, Lineage, Retention

- Automated discovery scope/coverage targets; sensitivity tag schema; owners/stewards recorded.
- Retention schedules and legal holds bound to tags; lineage updates on ingest/transform/export.
- Maps to §6: 6.1
- Tests: discovery coverage, tagging latency, retention enforcement.

Zero-Trust Data Access (ABAC/MFA/JIT)

Example Sub-Standard: ISAU-DS-DS-1050 – Purpose-Bound ABAC & Privileged Elevation

- Deny-by-default for sensitive classes; ABAC rules include purpose; MFA/JIT for privileged data actions.
- Maps to §6: 6.2
- Tests: out-of-policy denial rate, elevation approval/auto-revocation logs.

Encryption & Key Governance for Data (architectural)

Example Sub-Standard: ISAU-DS-DS-1020 – Storage/Transport Profiles & KMS Integration

- Encryption posture for at-rest/transport per CEK profiles; KMS policies for key use/rotation; encrypted restore proof; no plaintext export path.
- Maps to §6: 6.3
- Tests: transport scans, storage posture, KMS rotation/usage logs, restore drills.

DLP Policy Engineering & Enforcement

Example Sub-Standard: ISAU-DS-DS-1030 – Tag-Driven DLP Across Channels

- Endpoint/network/cloud/SaaS DLP policies keyed to tags and data types; FP/FN tolerances defined.
- Maps to §6: 6.4, 6.6
- Tests: exfil simulation results (email/web/cloud sync), SIEM correlation.

Obsolete and withdrawn documents should not be used; please use replacements.

Immutable Backup & Recovery Architecture

Example Sub-Standard: ISAU-DS-DS-1040 – WORM, Multi-Region, RTO/RPO Drills

- WORM retention enforcement; denied-alter/delete events; periodic restore exercises meeting targets.
- Maps to §6: 6.5
- Tests: denied alterations, drill attainment, and integrity verification jobs.

Logging, Audit, & Anomaly Analytics

Example Sub-Standard: ISAU-DS-DS-1060 – Standardized Access Events & MTDD SLOs

- Event schema: timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision; tamper-evident storage and SIEM onboarding.
- Maps to §6: 6.6
- Tests: schema conformance, immutability proof, anomaly detection latency (MTDD).

Table E-5. Example Future Sub-Standards

Sub-Standard ID	Sub-Standard Name	Focus Area
ISAU-DS-DS-1010	Catalog, Tags, Owners, Lineage, Retention	Classification & governance.
ISAU-DS-DS-1020	Storage/Transport Encryption & KMS Integration	Encryption posture (CEK-aligned).
ISAU-DS-DS-1030	DLP Policy Engineering & Enforcement	DLP across channels.
ISAU-DS-DS-1040	Immutable Backup & Recovery Architecture	WORM & recovery drills.
ISAU-DS-DS-1050	Zero-Trust Data Access (ABAC/MFA/JIT)	Purpose-bound access control.
ISAU-DS-DS-1060	Logging, Audit & Anomaly Analytics	Telemetry and MTDD SLOs.

Obsolete and withdrawn documents should not be used; please use replacements.

Development and Approval Process

ISAUnited uses an open, peer-driven annual process to propose, review, and publish sub-standards:

- Open Season Submission – Proposals Must cite Annex E §6 outputs and §7 principles they extend, plus NIST/ISO clauses from §8.
- Technical Peer Review – Evaluate engineering rigor, testability, scope clarity, and cross-domain consistency.
- Approval & Publication – Assign identifier, version, and publish as an actionable extension.

Sub-Standard Deliverables (normative)

Each sub-standard Must include:

- Inputs (Requirements): Preconditions (from Annex E §5) depend on.
- Outputs (Specifications): Concrete data-layer behaviors and thresholds (SLOs) tied to §6.
- Verification/Validation: Named tests and acceptance criteria tied to §12 (e.g., coverage, denial rates, DLP FP/FN, drill attainment).
- Evidence: Artifact list and storage location (Evidence Pack ID).
- Standards Mapping: DSR-ID/Spec to NIST/ISO clause (from §8) to Controls (from §9) to Test-ID to Evidence Pack ID.
- Interfaces: Explicit delineation of what is enforced in data platforms/policies (Annex E) vs. delivery (Annex J) and crypto parameters (Annex I).


	Practitioner Guidance: • Bind tags before tests. Define Sensitivity Tag to ABAC/DLP/retention/logging bindings first; if any tag lacks a bound control and named Test-ID (§12), halt and open a tracked risk. • Make SLOs explicit and provable. Pick 1–2 SLOs per sub-standard (e.g., discovery coverage $\geq 98\%$, tagging latency ≤ 24 hours; out-of-policy denials $\geq 99\%$; DLP block $\geq 95\%$; encrypted restore meets RTO/RPO) and point to the Evidence Pack ID that proves each. • Keep CEK separation and traceability. Say “per CEK cryptographic profiles” for encryption; verify via KMS/HSM logs and posture scans. In your mapping sheet, always include: §5 input(s) to §6 output(s) to NIST/ISO clause (§8) to control (§9) to Test-ID (§12) to Evidence Pack ID.
---	---

Obsolete and withdrawn documents should not be used; please use replacements.

Section 12. Verification and Validation (Tests)


The effectiveness and defensibility of data security architecture must be continuously verified and validated using structured, engineering-grade assessments. While detailed test requirements for specific platforms will live in Data sub-standards, this Parent establishes the gold-standard expectations below.

Figure 12-1 Verification confirms that the system has been implemented in accordance with the Requirements (Inputs) and Technical Specifications (Outputs) defined in this standard:



Obsolete and withdrawn documents should not be used; please use replacements.

Figure 12-2 Validation ensures that the system performs effectively in real-world operational conditions and withstands adversarial testing:



Core Verification Activities

- Confirm §6 controls exist and are enforced at data boundaries and platforms: automated discovery/tagging with SLOs; ABAC deny-by-default with purpose context; encryption by policy per CEK cryptographic profiles with KMS/HSM governance; tag-driven DLP across channels; standardized access/modify events (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision) with tamper-evident storage.
- Review platform and policy baselines: catalog/tag schemas and owners; ABAC rules and elevation flows (MFA/JIT); retention/WORM profiles; data egress allowlists; logging schema validators; confirm no plaintext export paths for protected classes.
- Verify integrations do not break data paths: gateway/policy engine ↔ stores; catalog/lineage ↔ pipelines; KMS/HSM ↔ platforms; DLP ↔ SIEM—and confirm enforcement points align with business-critical flows.

Obsolete and withdrawn documents should not be used; please use replacements.

Core Validation Activities

- Adversary-informed exercises: simulate out-of-policy access (deny-by-default), purpose misuse, excessive data exposure, and exfiltration via email/web/cloud sync; require DLP block/quarantine $\geq 95\%$ with FP/FN bounds.
- Runtime resilience: prove WORM denies alteration during retention; run encrypted restore drills to RTO/RPO; validate residency constraints; confirm anomaly detection MTTR ≤ 15 minutes on sensitive classes.
- Operational drills: classification drift detection (coverage $\geq 98\%$, tagging latency ≤ 24 hours), elevation approval/auto-revocation trails, SIEM correlation from data-access + DLP events, and end-to-end incident reconstruction using standard event fields.

Required Deliverables

1. Test Plans and Procedures — Scope, tooling, and methods for verification and validation phases, including Test IDs and owners.
2. Validation Reports — Pass/fail results, residual risk, and prioritized remediation actions tied to §6 outputs and DSR-IDs.
3. Evidence Artifacts — Discovery/coverage reports; tag exports/owners; ABAC policy/decision logs; MFA/JIT trails; transport/storage posture scans; KMS/HSM rotation/use logs; DLP simulations/incidents with FP/FN; WORM configs and deny-alter events; restore drill results; standardized event samples and immutability settings—each labeled with an Evidence Pack ID (EP-05.1/05.2/05.3, summarized in EP-05.0).
4. Corrective Action Plans — Time-bound remediation for findings that must be closed prior to acceptance.

Common Pitfalls to Avoid

- Treating classification as a one-time import (coverage/latency drift breaks every downstream control).
- Policy-only ABAC (no deny-by-default or purpose binding), or encryption without CEK-aligned key governance and encrypted-restore proof.
- DLP not tied to tags; logs missing trace_id/policy_decision; lack of immutability—investigations become guesswork.

Obsolete and withdrawn documents should not be used; please use replacements.

Table E-5. Traceability Matrix: Requirements (§5) to Verification/Validation (§12) to Technical Specifications (§6)

Requirement ID	Requirement (summary)	Verification (build-correct)	Validation (works-right)	Related §6 Outputs
5.1	Enterprise data catalog & classification framework.	Catalog, tag schema, owners/stewards present; lineage pointers set.	Coverage ≥ 98 %; tagging latency ≤ 24 hours on sampled changes.	6.1
5.2	Sensitivity tags bound to controls.	Documented tag to control bindings (ABAC, DLP, retention/erasure, masking, logging).	Bound controls trigger as expected (deny/allow, block/quarantine, enforce retention/masking).	6.1, 6.2, 6.4, 6.6
5.3	Zero-Trust Data Access baseline.	ABAC policies with deny-by-default; MFA/JIT configured.	Out-of-policy denials ≥ 99 %; privileged actions show approval and auto-revocation logs.	6.2
5.4	Encryption policy & KMS integration.	Encryption by policy at rest and in transit; KMS/HSM integrated (CEK-aligned).	Encrypted restore meets RTO/RPO; key rotation/usage logs meet policy.	6.3
5.5	DLP channel coverage & policy registry.	Endpoint/network/SaaS DLP deployed; SIEM correlation configured.	Simulated exfiltration blocked ≥ 95 % with FP/FN bounds.	6.4, 6.6
5.6	Immutable backup & recovery objectives.	WORM retention enforced; multi-region copies configured.	Alter/delete attempts denied; quarterly restores meet RTO/RPO; integrity checks pass.	6.5
5.7	Residency & processing constraints.	Residency/sovereignty constraints mapped to classes and systems.	The audit sample shows the required controls are active for regulated/region-bound datasets.	6.1, 6.2, 6.3, 6.6
5.8	Discovery & lineage coverage targets.	Discovery/lineage tooling with KPIs and owners.	KPI attainment report: lineage matches actual flows for sampled datasets.	6.1

Obsolete and withdrawn documents should not be used; please use replacements.

Requirement ID	Requirement (summary)	Verification (build-correct)	Validation (works-right)	Related §6 Outputs
5.9	Centralized logging & schema.	Standard event schema; tamper-evident storage.	MTTD $\leq$ 15 minutes on sensitive classes; end-to-end reconstruction succeeds.	6.6
5.10	Evidence & metrics readiness.	Dashboards/queries for coverage, denials, DLP, KMS, and drills.	Baseline metrics captured; post-change deltas recorded under Evidence Pack ID (EP-05.0).	6.1–6.6

Evidence guidance

Attach (per row) to the Evidence Pack ID: catalog/tag exports; coverage/latency reports; ABAC policies + deny/approval logs; MFA/JIT events; transport/storage scans; KMS rotation/use logs; DLP policies/incidents/tests with FP/FN; SIEM queries/dashboards; WORM configs and deny-alter logs; restore drill reports; logging-schema validators and anomaly timelines; residency/retention mappings. Use EP-05.1/05.2/05.3 as applicable; summarize in EP-05.0.

How to use this matrix

- **Plan:** For each §5 requirement, define ≥ 1 verification and ≥ 1 validation tied to §6 outputs.
- **Execute:** Run tests; record SLO met/not met with direct artifact links in the Evidence Pack.
- **Maintain:** When a requirement or enforcement changes, update the row and re-run impacted tests.

	Practitioner Guidance: • Test what the tags drive. Start with tags to control bindings (ABAC, DLP, retention, logging) and prove them with named Test-IDs; if any tag lacks a bound control and test, stop and fix. • Prove resilience, not just posture. Posture scans without encrypted-restore drills, WORM deny-alter evidence, and MTTD measurements are incomplete.
---	--

Obsolete and withdrawn documents should not be used; please use replacements.

	Keep CEK separation and traceability. Say “per CEK cryptographic profiles,” prove via KMS/HSM logs and transport/storage scans, and link every result to an Evidence Pack ID.
--	---

	Quick Win Playbook: Title: Classification + ABAC + DLP V&V Smoke Suite Objective: Prove that tags drive enforcement by validating deny-by-default on sensitive classes and blocking exfiltration on a single High-Value Dataset (HVD), with investigation-ready telemetry. Target: Stand up a “classification + ABAC + DLP” V&V smoke suite with fail-closed gates for one HVD (§6.1, §6.2, §6.4, §12). Component/System: Data catalog/classification jobs; policy engine (ABAC); endpoint/email/cloud-sync DLP; SIEM; append-only/WORM evidence store (if used). Protects: Sensitive data from out-of-policy access and exfiltration across common channels. Stops/Detects: Access without required attributes/purpose; exfiltration of tagged samples; missing tags on new data. Action: - Classify/tag the HVD; bind tags to ABAC (deny-by-default) and to DLP rules. - Run smoke: (1) in-policy access (allow; purpose attribute present), (2) out-of-policy access (deny), (3) email/cloud-sync of tagged sample (block/quarantine). - Confirm events carry trace_id/policy_decision and are stored append-only/WORM (if used); capture SIEM correlations. Proof (Evidence Packs): EP-05.1 (classification/tag artifacts), EP-05.2 (ABAC decision logs and tests), EP-05.3 (immutability settings if used for artifacts), summarized in EP-05.0. Metric: Out-of-policy denials ≥ 99%; exfil block ≥ 95%, with FP/FN bounds; 100% of relevant events include required schema fields. Rollback: Temporarily relax the DLP block rule and revert the ABAC policy bundle; record a time-bounded exception with owner/expiry in EP-05.0 and link affected artifacts in EP-05.1/05.2.
---	--

Obsolete and withdrawn documents should not be used; please use replacements.

Section 13. Implementation Guidelines

This section does not prescribe vendor-specific tactics. Parent Standards are stable, long-lived architectural foundations. Here, we define how sub-standards and delivery teams must translate the Parent's intent into operational behaviors that are testable, automatable, and auditable for Data Security Architecture (Annex E). Delivery mechanics (pipeline orchestration, provenance/SBOM, promotion/rollback) are governed by Annex J.

Purpose of This Section in Sub-Standards

Sub-standards must use Implementation Guidelines to:

- Translate Parent expectations into enforceable data-layer behaviors (for example, automated discovery/tagging, ABAC deny-by-default with purpose, tag-driven DLP, WORM retention, standardized events with trace_id/policy_decision).
- Provide stack-agnostic practices that improve adoption, reduce failure, and align with ISAUnited's defensible design philosophy.
- Highlight common failure modes and how to prevent them with measurable gates and checks.
- Offer repeatable patterns (as code) that enforce controls, trust models, and engineering discipline across catalogs/lineage, policy engines, data stores/warehouses/object stores, pipelines/streaming, gateways, KMS/HSM, and SIEM.

Open Season Guidance for Contributors

Contributors developing sub-standards Must:

- Align all guidance with this Parent's strategic posture and §6 outputs.
- Avoid vendor/product terms; express controls as requirements, tests, and evidence with an Evidence Pack ID.
- Include lessons learned (what fails, why, and how the test proves it).
- Focus on repeatable engineering patterns (policies-as-code / data-controls-as-code), not one-offs.
- Provide a minimal Standards Mapping (Spec/Control to NIST/ISO clause from §8 to Evidence Pack ID).

Obsolete and withdrawn documents should not be used; please use replacements.

Technical Guidance

A. Organizing Principles (normative)

1. Everything as code. Catalog/tag schemas, tag to control bindings, ABAC rules, DLP policies, retention/WORM profiles, logging schemas, egress allowlists, and runbooks Must be version-controlled, peer-reviewed, and promoted on protected branches.
2. Gated change. Every merge/release Must pass non-bypassable security gates tied to §6 and §12 (for example, discovery coverage $\geq 98\%$ and tagging latency ≤ 24 hours; deny-by-default tests on sensitive classes; DLP exfil simulations pass with FP/FN bounds; event-schema conformance = 100%; encrypted-restore drill meets RTO/RPO).
3. Immutable, reproducible releases. No manual policy changes post-build; releases are reproducible and verified at data boundaries (gateways/policy engines) and platforms.
4. Least privilege & JIT (data context). Service identities and automation accounts are scoped; MFA/JIT is required for privileged data actions; logs preserve confidentiality while remaining diagnostically useful.
5. Environmental parity. Staging mirrors production controls (catalog/tags, ABAC, encryption policy, DLP, logging schema, WORM posture) so tests are predictive; drift is monitored and reconciled.

B. Guardrails by Pipeline Stage (normative)

1. **Pre-commit / local**
 - Secrets scanning and signed commits required.
 - Hooks should validate catalog/tag schema, generate policy stubs, and lint ABAC/DLP definitions; block unsafe patterns (for example, plaintext export of protected classes).
2. **Pull request (PR) / code review**
 - CODEOWNERS approval required; record a “Data-Model Delta” for boundary/schema/tag changes.
 - Discovery/coverage gate for changed data domains; critical findings = 0.
 - ABAC deny-by-default tests for changed sensitive classes; DLP rule diffs; evidence pointers in PR (planned §12 Test-IDs and Evidence Pack ID stub).
3. **Build & package**
 - Deterministic artifacts (pinned policy bundles; no ad-hoc fetch at deploy); integrity checks for policies-as-code.
 - Package tag bindings, ABAC/DLP rules, retention/WORM profiles, and logging schema as deployable config.
4. **Pre-deploy / release**
 - Config-drift detection against approved policies; approvals “as code.”
 - Progressive rollout (staged/canary) for ABAC/DLP and logging schema with health SLOs and automatic rollback.
 - Positive/negative tests: tag to control bindings, out-of-policy access (deny), DLP exfil simulations, event-schema conformance.

Obsolete and withdrawn documents should not be used; please use replacements.

5. Deploy & runtime

- Enforce deny-by-default for sensitive classes with purpose context; block plaintext export paths for protected classes.
- Encryption by policy per CEK cryptographic profiles; prove via platform posture and KMS/HSM logs.
- Unified event schema (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision); append-only/tamper-evident storage and SIEM correlation.

6. Post-deploy validation & operations

- Continuous validation: discovery coverage/latency, ABAC suites, DLP exfil simulations, encrypted-restore drills, anomaly MTTD checks.
- Track Security SLOs: coverage $\geq 98\%$, tagging latency ≤ 24 hours; out-of-policy denials $\geq 99\%$; DLP block $\geq 95\%$ with FP/FN bounds; encrypted-restore meets RTO/RPO; schema conformance = 100 %.
- Auto-generate a Data Evidence Pack per release (policy diffs, validation results, deny/allow logs, DLP tests, encrypted-restore results, anomaly timelines, ADR links) to be summarized in EP-05.0 and linked to EP-05.1/05.2/05.3 as applicable.

C. Identity, Access, and Secrets (normative alignment to §6.2–§6.6)

- Enforce purpose-bound ABAC at data gateways and platforms; require MFA/JIT for privileged data actions.
- Secrets never in repos or images; inject via approved services with audit trails; redact sensitive fields in logs per class.
- Error/deny semantics are deterministic and include trace_id; telemetry meets the required schema.

D. Data Supply-Chain Integrity (normative; mechanics in Annex J)

- Only deploy policy bundles and data jobs whose tests passed gates; restrict sources/namespaces for policy artifacts.
- Quarantine and verify third-party connectors/plugins; enforce integrity and license checks.
- Separate build and deploy identities; forbid production writes from build jobs; treat policy tampering as release-blocking.


E. Measurement & Acceptance (aligned to §6 and §12)

- **Classification & Governance:** coverage $\geq 98\%$; tagging latency ≤ 24 hours; retention bound to class; lineage updates proven.
- **Access:** deny-by-default and purpose decisions on sensitive classes; MFA/JIT trails for privileged actions.
- **Protection:** encryption by policy per CEK; no plaintext export paths; WORM retain/deny-alter.
- **Monitoring:** DLP block $\geq 95\%$ with FP/FN bounds; event schema at ingest = 100 %; MTTD ≤ 15 minutes on sensitive classes.
- **Evidence:** every change links §5 to §6 to §12 via an Evidence Pack ID (EP-05.1/05.2/05.3, summarized in EP-05.0).

Obsolete and withdrawn documents should not be used; please use replacements.

Figure 13-1. Technical Implementation Workflow:

This illustrates the D05 technical implementation workflow for Defensible Data Security Architecture. The workflow shows how practitioners translate data security implementation guidance into controlled engineering actions across organizing principles, pipeline-stage guardrails, identity, access, and secrets management, data supply-chain integrity, gated release controls, measurement, and evidence production. The figure emphasizes that data security changes must be version-controlled, peer-reviewed, tested through non-bypassable security gates, validated through measurable acceptance criteria, and linked to EP-05.x Evidence Pack artifacts for traceability across §5 Requirements, §6 Technical Specifications, and §12 Verification and Validation.



Obsolete and withdrawn documents should not be used; please use replacements.

Common Pitfalls (and the engineered countermeasure)

1. Controls without classification. Require discovery/coverage gates and a tag to control bindings before rollout.
2. Policy-only ABAC. Test deny-by-default and purpose; block promotions if any sensitive class is allowed by default.
3. Encryption posture without resilience. Require encrypted-restore drills and KMS/HSM rotation/use logs (CEK-aligned).
4. DLP is not tied to tags. Enforce tag-driven rules and SIEM correlation; measure FP/FN bounds.
5. Logs without immutability or schema. Store append-only/WORM; validate required fields, especially trace_id/policy_decision.

	Practitioner Guidance: • Start from tags. For each sensitivity class, list bound controls (ABAC, DLP, retention/erasure, masking/tokenization, logging) and the Test-IDs that prove them; file artifacts under the Evidence Pack ID (EP-05.1/05.2/05.3, summary EP-05.0). • Measure the core four weekly. Coverage/latency, out-of-policy denial rate, DLP block/quarantine rate (with FP/FN bounds), encrypted-restore attainment (RTO/RPO). • Keep CEK separation and traceability. Say “per CEK cryptographic profiles”; prove with KMS/HSM logs and posture scans; reference the Evidence Pack ID everywhere.
---	---

	Quick Win Playbook: Title: Fail-Closed PR and Pre-Deploy Gates for One Data Domain Objective: Enforce “classification to ABAC deny-by-default to DLP block to logging schema” gates that fail closed before promotion, with artifacts linked to EP-05.x. Target: Wire non-bypassable PR and pre-deploy gates for one high-value data domain (§13.A.2, §13.B.2–4; §6.1, §6.2, §6.4, §6.6; §12). Component/System: Repo (policies-as-code), catalog/classification jobs, policy engine (ABAC), endpoint/email/cloud-sync DLP, SIEM, append-only/WORM evidence store (if used). Protects: Prevents unclassified data, permissive access, out-of-policy exfiltration, and non-conforming telemetry from reaching production.
---	--

Obsolete and withdrawn documents should not be used; please use replacements.

	Stops/Detects: Missing tags, allow-by-default on sensitive classes, exfil of tagged samples, and events lacking trace_id/policy_decision. Action: • Add CODEOWNERS and a “Data-Model Delta” template to PRs that change schemas/tags/policies. • Require discovery/coverage gate, ABAC deny-by-default tests (purpose attribute present), DLP exfil simulations, and event-schema checks to pass (fail-closed). • Run policy and schema drift detection against the approved baseline before promotion; block on drift. • Canary the policy bundle; on pass, promote. Proof (Evidence Packs): • EP-05.1 — Catalog/tag diffs, coverage/latency outputs. • EP-05.2 — ABAC policy bundle, deny/allow logs, DLP simulation outputs, SIEM correlation views, schema validator results. • EP-05.3 — Immutability settings (if used) and any WORM-related evidence for artifacts or recovery evidence. • EP-05.0 — Summary linking the above packs and gate results. Metric: Coverage ≥ 98 % with latency ≤ 24 hours; out-of-policy denials ≥ 99 %; exfil block ≥ 95 % with FP/FN bounds; event-schema conformance = 100 %. Rollback: Temporarily set gates to warn-only and revert the policy bundle; record an exception with owner and expiry in EP-05.0 and link affected artifacts in EP-05.1/EP-05.2.
--	--

Obsolete and withdrawn documents should not be used; please use replacements.

Appendices

Appendix A: EP-05 Engineering Traceability Matrix (ETM)

The ETM is the single record sheet that ties this Parent Standard together. It shows, row by row, how a Requirement (Inputs, §5) is realized as a Technical Specification (Outputs, §6), anchored by the relevant Core Principles (§7), backed by external Control Mappings (§9), and proven through named Verification (build-correct) and Validation (works-right) tests (§12) with links to the exact Evidence Pack ID (EP-05.x) holding artifacts. Use it to plan, execute, and audit work without searching multiple documents.

Table A-1. Inputs (§5) to Outputs (§6) to Principles (§7) to Controls (§9) to V&V (§12) to Evidence Packs

Req ID	Requirement (Inputs) (§5)	Technical Specification (Outputs) (§6)	Core Principles (§7)	Control Mappings (§9)	Verification (Build-Correct) (§12)	Validation (Works-Right) (§12)	Evidence Pack ID
R5.1	Enterprise data catalog & classification framework	TS6.1 Data Classification & Governance	RP-05 Secure by design; RP-10 Secure defaults; RP-15 Evidence production	CSA CCM DSI-01; CIS 8.x (logging scope for catalog events)	Catalog/tag schema present; owners/stewards recorded; lineage pointers set	Coverage ≥ 98 %; tagging latency ≤ 24 hours on sampled changes	EP-05.1
R5.2	Sensitivity tags bound to controls	TS6.1 Data Classification & Governance; TS6.2 Secure Data Access; TS6.4 DLP; TS6.6 Logging	RP-01 Least privilege; RP-03 Complete mediation; RP-15 Evidence production	CSA CCM DSI-03; CIS 3.4; CIS 13.7; OWASP API3	Tag to control binding sheet (ABAC/DLP/retention/logging) approved	Bound controls trigger as expected (deny/allow; block/quarantine; retention/masking enforced)	EP-05.1
R5.3	Zero-Trust Data Access baseline (ABAC/MFA/JIT)	TS6.2 Secure Data Access (ZTDA); TS6.6 Logging	RP-01 Least privilege; RP-02 Zero Trust; RP-03 Complete mediation	CIS 3.4; CSA CCM IAM-05; OWASP ASVS V2.1	ABAC policies with deny-by-default; MFA/JIT configured; purpose attribute present in decisions	Out-of-policy denials ≥ 99 %; elevation approval + auto-revocation logs present	EP-05.2

Obsolete and withdrawn documents should not be used; please use replacements.

Req ID	Requirement (Inputs) (§5)	Technical Specification (Outputs) (§6)	Core Principles (§7)	Control Mappings (§9)	Verification (Build-Correct) (§12)	Validation (Works-Right) (§12)	Evidence Pack ID
R5.4	Encryption policy & KMS integration (architectural)	TS6.3 Encryption & Data Protection (per CEK)	RP-18 Protect confidentiality; RP-19 Protect integrity; RP-10 Secure defaults	CIS 3.11; CIS 3.12; CSA CCM EKM-02; OWASP ASVS V9.1	Encryption by policy at rest/in transit; KMS/HSM integrated; posture scans available	Encrypted restore meets RTO/RPO; key rotation/usage logs meet policy	EP-05.3
R5.5	DLP channel coverage & policy registry	TS6.4 DLP & Monitoring; TS6.6 Logging/Correlation	RP-04 Defense in depth; RP-18 Protect confidentiality	CIS 13.7	DLP policies deployed (endpoint, network, cloud/SaaS); SIEM correlation configured	Exfil simulations blocked ≥ 95 % with FP/FN bounds; correlated to access events.	EP-05.1 (tags) + EP-05.2 (decision/correlation logs)
R5.6	Centralized data access logging & schema	TS6.6 Logging, Auditing & Anomaly Detection	RP-03 Complete mediation; RP-15 Evidence production	CIS 8.x; OWASP API3 (exposure context)	Standard event schema enforced (timestamp, subject, source, object, action, result, purpose, trace_id, policy_decision); tamper-evident storage	MTTD ≤ 15 minutes on sensitive classes; end-to-end reconstruction succeeds	EP-05.2 (decision/logs) and/or EP-05.3 (immutability proof)
R5.7	Immutable backup & recovery objectives	TS6.5 Data Integrity & Resilience	RP-19 Protect integrity; RP-20 Protect availability	CIS 11.3	WORM retention configured; multi-region copies set; integrity jobs scheduled	Alter/delete attempts are denied; quarterly restores meet RTO/RPO; hash verifications pass.	EP-05.3
R5.8	Residency & processing constraints (technical)	TS6.1 Governance (class/residency flags); TS6.2 Access; TS6.3 Encryption	RP-05 Secure by design; RP-04 Defense in depth	ISO/IEC 27001/27002 (clause-mapped in sub-standard s)	Residency/sovereignty flags mapped to classes/systems; routings documented	Audit sample shows required controls active for region-	EP-05.1 (mappings) + EP-05.3 (posture)

Obsolete and withdrawn documents should not be used; please use replacements.

Req ID	Requirement (Inputs) (§5)	Technical Specification (Outputs) (§6)	Core Principles (§7)	Control Mappings (§9)	Verification (Build-Correct) (§12)	Validation (Works-Right) (§12)	Evidence Pack ID
		(per CEK); TS6.6 Logging				bound datasets	
R5.9	Discovery & lineage coverage targets	TS6.1 Data Classification & Governance	RP-05 Secure by design; RP-15 Evidence production	CSA CCM DSI-01	Discovery tooling enabled with KPIs/owners; lineage captures set	KPI report shows coverage ≥ 98 %; lineage matches observed flows	EP-05.1
R5.10	Evidence & metrics readiness	TS6.1–6.6 (dashboarded SLOs)	RP-15 Evidence production	(Framework-agnostic; reporting)	Dashboards/queries exist for coverage, denials, DLP, KMS, and drills	Baselines captured; post-change deltas recorded under EP.	EP-05.0 (summary), links to EP-05.1/05.2/05.3

Notes for practitioners

- Use this ETM as the single sheet of record: every row must name the Output (§6), the Principle (§7), the external Control(s) (§9), and the exact Test-IDs and Evidence Pack ID(s) where artifacts live.
- Where a row references two packs (for example, EP-05.1 and EP-05.2), list both in the Evidence column of your working ETM and hyperlink to the sub-folders for the release.
- Keep metric phrasing and thresholds identical to §6 and §12 (for example, ≥ 98 %, ≤ 24 hours, ≥ 95 %, RTO/RPO, MTTR ≤ 15 minutes).

Obsolete and withdrawn documents should not be used; please use replacements.

Appendix B: EP-05 Evidence Pack Matrix

This appendix helps practitioners assemble, version, and audit Evidence Packs for Data Security Architecture (Annex E). The matrix lists the parent Evidence Pack for the annex and the current sub-packs aligned to §5 Requirements and §6 Technical Specifications. Evidence Packs are versioned per release and referenced in §12 traceability.

Table B-1. EP-05 Summary Matrix

Layer	EP Identifier	Purpose	Evidence Categories Included
Parent EP	EP-05	Serves as the master Evidence Pack for the Data Security Architecture Parent Standard (ISAU-DS-DS-1000). Stores annex-level artifacts, global V&V evidence, and cross-cutting design documentation supporting §5, §6, §10, and §12.	• Data security architecture diagrams • Data trust boundaries and gateways • Catalog and lineage overview maps • Invariants register • Interface Control Documents (ICDs) for data paths • Architecture Decision Records (ADRs) • Annex-level V&V summaries • Cross-domain logs/scan manifests • Evidence index linking EP-05.x packs
Sub-EP	EP-05.0	Readiness Gate summary for the program. One-pager used before design to confirm §5 inputs are green and to link all sub-packs for the current release.	• Readiness checklist (5.1–5.10) with owners/status • Controls to Outputs to Tests sheet (domain summary) • SLO baselines (coverage, denial rate, DLP efficacy, MTTD, RTO/RPO) • Links to EP-05.1/05.2/05.3 for artifacts
Sub-EP	EP-05.1	Catalog, Tags, Owners, Lineage, Retention — backs §5.1, §5.2, §5.9 and maps to §6.1. Establishes a tag to control bindings that drive enforcement.	• Classification policy and tag schema • Catalog exports (dataset, owner/steward) • Lineage snapshots and change deltas • Retention/hold flags mapped to classes • Discovery coverage and tagging latency reports

Obsolete and withdrawn documents should not be used; please use replacements.

Layer	EP Identifier	Purpose	Evidence Categories Included
			• Tag to control binding sheets (ABAC, DLP, retention, logging)
Sub-EP	EP-05.2	Zero-Trust Data Access (ABAC/MFA/JIT) — backs §5.3 and maps to §6.2/§6.6. Proves deny-by-default and purpose-bound decisions for sensitive classes.	• ABAC policy bundles and evaluation logs • Deny/allow decision logs mapped to tags • MFA/JIT elevation approvals and auto-revocation trails • Purpose attributes and policy_decision samples • Test results for out-of-policy denial SLO • SIEM correlation snapshots for access events
Sub-EP	EP-05.3	Encryption & Recovery (CEK-aligned) — backs §5.4/§5.7 and maps to §6.3/§6.5. Demonstrates encryption by policy, KMS/HSM governance, immutability, and encrypted restore.	• Transport/storage posture scans (per CEK profiles) • KMS/HSM key rotation and key-use logs • WORM configuration exports and deny-alter events • Hash/integrity verification job outputs • Encrypted restore drill reports with RTO/RPO attainment • No-plaintext-export path validations
Sub-EP (Reserved)	EP-05.4	Reserved for DLP Policy Engineering & Monitoring (if the program elects to split DLP evidence). When used, maps primarily to § 6.4 or § 6.6.	• DLP policy registry keyed to tags • Exfiltration simulation results with FP/FN bounds • Channel coverage evidence (endpoint, network, cloud/SaaS) • SIEM correlation dashboards and alert timelines
Sub-EP (Reserved)	EP-05.5	Reserved for Logging Schema & Anomaly Detection (if separated). Maps to §6.6.	• Event-schema validators and conformance reports • Append-only/immutability settings • MTTD measurement runs and investigation timelines • End-to-end reconstruction samples

Obsolete and withdrawn documents should not be used; please use replacements.

Layer	EP Identifier	Purpose	Evidence Categories Included
Sub-EP (Reserved)	EP-05.6–EP-05.9	Reserved for future Data sub-standards published through Open Season. Follows this table's structure and inherits annex evidence rules.	To be defined per sub-standard scope; will include Inputs to Outputs to Tests mappings and named Test-IDs with SLOs

Usage notes

- Reference EP-05 (parent), EP-05.0 (readiness), and the applicable EP-05.x sub-packs in §4 Use Case Evidence, §6 Evidence lines, and §12 Traceability Matrix.
- Each Evidence Pack must be stored in append-only/WORM locations where specified, with linked Test-IDs and SLO results.
- For audits, begin with EP-05.0 to see status and links, then drill into EP-05.1/05.2/05.3 for artifacts proving each §6 output.

Adoption References

NOTE: ISAUnited Charter Adoption of External Organizations.

ISAUnited formally adopts the work of the International Organization for Standardization / International Electrotechnical Commission (ISO/IEC) and the National Institute of Standards and Technology (NIST) as foundational standards bodies, and the Center for Internet Security (CIS), the Cloud Security Alliance (CSA), and the Open Worldwide Application Security Project (OWASP) as security control–framework organizations. This adoption aligns with each organization's public mission and encourages use by practitioners and institutions. ISAUnited incorporates these organizations into its charter so that every Parent Standard and Sub-Standard is grounded in a common, defensible foundation.

a) Foundational Standards (Parent level).

ISAUnited adopts *ISO/IEC* and *NIST* as foundational standards organizations. Parent Standards align with these bodies for architectural grounding and auditability and extend that foundation through ISAUnited's normative, testable specifications. This alignment does not supersede *ISO/IEC* or *NIST*.

b) Security Control Frameworks (Control level).

ISAUnited adopts *CIS*, *CSA*, and *OWASP* as control framework organizations. Control mappings translate architectural intent into enforceable technical controls within Parent Standards and Sub-Standards. These frameworks provide alignment at the implementation level rather than at the foundational level.

Obsolete and withdrawn documents should not be used; please use replacements.

c) **Precedence and scope.**

Foundational alignment (*ISO/IEC*, *NIST*) establishes the architectural baseline. Control frameworks (*CIS*, *CSA*, *OWASP*) provide enforceable mappings. ISAUnited's security invariants and normative requirements govern implementation details while remaining consistent with the organizations it has adopted.

d) **Mapping.**

Each cited control mapping is tied to a defined output, an associated verification and validation activity, and an Evidence Pack ID, thereby maintaining end-to-end traceability from requirements to controls, tests, and evidence.

e) **Attribution.**

ISAUnited cites organizations by name, respects attribution requirements, and conducts periodic alignment reviews. Updates are recorded in the Change Log with corresponding evidence.

f) **Flow-downs.**

(Parent to Sub-Standard). Parent alignment to the International *ISO/IEC* and *NIST* flows down as architectural invariants and minimum requirements that Sub-Standards must uphold or tighten. Parent-level mappings to *CIS*, *CSA*, and *OWASP* flow down as implementation control intents that Sub-Standards must operationalize as controls-as-code, tests, and evidence. Each flow-down shall reference the Parent clause, the adopted organization name, the Sub-Standard clause that implements it, the associated verification/validation test, and an Evidence Pack ID for traceability. Any variance requires a written rationale, compensating controls, and a time-bounded expiry recorded with an Evidence Pack ID.

Obsolete and withdrawn documents should not be used; please use replacements.

Change Log and Revision History

Review Date	Changes	Committee	Action	Status
June 2026	Diagram layer integration	Task Group ISAU-TG39-2024	Published	Complete
March 2026	Standards v 1.0 Submitted for Peer Review	Technical Fellow Society	Peer review	Complete
January 2026	Standards v 1.0 Published Draft	Task Group ISAU-TG39-2024	Published	Complete
December 2025	Standards Revision	Standards Committee	Submitted	Complete
October 2025	Standards Revision	Task Group ISAU-TG39-2024	Submitted	Complete
December 2024	Standards Development (Parent D05)	Task Group ISAU-TG39-2024	Draft Complete	Complete

End of Document
IO.

Obsolete and withdrawn documents should not be used; please use replacements.