



Cyber Phenomenon Series

**TRUST
OPERATIONS
(TRUSTOPS)**

INTEGRATED OPERATING MODEL
MEASURABLE OUTCOMES: TRUST

Enterprise Convergence toward Consolidated Cyber Trust Operations™

Predicting the Convergence of IT, Security, Privacy, and AI Operations

Scott Foote

Last Updated: 12 February 2026

Phenomenati Consulting
www.phenomenati.com

6 Liberty Square, #2736
Boston, MA 02109
(508) 709-7990 (office)

CONFIDENTIALITY NOTICE: The contents of this document, including any attachments, are intended solely for stakeholders of Phenomenati Consulting, may contain confidential and/or privileged information, and are legally protected from disclosure.

<this page is intentionally blank>

CONFIDENTIAL

Contents

1	Executive Summary.....	1-1
1.1	Key Predictions.....	1-1
2	What are Cyber Trust Operations?.....	2-1
2.1	Principles.....	2-1
3	Why Convergence is Happening Now.....	3-1
3.1	Shared Data, Shared Failures.....	3-1
3.2	AI-driven Operations and Operations-driven AI.....	3-1
3.3	'Shift-Left' Becomes 'Shift-Everywhere'.....	3-1
3.4	Risk Frameworks are Converging around Governance + Lifecycle Management.....	3-1
3.5	Regulatory Obligations Tie AI, Security, and Privacy Together.....	3-1
3.6	A Unified Incident Lifecycle is Becoming Mandatory.....	3-1
4	The Convergence Map: From Siloed Operations to TrustOps.....	4-1
5	The TrustOps Operating Model.....	5-1
5.1	The Trust Platform: Telemetry, Control, Automation, Evidence.....	5-1
5.2	Core TrustOps Functions – a Capability View.....	5-1
5.3	What Changes: From 'Ticket Factories' to Productized Operations.....	5-1
6	Where GRC Fits and Why Independence Must Be Preserved.....	6-1
6.1	A Workable Separation-of-Duties Pattern.....	6-1
7	Trust Incidents: One Command Structure, Multiple Playbooks.....	7-1
7.1	A Shared Incident Lifecycle.....	7-1
7.2	Example: Combined Security + Privacy + AI Incident Workflow.....	7-1
8	Reference Architecture: the TrustOps Control-and-Evidence Fabric.....	8-1
8.1	Telemetry and Correlation.....	8-1
8.2	Detection, Investigation, and Response.....	8-1
8.3	Governance and Control Automation.....	8-1
8.4	Evidence-as-Code and Audit Readiness.....	8-1
9	A Practical Adoption Roadmap (0-24 months).....	9-1
9.1	0-90 days: Establish the TrustOps Nucleus.....	9-1
9.2	3-12 months: Build the Shared Platform.....	9-1
9.3	12-24 months: Continuous Controls and “Trust as a Product”.....	9-1
10	Risks and Anti-Patterns (and How to Avoid Them).....	10-1
10.1	Over-Centralization.....	10-1
10.2	Tool Sprawl Disguised as Strategy.....	10-1
10.3	Evidence without Meaning.....	10-1

10.4	Blurring Assurance Independence	10-1
10.5	AI 'Shadow Operations'	10-1
11	Conclusion	11-1
12	References	12-1

CONFIDENTIAL

1 Executive Summary

Over the next 3-5 years, most medium-large enterprises will face a practical choice: either invest in and coordinate five separate operational disciplines... **IT Operations**, **Security Operations**, **Privacy Operations**, **AI Operations**, and **GRC Operations**... or converge the overlapping parts into a single corporate capability focused on measurable stakeholder outcomes: **Trust**.

This paper predicts the emergence of a corporate-wide entity called **Cyber Trust Operations (TrustOps)**: an *integrated* operating model that unifies day-to-day information technology operations with information security, privacy, and AI lifecycle operations, built on a shared telemetry, control, and automation plane. The goal is not organizational vanity; it is to **maximize stakeholder trust** by reducing time-to-*detect*, time-to-*fix*, time-to-*comply*, and time-to-*prove* in an environment where outages, cyber incidents, privacy breaches, and AI failures increasingly intertwine.

The prediction is grounded in observable convergence *patterns* already underway: DevSecOps explicitly integrates security into development and operations workflows and generates security and compliance artifacts automatically [3]; AIOps is commonly defined as using big data and machine learning to automate IT operations tasks such as event correlation and anomaly detection [1]; and MLOps is explicitly described as unifying ML development with deployment and operations [2]. Meanwhile, leading frameworks increasingly treat cybersecurity, privacy, and AI as enterprise risk disciplines that require governance, continuous monitoring, and repeatable controls (e.g., NIST CSF 2.0 adds a core 'Govern' function [4], NIST Privacy Framework emphasizes enterprise risk management [5], and NIST AI RMF structures AI risk management around functions including 'Govern' [6]).

1.1 Key Predictions

- **TrustOps** becomes the **de facto** home for a **unified telemetry + automation platform** spanning **reliability**, **security** detection/response, **privacy** incident handling, and **AI** model/runtime monitoring (enabled by vendor-neutral telemetry standards such as OpenTelemetry [8]).
- The **incident lifecycle converges**: security incident response, privacy breach response (including notification obligations), and AI incident response will share triage, communications, evidence capture, and post-incident learning loops [12].
- Customers and regulators increasingly measure '**trust outcomes**' that cut **across domains** (security, availability, confidentiality, processing integrity, privacy) as reflected in widely used assurance criteria such as the AICPA Trust Services Criteria [13].
- **GRC capabilities** partially **converge operationally** (evidence collection, continuous control monitoring, control-as-code), but **independence must be preserved** for assurance roles consistent with the IIA's Three Lines Model, where internal audit remains independent and objective [14].
- Responsible AI **regulation accelerates** the need for AI operations integrated with security and privacy operations, including risk-based controls and governance (e.g., EU AI Act phased implementation timelines) [15].

2 What are Cyber Trust Operations?

Cyber Trust Operations is a corporate **operating capability** that delivers and sustains trustworthy digital services by integrating four operational disciplines:

- **IT Operations** (reliability, performance, capacity, change, service management)
- **Security Operations** (detection, response, vulnerability management, identity-centric operations)
- **Privacy Operations** (data protection operations, data subject request workflows, breach handling, privacy engineering collaboration)
- **AI Operations** (model lifecycle operations, AI risk controls, monitoring, incident management for AI systems)

TrustOps is not a rebranding of a Security Operations Center (SOC) or Network Operations Center (NOC). It is a broader operational system that manages **Risk** and **Resilience** across technology, data, and AI. In practice, TrustOps operates a *shared 'trust platform'* for telemetry, control enforcement, automation, and evidence generation, while **embedding trust requirements** into product, engineering, and business workflows.

TrustOps is aligned with outcomes that customers already recognize as **'trust'**: **security, availability, confidentiality, processing integrity, and privacy**, which are explicitly represented in the AICPA Trust Services Criteria used in SOC 2 examinations [13].

2.1 Principles

- **One Telemetry Fabric**: shared traces, metrics, logs, and contextual metadata enabling cross-domain correlation and root-cause analysis (e.g., OpenTelemetry as a vendor-neutral standard) [8].
- **One Control-and-Evidence Pipeline**: policy-as-code and control-as-code where feasible, producing audit-ready evidence continuously rather than quarterly scrambles.
- **One Incident Lifecycle with Multiple Playbooks**: security, privacy, and AI incidents share triage and communications patterns, while retaining specialist workflows and legal/regulatory steps (e.g., privacy breach notifications).
- **Federated Execution, Centralized Enablement**: TrustOps provides the platform, standards, and operational muscle; product and engineering teams remain accountable for their services.
- **Separation of Duties for Assurance**: audit independence and objective oversight remain outside TrustOps, consistent with the Three Lines Model [14].

3 Why Convergence is Happening Now

3.1 Shared Data, Shared Failures

Modern services are built on shared infrastructure, shared identities, shared data pipelines, and shared third-party platforms. As a result, an **outage**, a cyber **intrusion**, a privacy **breach**, and an AI **failure** are increasingly different facets of the same underlying event. That reality pushes organizations toward shared **detection**, shared **response** mechanics, and shared **evidence capture**.

3.2 AI-driven Operations and Operations-driven AI

AIOps and **MLOps** are explicit linguistic signals of convergence. Gartner defines *AIOps* as combining big data and machine learning to automate IT operations processes (including event correlation and anomaly detection) [1]. In parallel, *MLOps* is commonly defined as practices that unify ML development with deployment and operations [2]. These disciplines create an **operational bridge** between classic **ITOps** and modern **AI lifecycle management**.

3.3 'Shift-Left' Becomes 'Shift-Everywhere'

DevSecOps extends *DevOps* by integrating **security** practices into the pipeline and automatically generating security and compliance artifacts across **build**, **test**, **distribution**, and **deployment** workflows [3]. Once **compliance** and **evidence** are generated **continuously**, the boundary between operational controls and GRC evidence collection starts to blur (though the independence boundary for assurance must remain intact).

3.4 Risk Frameworks are Converging around Governance + Lifecycle Management

NIST CSF 2.0 elevated **governance** to a core function, signaling that cybersecurity **risk management** is **inseparable** from **enterprise governance** [4]. NIST's Privacy Framework positions **privacy** as an **enterprise risk management discipline** [5]. NIST AI RMF structures **AI Risk Management** around core functions, including 'Govern' across the **AI Lifecycle** [6]. Standards bodies are also codifying **AI Management Systems** (e.g., ISO/IEC 42001 as an AI Management System or "AIMS" standard) [7]. Together, these moves normalize an **integrated operational posture** where trust is managed continuously rather than audited retrospectively.

3.5 Regulatory Obligations Tie AI, Security, and Privacy Together

Privacy regulation already requires **operationalized breach handling** and **communications**. For example, GDPR requires **notifying supervisory authorities** of certain personal data breaches without undue delay and, where feasible, within 72 hours of awareness [16]. *Emerging* AI regulation adds **new operational duties** for AI systems. The European Parliament summarizes the EU AI Act as a Risk-Based Regulatory regime with phased applicability following adoption in June 2024 [15]. Operationally, this pushes organizations to treat **AI controls**, **Security controls**, and **Privacy controls** as a **single control fabric** with differentiated playbooks.

3.6 A Unified Incident Lifecycle is Becoming Mandatory

Incident Response already includes **detection**, **containment**, **eradication**, **recovery**, and **communications**. Not surprisingly, NIST SP 800-61 Revision 3 explicitly links Incident Response to **Risk Management** and includes incident reporting and other incident-related communications as part of **Response** and **Recovery** activities [12]. When Privacy notification and AI incident obligations are layered on top, a **shared operational backbone** becomes *the only scalable path*.

4 The Convergence Map: From Siloed Operations to TrustOps

The table below summarizes the practical overlap that is already driving convergence. The center column ('Shared TrustOps Backbone') is the nucleus of TrustOps: common data, workflows, and automation that reduce duplication.

Operational Domain	Traditional Focus	Shared TrustOps Backbone (converging capabilities)	Illustrative Signals and Artifacts
IT Operations (ITOps)	Service Reliability, Performance Management, Change Management, Incident/Problem Management	Unified Telemetry; SLOs/Error Budgets; Automated Remediation; Change Risk Scoring	Traces/metrics/logs (OpenTelemetry) [8]; Specialized Runbooks; SLO Dashboards
Security Operations (SecOps)	Threat Detection and Response; Vulnerability Management; Identity Operations	Telemetry Correlation; Automated Triage/Response (SOAR); Identity-Centric Controls; Shared Incident Lifecycle	SOAR automates and coordinates Security Tools and Incident Workflows [17]; Zero Trust Architectures emphasize protecting resources and <i>continuous verification</i> [18]
Privacy Operations (PrivacyOps)	Privacy Engineering coordination; Consent/Retention Operations; DSAR Workflows ; Privacy Incident Response	Evidence-as-Code; Data Inventory/lineage/provenance; Breach Triage and Notification workflows integrated with incident response	Privacy treated as Enterprise Risk Management/ERM (NIST PF) [5]; Breach Notification obligations (GDPR) [16]
AI Operations (AIOps/MLOps/ModelOps)	Model Lifecycle Operations; monitoring; Model Risk Controls; AI Incident Response	Centralized Model Registry and <i>approvals</i> ; Control Checks in pipelines; Shared Telemetry + Governance; Model Monitoring and <i>Drift Detection</i>	AIOps Definition highlights <i>Automation</i> of ITOps tasks [1]; MLOps unifies ML Development with Deployment/Operations [2]; AI Risk Management functions include 'Govern' [6]
GRC Operations (governance, risk, compliance)	Control Design/Testing; Risk Assessments; Compliance Reporting; Oversight and Challenge	Continuous Control Monitoring; Evidence Collection Automation; Policy-as-Code; Risk Reporting fed by Operational Telemetry	Cybersecurity Governance emphasized in NIST CSF 2.0 [4]; Assurance Criteria for Trust Outcomes (SOC 2/TSC) [13]

Note on GRC: TrustOps can (and should) automate evidence capture and enable continuous control monitoring, but independent assurance and audit objectivity must remain separated (see Section 5) [14].

5 The TrustOps Operating Model

5.1 The Trust Platform: Telemetry, Control, Automation, Evidence

TrustOps is best understood as an operating system, not only an organization chart. The Trust Platform is the minimum shared foundation that allows multiple teams to act as one during normal operations and during incidents.

- **Telemetry Plane:** standardized collection of traces, metrics, logs, and context across applications, infrastructure, and AI services (e.g., OpenTelemetry) [8].
- **Control Plane:** policy-as-code and control-as-code for security, privacy, and AI controls wherever feasible; centralized identity and access patterns aligned to zero trust principles [18].
- **Automation Plane:** orchestration of response and remediation workflows (SOAR for security, but extended to privacy and AI playbooks) [17].
- **Evidence Plane:** continuous generation of audit-ready artifacts from pipelines and runtime controls (a DevSecOps pattern explicitly includes automatic compliance artifact generation) [3].

5.2 Core TrustOps Functions – a Capability View

In mature implementations, TrustOps typically owns or coordinates these capabilities:

- **Trust Telemetry Engineering** (instrumentation standards, pipelines, correlation rules, data quality)
- **SRE-style Reliability Operations** (SLOs, incident command, post-incident learning; reliability and security are interdependent as highlighted in Google's SRE guidance) [19]
- **Detection and Response Engineering** (triage, automation, threat hunting, purple-team feedback loops)
- **Privacy Operations Enablement** (DSAR workflows, retention execution, breach operations and communications)
- **AI Operations** (model inventory/registry, deployment gates, monitoring for drift and misuse, AI incident response; aligned to AI governance/risk frameworks) [6][7]
- **Continuous Control Monitoring** and **Evidence Automation** to support GRC reporting and audits [4][13]

5.3 What Changes: From 'Ticket Factories' to Productized Operations

TrustOps should be **product-managed**. That means: *documented* service offerings, *published* SLOs for internal services, self-service where possible, and an *explicit* roadmap. Gartner's discussion of **Platform Engineering** emphasizes building internal platforms and '*paved roads*' that can **mandate security and architecture controls** through better defaults rather than coercion [20].

6 Where GRC Fits and Why Independence Must Be Preserved

TrustOps benefits from bringing **operational control enforcement** and **evidence generation** closer to the systems being controlled. However, the organization must preserve *independent* challenge and assurance to **avoid self-assessment** and *conflicts of interest*.

The **IIA's Three Lines Model** explicitly describes Internal Audit as providing *independent* and *objective* assurance on the adequacy and effectiveness of Governance and Risk Management, and emphasizes **third-line independence** [14]. In practical terms: TrustOps can operate controls and generate evidence, but should not be the final authority on its own effectiveness.

6.1 A Workable Separation-of-Duties Pattern

A common pattern is to separate responsibilities as follows:

Line	Typical Roles	TrustOps Interaction
First Line (ownership)	Product, Engineering, IT Operations, Data/AI teams	Own Risk in their services; Consume TrustOps platforms; Implement controls-as-code and reliability practices
Second Line (oversight and guidance)	Risk Management, Compliance, Privacy Office, Security Governance (parts of GRC)	Define policies and control objectives; Validate measurement approaches; Review exceptions; Leverage evidence generated by TrustOps
Third Line (independent assurance)	Internal Audit	Independently assesses design and operating effectiveness; validates TrustOps evidence and sampling approaches; reports objectively to governing body [14]

This model allows convergence where it reduces friction (shared telemetry, shared incident command, evidence automation), without collapsing the independence required for objective assurance.

7 Trust Incidents: One Command Structure, Multiple Playbooks

In a TrustOps model, '*incident*' is a single umbrella with multiple specialized playbooks. For example, a single intrusion can trigger: *availability* impacts, *security containment* actions, *privacy breach assessment* and *notification*, and AI model *misuse investigation*.

7.1 A Shared Incident Lifecycle

NIST **Incident Response** guidance describes a *lifecycle* that includes *preparation*, *detection/analysis*, *containment/eradication/recovery*, and *post-incident* activity, and highlights incident reporting and communications as part of Response [12]. TrustOps uses this lifecycle for every major incident, then attaches domain playbooks.

7.2 Example: Combined Security + Privacy + AI Incident Workflow

- **Triage and Declare:** single incident commander, shared timeline, initial severity and business impact.
- **Stabilize Service:** reliability actions first to stop harm propagation (traffic shaping, feature flags, failover).
- **Contain Threat:** security containment steps and forensic preservation; orchestration and automation reduce manual work (SOAR concepts) [17].
- **Assess Privacy Impact:** determine whether personal data is involved, scope data subjects, and prepare notification decisioning; GDPR breach notification timing is a key operational constraint [16].
- **Assess AI Impact:** determine whether AI systems contributed (e.g., prompt injection, data poisoning, model inversion, unsafe outputs), and whether AI governance controls were bypassed; align risk handling to AI RMF governance practices [6].
- **Evidence Capture:** automatically preserve logs, access trails, model versions, approvals, and change records; generate a defensible record.
- **Post-incident Learning:** one blameless review, with separate action streams for reliability fixes, security hardening, privacy remediation, and AI control improvements.

8 Reference Architecture: the TrustOps Control-and-Evidence Fabric

TrustOps is accelerated when the enterprise standardizes on a small set of shared building blocks. This section provides a pragmatic **Reference Architecture** that organizations can adopt.

8.1 Telemetry and Correlation

A *vendor-neutral telemetry standard* reduces friction between ITOps, SecOps, and AI operations. OpenTelemetry positions itself as open source and vendor-neutral, supporting unified collection of traces, metrics, and logs [8]. With consistent **context** propagation, operational data can be *correlated* across performance, security signals, and AI runtime behavior.

8.2 Detection, Investigation, and Response

Security tooling is already converging around **cross-domain telemetry correlation**. XDR is commonly described as collecting and correlating data across endpoints, networks, cloud workloads, email, and identity systems to improve detection and response [21]. SOAR platforms coordinate security tools and streamline incident response workflows by automating repetitive tasks [17]. TrustOps *extends* these patterns to **Privacy** and **AI** playbooks (e.g., automatic breach assessment checklists; model rollback and quarantine actions).

8.3 Governance and Control Automation

TrustOps *operationalizes governance frameworks* through automation. For cybersecurity, **CSF 2.0's 'Govern' function** formalizes governance as part of the framework core [4]. For AI, **ISO/IEC 42001** defines requirements for an AI management system to manage **Risks** and **Opportunities** across the **AI Lifecycle** [7]. For privacy, NIST's Privacy Framework provides a Risk Management approach for Privacy Risk within Enterprise Risk Management [5]. The common move is to express control intent in ways that can be tested and evidenced *continuously*.

8.4 Evidence-as-Code and Audit Readiness

The competitive advantage of TrustOps is *'time-to-prove'* - the ability to demonstrate control effectiveness quickly during **customer due diligence** or **regulatory inquiry**. This aligns naturally with assurance frameworks that customers already ask for, such as SOC 2 based on the AICPA Trust Services Criteria [13].

9 A Practical Adoption Roadmap (0-24 months)

9.1 0-90 days: Establish the TrustOps Nucleus

- Define TrustOps **charter**, **scope**, and **boundaries** (including *GRC independence* and *decision rights*) [14].
- Standardize **Incident Command**: a single major-incident process with security, privacy, and AI playbooks attached [12].
- Establish **telemetry standards** and **minimum instrumentation requirements** (start with critical services) [8].
- Create a **unified Risk-and-Control backlog**: *prioritize* top Operational Risks that span **reliability**, **security**, **privacy**, and **AI**.

9.2 3-12 months: Build the Shared Platform

- Deploy a **unified telemetry pipeline** and **correlation capability**; *integrate* IT **observability** with **security detections** and **AI runtime monitoring** [8][21].
- Implement **automation** for repetitive response tasks (SOAR patterns) and **measurable reduction** in mean time to detect/respond [17].
- Adopt DevSecOps practices that generate **security and compliance artifacts** *through* the pipeline [3].
- Create an **AI system inventory** and **model registry**; add **deployment gates** aligned to AI RMF/ISO 42001 governance requirements [6][7].
- **Operationalize privacy workflows** (deletion/retention execution, DSAR, breach response) and integrate with **incident processes** [10][16].

9.3 12-24 months: Continuous Controls and “Trust as a Product”

- Move to continuous **control monitoring** and **evidence automation** for high-value controls used in audits and customer assurances (SOC 2, ISO 27001, etc.) [13].
- **Institutionalize reliability + security engineering** loops (secure-by-design baselines, post-incident action tracking, error budgets) [19].
- Integrate **Regulatory Readiness** for **emerging AI rules** using **Risk-Based control packs** and **operational reporting** (e.g., EU AI Act phased obligations) [15].
- Publish TrustOps **Service Catalog** and 'paved roads' that make secure, private, and compliant delivery the **default path** (platform engineering approach) [20].

10 Risks and Anti-Patterns (and How to Avoid Them)

10.1 Over-Centralization

- TrustOps should **enable** teams, not become a bottleneck.
- Use federated execution with **self-service platforms** and **clear service levels**.

10.2 Tool Sprawl Disguised as Strategy

- Converge on a small number of platforms, and **enforce telemetry standards** and integration requirements [8].

10.3 Evidence without Meaning

- Automate evidence only for **controls that are tied to risk and business outcomes**; align metrics to recognized criteria (e.g., trust outcomes) [13].

10.4 Blurring Assurance Independence

- Keep **Internal Audit independent** and **preserve second-line challenge functions** per the Three Lines Model [14].

10.5 AI 'Shadow Operations'

- Require **model/system inventory, approvals, and monitoring**;
- Align to AI governance frameworks and standards [6][7].

11 Conclusion

The convergence of ITOps, SecOps, PrivacyOps, and AIOps is already visible in the language, tooling, and frameworks shaping modern operations (DevSecOps [3], AIOps [1], MLOps [2], and governance-centric risk frameworks [4][5][6]). TrustOps is the logical endpoint: an operating entity that measures and improves trust outcomes continuously, powered by shared telemetry, automation, and evidence.

Organizations that operationalize TrustOps early will reduce duplication, shorten incident cycles, and accelerate their ability to demonstrate trust to customers, partners, and regulators.

CONFIDENTIAL

12 References

- [1] Gartner. “*AI Ops (Artificial Intelligence for IT Operations)*.” Gartner Glossary. Accessed Feb. 20, 2026. <https://www.gartner.com/en/information-technology/glossary/aiops-artificial-intelligence-operations>
- [2] Amazon Web Services. “*What is MLOps?*” AWS. Accessed Feb. 20, 2026. <https://aws.amazon.com/what-is/mlops/>
- [3] NIST NCCoE. “*Secure Software Development, Security, and Operations (DevSecOps) Practices*.” Accessed Feb. 20, 2026. <https://www.nccoe.nist.gov/projects/secure-software-development-security-and-operations-devsecops-practices>
- [4] NIST. “*NIST Releases Version 2.0 of Landmark Cybersecurity Framework*.” Feb. 26, 2024. Accessed Feb. 20, 2026. <https://www.nist.gov/news-events/news/2024/02/nist-releases-version-20-landmark-cybersecurity-framework>
- [5] NIST. “*NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management, Version 1.0*.” NIST CSWP 01162020. Jan. 16, 2020. <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.01162020.pdf>
- [6] NIST. “*NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0)*.” NIST AI 100-1. 2023. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- [7] ISO. “*ISO/IEC 42001: AI management systems*.” Accessed Feb. 20, 2026. <https://www.iso.org/standard/42001>
- [8] OpenTelemetry. “*What is OpenTelemetry?*” Documentation. Accessed Feb. 20, 2026. <https://opentelemetry.io/docs/what-is-opentelemetry/>
- [9] ISO, “*ISO/IEC 27701: Privacy Information Management Systems (PIMS)*,” ISO standard page (noting newer revision availability). Available: <https://www.iso.org/standard/71670.html>
- [10] IAPP. “*Measuring Privacy Operations Report*.” Accessed Feb. 20, 2026. <https://iapp.org/resources/article/measuring-privacy-operations>
- [11] Phenomenati, “*Phenomenati’s Taxonomy of a SOC™*: A Reference Model of operational needs to guide the evolution of Security Operations,” PDF (2019-04-05). Available: <https://img1.wsimg.com/blobby/go/3ad13048-c1b5-4403-aff5-ab0dcf0c2e01/downloads/Phenomenati-SOC-Taxonomy%202019-04-05.pdf>
- [12] NIST. “*Incident Response Recommendations and Considerations for Cybersecurity Risk Management (SP 800-61 Rev. 3)*.” 2025. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r3.pdf>
- [13] AICPA. “*2017 Trust Services Criteria (with Revised Points of Focus - 2022)*.” Accessed Feb. 20, 2026. <https://www.aicpa-cima.com/resources/download/2017-trust-services-criteria-with-revised-points-of-focus-2022>
- [14] The Institute of Internal Auditors (IIA). “*The IIA’s Three Lines Model*.” July 2020. <https://www.theiia.org/globalassets/documents/resources/the-iias-three-lines-model-an-update-of-the-three-lines-of-defense-july-2020/three-lines-model-updated-english.pdf>
- [15] European Parliament. “*EU AI Act: First Regulation on Artificial Intelligence*.” Updated Feb. 19, 2025. <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
- [16] legislation.gov.uk. “*Regulation (EU) 2016/679 (GDPR), Article 33 - Notification of a personal data breach to the supervisory authority*.” <https://www.legislation.gov.uk/eur/2016/679/article/33>
- [17] IBM. “*What is SOAR (Security Orchestration, Automation and Response)?*” Accessed Feb. 20, 2026. <https://www.ibm.com/think/topics/security-orchestration-automation-response>
- [18] NIST. “*Zero Trust Architecture (SP 800-207)*.” 2020. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>
- [19] Google. “*Site Reliability Engineering (SRE) Book*.” Accessed Feb. 20, 2026. <https://sre.google/books/>
- [20] Gartner. “*Platform Engineering*.” Gartner Topics. Accessed Feb. 20, 2026. <https://www.gartner.com/en/infrastructure-and-it-operations-leaders/topics/platform-engineering>
- [21] Microsoft. “*What is XDR?*” Microsoft Security. Accessed Feb. 20, 2026. <https://www.microsoft.com/en-us/security/business/security-101/what-is-xdr>