The 2026 Institute for Digital Integrity (IDI)

Data Integrity Toolkit

April 2026

Carole House with Noah Ringler and Robert Rueter

Institute for Digital Integrity Data Integrity Toolkit V1.01
IDI is a licensed 501C3 Tax-Exempt Charitable Organization
theidi.org
CC BY-NC 4.0

License

The Data Integrity Toolkit is published under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0).

We built this framework as public infrastructure, and we want it to be used. You are free to share, adapt, translate, and build on the Toolkit, including under your own organization's branding, provided you credit the Institute for Digital Integrity (IDI) and link back to the original work. Commercial use, uses to generate fees or revenue, requires separate permission; contact us at Robert@TheIDI.org to discuss licensing.

Suggested attribution: "Adapted from the Data Integrity Toolkit, Institute for Digital Integrity (2026), licensed under CC BY-NC 4.0."

Full license terms: https://creativecommons.org/licenses/by-nc/4.0/

A LETTER FROM THE EXECUTIVE DIRECTOR

Introducing the Data Integrity Toolkit

Dear friends and colleagues,

Something fundamental is shifting beneath our shared digital life. AI now writes roughly half of the articles published online. Deepfakes have moved from novelty to commonplace. Photographs are forged, scientific datasets are quietly altered, and the historical record itself has become editable by anyone with the right tools. It is now cheaper and faster to generate synthetic information than to create authentic content, and the gap widens every month.

This is not a reason for despair. It is a reason to build.

Today the Institute for Digital Integrity (IDI) is proud to release the **Data Integrity Toolkit**, a free, openly licensed resource that puts practical, standards-based data integrity within reach of everyone, not just those who can afford consultants and proprietary systems. Authored by Carole House with Noah Ringler and Robert Rueter, the Toolkit organizes the work of protecting data into nine functional domains, from cryptographic hashing and provenance metadata to privacy-preserving verification, long-term preservation, and governance. For each domain it names the standards that matter, the tools that implement them, and the practices that make data provably authentic, tamper-evident, and durable.

I want to be candid about what this first edition is and is not. It is a foundation, and it is a technical one by design. We began with standards because standards move the slowest and everything else must be built on them. What you hold is the reference layer: rigorous, comprehensive, and yes, wonky. The Toolkit is a living work, and the next phase is where it becomes operational. Our ongoing research is focused on mapping the ecosystem of research and implementation partners, and on developing concrete use cases and plain-language explainers for the communities that need them most: universities and academic researchers, NGOs, libraries, museums, newsrooms, and archives. If this edition tells you what good data integrity looks like, the next will show you, step by step, how to build it in your own context.

We built the Toolkit on a simple conviction: an informed democracy requires a trustworthy digital record, and the infrastructure that protects that record must be public, transparent, and accountable to all of us. Truth is not a partisan value. It is the ground we all stand on.

That is why the Toolkit is published under a Creative Commons license (CC BY-NC 4.0). You are free to use it, adapt it, translate it, and carry it into your own community under your own banner, with attribution. We want journalists, researchers, archivists, policymakers, technologists, and civil society organizations to make this framework their own.

CC BY-NC 4.0

Here is how you can help.

Adopt it. Read the Toolkit, run the assessment, and put its practices to work in your organization. Then tell us what worked, what didn't, and where you got stuck. Your experience directly shapes the use cases and implementation guides we are building next.

Join us. The ecosystem map is being drawn now, and we are seeking research and implementation partners: institutions willing to pilot the framework, contribute a sector use case, or co-develop the guides that will carry this work from standards to practice. If your university, library, museum, newsroom, or organization wants a place on that map, we want to hear from you.

Fund the mission. IDI is a nonprofit. Everything we build is public infrastructure, and public infrastructure requires public support. Your contribution funds the next phase of this work: the ecosystem map, the sector-specific implementation guides, the preservation of cultural memory under threat, and the Center for Information Integrity. A gift of any size moves this work forward, and we welcome conversations with foundations and institutional funders about deeper partnership.

The threats to our shared reality are real and accelerating. But so is the coalition forming to meet them. The standards exist. The tools exist. What has been missing is a clear, accessible path from principle to practice, and that is exactly the path we are now building, together.

Truth is worth defending. Let's build its infrastructure.

With gratitude and resolve,

Robert Rueter



Executive Director
Institute for Digital Integrity
Robert@TheIDI.org

CC BY-NC 4.0

EXECUTIVE SUMMARY

Most individuals and organizations lack the resources or expertise to implement robust data integrity practices. Expensive consultants and proprietary solutions are out of reach for many, leaving the majority exposed to manipulation and exploitation. Individuals and marginalized communities continue to face challenges securing and ensuring the integrity of the data they create, use, and share. Well-resourced institutions and AI model laboratories struggle with many of these same questions.

The **IDI's Data Integrity Toolkit** is a resource designed to help organizations of all sizes implement practical, scalable, and standards-based data integrity practices. The toolkit is built to provide of operationally useful resources built for ease of use by a wide variety of stakeholders. It provides step-by-step guidance, templates, and tools to ensure that data remains accurate, consistent, verifiable, and secure throughout its lifecycle. The toolkit equips organizations with the knowledge and resources to deploy cryptographic safeguards, governance controls, and lifecycle management practices, enabling them to protect their data, build trust, and unlock economic value in an increasingly complex digital landscape.

Data integrity has become a strategic imperative in the face of accelerating digital transformation, the rise of generative and agentic AI, and an evolving threat landscape. Roughly half of newly published online articles are now primarily AI-generated (49.9% in Q1 2026), a share that has held near 50% for more than a year; whether it resumes climbing as agentic AI matures is an open question.[1] More than half of web requests for HTML content now come from automated systems rather than humans (roughly 57.5%), with agentic AI the fastest-growing category, creating unprecedented challenges for core elements that support data integrity: trust, authenticity, and governance.[2] On current trajectories, humans online will be vastly outnumbered, acting with and through multi-agent systems at scale, both human-authorized and AI-agent authorized. Without robust data integrity frameworks, organizations risk falling victim to personalized misinformation, data poisoning, and operational failures, while missing opportunities to unlock the latent value of their data assets.

The economic stakes are equally high. Gartner research estimates that poor data quality costs the average organization $12.9 million annually, and organizations face mounting costs from breaches, regulatory penalties, and inefficiencies.[3] AI systems are only as good as they are trained on and interact with. Conversely, investments in data integrity drive measurable returns, including cost savings, operational efficiency, and competitive advantage, such as saving an estimated 1,345 working years annually for one nation through interoperable data exchange (per Estonia's self-reported assessment of its X-Road platform) and delivering a 355% three-year return on investment in an independent assessment of one major cloud provider's combined data security, governance, and compliance platform.[4,5]

The IDI Data Integrity Toolkit is designed to help address these challenges by providing a practical, adaptable framework for implementing data integrity across diverse organizational contexts. Its key features include strategic and economic framing grounded in real-world examples and use cases

CC BY-NC 4.0



that demonstrate the value of data integrity investments, and a comprehensive set of fully indexed, AI-digestible resources and interactive tools built upon authoritative standards and tools for implementing data integrity controls, from privacy-preserving technologies to cryptographic hashing to data lifecycle management.

The toolkit organizes data integrity into nine functional domains, each addressing a critical aspect of data integrity and security:

1. **Data Preparation, Minimization, and Digitization.** Focuses on ensuring data is collected, structured, and validated at the point of creation. Includes guidance on data minimization, schema validation, and analog-to-digital chain-of-custody documentation.

2. **Cryptographic Integrity: Hashing, Signing, and Verification.** Provides tools and standards for tamper detection, digital signatures, and post-quantum cryptography to ensure data authenticity and prevent unauthorized modifications.

3. **Provenance, Metadata, and Digital Attestation.** Covers embedding and verifying provenance metadata, digital attestation, and timestamping to establish trust in data origin and history.

4. **Data Lineage, Versioning, and Lifecycle Management.** Enables organizations to track data transformations, maintain version control, and ensure accountability across the data lifecycle.

5. **Access Control, Authentication, and Secure Delegation.** Focuses on role-based and attribute-based access controls, zero-trust architecture, and secure delegation to prevent unauthorized access and ensure accountability.

6. **Privacy-Preserving Verification and Consent.** Introduces advanced cryptographic techniques like zero-knowledge proofs, homomorphic encryption, and secure multi-party computation to enable verification without exposing sensitive data.

7. **Data Quality, Anomaly Detection, and AI Integrity.** Provides tools for real-time validation, anomaly detection, and AI training data governance to ensure data remains accurate, consistent, and fit for purpose.

8. **Archiving, Preservation, and Immutable Records.** Offers guidance on long-term digital preservation, decentralized archival methods, and tamper-evident storage of records to ensure data remains accessible and verifiable over decades.

9. **Regulatory Compliance and Governance Frameworks.** Aligns organizations with global standards and regulations, including GDPR, CCPA, ISO/IEC 27001, and the EU AI Act, to ensure compliance and accountability.

Domains to be Addressed in Future Toolbox Versions:

CC BY-NC 4.0

10. **Payments and Intellectual Property**: Enabling fair data economies that enable users' rights over their data.

11. **Data Institutions**: Supporting coops, trusts, unions, and new data collectives

12. **Innovation Ecosystems and Incentive Alignments**: Seed funding and strategic support for case studies, incubation, and acceleration of freedom technologies that embody the preceding principles.

The IDI Data Integrity Toolkit reframes data integrity as **shared critical infrastructure** essential for economic competitiveness, national security, and democratic resilience. By providing a practical and verifiable multi-layered framework for implementing state-of-the-art integrity controls, the toolkit empowers organizations to protect their data, build trust, and unlock innovation without sacrificing information security. Whether addressing the challenges of AI, complex agentic systems, cybersecurity, or regulatory compliance, the toolkit serves as a critical and practical public resource. The toolbox is a "How to" guide for integrating data integrity into AI enablement strategies and aligning with emerging global standards.

Thank You to Our Organizational Partners

- Project Liberty Institute
- Starling Lab for Data Integrity
- Human Rights Foundation
- Stanford University Human AI Institute
- Georgetown University's Schools of Public Policy, Law, and Business
- MIT Media LabComputer Science and AI Lab (CSAIL)
- CyberSMART: Georgetown, Virginia Tech, and Notre Dame
- Advanced AI Society
- Atlantic Council
- Linux Foundation
- First Person Project
- USC Marshall School's VanEck Digital Asset Initiative
- Milken Institute

CC BY-NC 4.0

Key Terms

Integrity: Assurance that data or a document (file) has not been altered, accidentally or deliberately, from the moment it was created or last verified. Integrity can be maintained during modifications to the underlying content if alterations are cataloged, and original versions are saved. *Definition adapted from NIST SP 800-175A, ISO/IEC 27000:2018.*[1]

Trust: Confidence that a system, organization, or identity will behave as expected and that the credentials or claims it presents are genuine and have not been fabricated. *Definition adapted from NIST SP 800-63 and ISO/IEC 38505-1.*[2]

Digital Identity: A verified, machine-readable representation of who or what an entity is via a person, organization, AI agent, or device used to establish that the verified parties are accessing verified resources. Definition *adapted from NIST SP 800-63.*[3]

Privacy Enhancing Technology: A tool or technique that allows data to be used, shared, or analyzed in ways that serve a legitimate purpose while ensuring that individuals' personal information is not exposed beyond what is strictly necessary. Definition *adapted from ISO/IEC 20889, NIST Privacy Framework, and ITIF Technology Explainer on Privacy Enhancing Technologies.*[4]

AI System: A machine-based system that processes inputs, such as data, text, or sensor readings, to generate outputs such as predictions, recommendations, or decisions that can influence real-world outcomes, and which operates with varying degrees of autonomy. Definition adapted from the *OECD Principles on AI (2019), updated 2024;*[5] *adopted by the EU AI Act (2024), Article 3(1)*[6] *and the NIST AI RMF (2023).*[7]

AI Agent: An AI system that can independently plan and execute a sequence of actions including using tools, querying external services, or directing other systems in order to complete a goal, without requiring a human to approve each individual step. As defined by the *NIST AI RMF Playbook (2023)*[8] *and ISO/IEC JTC 1/SC 42.*[9]

---

[1] SP 800-175A, Guideline for Using Cryptographic Standards in ...NIST Computer Security Resource Center (.gov)https://csrc.nist.gov › pubs › sp › final

[2] NIST SP 800-63 Digital Identity GuidelinesNational Institute of Standards and Technology (.gov)https://pages.nist.gov › 800-63-3

[3] Ibid.

[4] Retrieved from ITIF, https://itif.org/publications/2025/09/02/itif-technology-explainer-privacy-enhancing-technologies/.

[5] OECD Principles on AI (2019), updated 2024. Retrieved from https://oecd.ai/en/ai-principles

[6] EU AI Act (2024), Article 3(1). Retrieved from https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

[7] NIST AI RMF (2023). Retrieved from https://airc.nist.gov/RMF

[8] NIST AI RMF Playbook (2023). Retrieved from https://airc.nist.gov/Docs/1

[9] ISO/IEC JTC 1/SC 42. Retrieved from https://www.iso.org/committee/6794475.html

CC BY-NC 4.0

INTRODUCTION

Data integrity has become foundational infrastructure for economic competitiveness, national security, and democratic resilience. As societies digitize and deploy AI systems into core public and private sector services such as finance, governance, healthcare, justice, and elections, the trustworthiness of underlying data determines whether institutions function, fail, or are captured. Data integrity is no longer a narrow compliance concern; it is the substrate on which trust, accountability, and legitimacy rest. Data integrity determines the usefulness, viability, and lens through which AI systems and their operators view the world.

In recent years, converging and accelerating forces have elevated data integrity to a strategic imperative. Generative and agentic AI systems now operate at speeds and scales that magnify the consequences of corrupted, manipulated, or unverified data. Geopolitical adversaries increasingly target data environments as instruments of statecraft. Meanwhile, regulatory divergence across jurisdictions has turned data integrity capabilities into a source of competitive advantage (or exclusion) from global markets.

Critically, data is not only an operational input. It is the primary source of institutional **trust** in digital societies. When data provenance cannot be established, records cannot be authenticated, or errors cannot be corrected, public confidence erodes. This erosion precedes broader democratic decay. Trust in elections, courts, markets, and public administration depends on the integrity of the information those systems produce.

Data integrity is also a precursor to **digital dignity**. Individuals cannot exercise meaningful agency, consent, or rights in digital environments if the data describing them is inaccurate, opaque, or manipulable without recourse. The ability to inspect, correct, own, and verify personal data is a prerequisite for fairness, due process, and participation in modern economic and civic life.

This toolkit reframes data integrity as a **shared responsibility** that is simultaneously economic, strategic, and democratic. It provides policymakers, executives, and technical leaders with a practical, layered framework for understanding and implementing state-of-the-art integrity controls across the full data lifecycle, while highlighting where governance and regulatory frameworks lag behind technical capability. Data integrity is scaffolding for trust itself; when it fails, nothing above it can hold.

WHAT IS DATA INTEGRITY?

“Data integrity” means that data can be relied on for decisions, automation, accountability, and trust. It means that data remains accurate, authentic, and usable across its full lifecycle: from the moment it is created, through processing and sharing, to long-term storage or deletion. It is not a single control or technology, but a system-level outcome shaped by technical safeguards, governance choices, and institutional discipline.

At a basic level, data with integrity must be **accurate** (i.e., it correctly reflects reality), **consistent** (i.e., it does not conflict across systems or time), and **complete** (i.e., it is fit for its intended

CC BY-NC 4.0

purpose).[10] But integrity also depends on less visible properties: **authenticity** (i.e., knowing who created or attested to the data), **provenance** (i.e., being able to trace where it came from and how it has changed), **timeliness** (i.e., remaining valid in context), and **auditability** (i.e., being independently verifiable when it matters).[11]

Data integrity is closely related to concepts of cybersecurity, privacy, and availability, though it is not the same thing. A system can be secure and still produce wrong or misleading data. Data can be private and still be corrupted. And highly available data can still be untrustworthy. Integrity is the dimension that determines whether data deserves confidence in consequential use.

In digital societies, data integrity functions as critical infrastructure for trust.[12] Institutions make claims through data – about identity, eligibility, risk, value, and truth. When those claims cannot be verified, corrected, or explained, trust erodes. That erosion is rarely sudden, but it is cumulative, affecting markets, governance, and democratic legitimacy over time. The rise of agentic systems and services may even necessitate the creation of new data institutions and structures, as seen in Virginia in SB 384.[13]

Finally, data integrity is not binary. Different contexts demand different levels of assurance: archival evidence, financial transactions, AI training data, and public statistics each require tailored controls. This toolkit treats data integrity as a governed capability, scaled to risk and impact, not a one-size-fits-all standard.

DATA INTEGRITY IS CRITICAL INFRASTRUCTURE

Policymakers and executives have long treated data integrity primarily through the lens of regulatory compliance: a cost center that exists to satisfy audit requirements and avoid penalties. This framing, while understandable, is increasingly insufficient in the age of artificial intelligence. A more apt comparison is to physical infrastructure. Like bridges and power grids, data integrity frameworks are invisible when functioning well and catastrophically visible when they fail. And like infrastructure, the costs of underinvestment accumulate silently, manifesting in systemic crises that are difficult to attribute to their root cause.

Infrastructure is also characterized by network effects – a single bridge benefits all who cross it, and its absence imposes costs on the entire network. The same is true of data integrity. When financial systems exchange unverified data, when AI systems train on corrupted datasets, when supply chain partners operate from inconsistent records, the failures propagate. The externalities of poor data

[10] As defined in ISO/IEC 25012. Retrieved from https://www.iso.org/standard/35736.html
[11] Content Authenticity Initiative Glossary, Retrieved from https://opensource.contentauthenticity.org/docs/getting-started/glossary.
[12] https://www.atlanticcouncil.org/blogs/econographics/designing-a-blueprint-for-open-free-and-trustworthy-digital-economies/
[13] VA SB384, Studying IVOs, April 4, 2026, https://legiscan.com/VA/text/SB384/id/3349373.

CC BY-NC 4.0

integrity, much like the externalities of a crumbling bridge, are not fully borne by those who chose not to invest. In digital systems, trust is not an abstract value; it is an emergent property of verifiable data

The Changing Threat Landscape

Three convergent developments have elevated data integrity from a technical concern to a strategic priority in the period from 2023 to 2026.

Generative AI has fundamentally altered the threat landscape[14]. The same capabilities that allow AI systems to synthesize realistic text, images, audio, and video can be weaponized to produce synthetic data at industrial scale, in the forms of fabricated evidence[15], manipulated records, spread convincing deepfakes, and algorithmically generated information that has permeated the information ecosystem. The asymmetry is stark: the cost of producing synthetic or manipulated data is collapsing, while the cost of detecting and authenticating it remains high. Without robust provenance and integrity frameworks, AI amplifies the capacity for deception and degradation of our information systems (including AI) far faster than institutions can adapt.

Agentic AI systems – autonomous software agents that transact, negotiate, and execute decisions on behalf of human principals – introduce systemic vulnerabilities[16] of a different kind. These agents must be able to trust the data environments in which they operate. When they cannot, autonomous systems that behave correctly under clean data conditions may fail catastrophically when fed manipulated inputs. Further, AI agents must be able to trust and verify each other at scale in order to interact safely. The financial and operational risks of agentic systems operating on unverifiable or compromised data have not yet been fully priced into enterprise risk frameworks.

Geopolitical actors have, meanwhile, operationalized data integrity attacks as an instrument of statecraft. Supply chain compromises, the manipulation of scientific records, the corruption of open-source software repositories, and coordinated influence operations targeting electoral systems all represent data integrity failures.[17] While these all have documented precedents in the recent past, what distinguishes the current environment is the scale, sophistication, and strategic intentionality of these attacks.

Societal Benefits of Data Integrity Infrastructure

Data integrity is not just a technical concern; it enables the trust society needs to function. Its benefits extend across all sectors, from protecting the rights of vulnerable populations to driving operational efficiency for businesses and ensuring trust in public institutions. For vulnerable populations, data integrity safeguards against exploitation, misinformation, and systemic bias, enabling access to fair services and protections. For businesses, it reduces operational risks, enhances decision-making, and builds customer trust. For governments and institutions, it ensures

[14] Retrieved from https://www.europol.europa.eu/cms/sites/default/files/documents/Europol_Innovation_Lab_Facing_Reality_Law_Enforcement_And_The_Challenge_Of_Deepfakes.pdf

[15] Retrieved from https://www.ncsc.org/resources-courts/ai-generated-evidence-threat-public-trust-courts.

[16] Retrieved from https://www.cftc.gov/media/10626/TAC_AIReport050224/download.

[17] https://carnegieendowment.org/research/2025/01/managing-the-risks-of-chinas-access-to-us-data-and-control-of-software-and-connected-technology

CC BY-NC 4.0

transparency, accountability, and resilience in democratic processes. By embedding data integrity into digital systems, we create a foundation for equitable participation, innovation, and trust in the digital age.

Table 1: Societal Benefits of Data Integrity

Group	Key Benefits	Example
Small and Medium Businesses (SMBs)	Reduces operational risks and costs	Prevents data breaches that could cripple small businesses
	Enhances decision-making through accurate, reliable data	Improves supply chain efficiency with verified data
	Builds customer trust and regulatory compliance	Enables GDPR/CCPA compliance without expensive consultants
Large Enterprises	Supports scalability and interoperability across global operations	Ensures cross-border data compliance for multinational corporations
	Reduces downtime and operational inefficiencies	Prevents costly errors in financial reporting or AI model training
	Protects intellectual property and trade secrets	Secures proprietary datasets and algorithms
Governments and Institutions	Ensures transparency and accountability in public services	Protects election systems from tampering or misinformation
	Builds trust in democratic processes	Verifiable public health data during crises
	Strengthens national security by protecting critical infrastructure	Prevents manipulation of scientific records or supply chain data
Civil Society and Non-Profits	Protects sensitive data in high-risk environments	Safeguards human rights documentation from tampering
	Enables secure collaboration across borders	Cryptographically verified evidence for legal or advocacy purposes
	Builds trust with donors and stakeholders	Transparent reporting of financial and operational data
Vulnerable Populations	Protection from data manipulation and exploitation	Ensures accurate health records for underserved communities
	Access to fair and unbiased services	Prevents algorithmic bias in housing or loan applications
	Safeguards against misinformation and identity theft	Protects against deepfake-based scams targeting marginalized groups

ECONOMIC COSTS AND COMPETITIVE STAKES

The Cost of Data Failure and Value Creation from Data Integrity Investment

CC BY-NC 4.0

The economic case for data integrity extends beyond risk avoidance[18] to value creation, efficiency gains, and strategic optionality. In practical and quantifiable costs for integrity failures, organizations face direct costs in forms like breach remediation, regulatory penalties, and litigation that carry significant dollar signs. Indirect costs also carry a hefty price tag, like from the costs of decisions (whether on the trading room floor or in a voting booth) based on incorrect data, the foregone economic value of data assets that cannot or are not being monetized, reputational costs that can erode customer trust and market position, and the opportunity and innovation cost when organizations spend value resources – people, time, and tools – on data cleaning and validation rather than product development.

The best available empirical evidence is sobering: an analysis of 75 data quality assessments across a wide range of organizations found that 47% of recently created data records contained at least one critical error, and only 3% met basic quality standards, while Gartner estimates poor data quality costs the average organization $12.9 million annually[19]. These assessments predate the growing integration of AI-enabled tools into enterprises and decision making, and those costs may continue to balloon without addressing impacts of data quality underlying these models as well as the effects on human reasoning and potential "cognitive surrender"[20] being observed.

Trustworthy data environments reduce transaction costs, enable automation, and unlock interoperability, which are core drivers of modern economic growth.

Table 2: Select Examples of Data Integrity Investment Economic Impact

Organization / Data Governance Initiative	Reported Outcome	Mechanism
Estonia X-Road (e-Estonia / Cybernetica)[21]	1,345 working-years saved annually; 100% of government	Cryptographically secured, decentralized interoperability layer

[18] Retrieved from https://www.atlanticcouncil.org/blogs/geotech-cues/why-data-governance-matters/.
[19] Nagle, T., Redman, T., & Sammon, D. (2020). Assessing data quality: A managerial call to action. Business Horizons, 63(3), 325–337. https://doi.org/10.1016/j.bushor.2020.01.006; Gartner (2020), as reported in "Data Quality: Why It Matters and How to Achieve It," https://www.gartner.com/en/data-analytics/topics/data-quality.
[20] Shaw, S., & Nave, G. (2026). Thinking – Fast, Slow, and Artificial: How AI Is Reshaping Human Reasoning and the Rise of Cognitive Surrender. SSRN. https://doi.org/10.2139/ssrn.6097646
[21] Estonia X-Road (e-Estonia / Cybernetica). Retrieved from https://e-estonia.com/solutions/interoperability-services/x-road/

CC BY-NC 4.0

	services digitalized by December 2024; platform now deployed in over 20 countries	enabling real-time cross-agency data exchange without centralizing personal data
Snowflake AI Data Cloud[22]	354% ROI, $19.45M total benefits over 3 years; financial services firm improved reporting for over 15 global regulators	Unified data platform replacing fragmented, lack of auditable workflows; integrity controls enabling compliant cross-border data sharing
Intel[23]	25% reduction in new chip design cycle time following maturation of enterprise data governance model	Shift from overly restrictive governance framework; improved data accessibility for engineering teams without compromising integrity controls
ITIC 2024 Hourly Cost of Downtime Survey[24]	91%+ of mid-size and large enterprises report downtime costs exceeding $300,000/hr.; 44% report costs roughly $1-5+ million/hr.	Establishes the baseline cost avoided by automated data integrity validation and high-availability architecture; a quantifiable floor of integrity investment ROI
JAMA Network Estimated Costs of Waste in the US Healthcare System[25]	$760B–$935B in annual U.S. healthcare waste (~25% of total health spending); $191B–$286B in addressable savings	Failure of care coordination, redundant testing, and administrative complexity – all directly reducible through data integrity, interoperability standards, and accurate longitudinal records
UK Open Data[26]	Estimated up to GBP £90M annual travel time savings; substantial GDP contribution from open/public data	Data sharing and interoperable access facilitating new services, reduced costs
AWS/Forrester - Impact of Public Sector Data Integration[27]	Composite public sector org saved ~$23.3M over 5 years; 33% ROI	Interoperability and data integration reduced duplicative processes, improved service delivery

The Competitive Geography of Data Trust

Data integrity is not just a bureaucratic or organizational imperative, it is also a source of international competitive advantage as well as a point of regulatory and set of national interests that affect global

[22] Snowflake AI Data Cloud. Retrieved from https://tei.forrester.com/go/Snowflake/AIDataCloud/?lang=en-us
[23] Intel. Retrieved from https://www.tandfonline.com/doi/full/10.1080/10580530.2025.2477459
[24] ITIC 2024 Hourly Cost of Downtime Survey. Retrieved from https://itic-corp.com/itic-2024-hourly-cost-of-downtime-report/
[25] https://jamanetwork.com/journals/jama/article-abstract/2752664
[26] UK Open Data. Retrieved from https://www.oecd.org/en/publications/2019/11/enhancing-access-to-and-sharing-of-data_070835df/full-report/economic-and-social-benefits-of-data-access-and-sharing_836734cb.html
[27] AWS/Forrester - Impact of Public Sector Data Integration. Retrieved from https://aws.amazon.com/blogs/publicsector/forrester-study-examines-data-integration-roi-for-public-sector-organizations/

CC BY-NC 4.0

trade. The European Union's Data Act of 2025[28] and AI Act of 2024[29] establish mandatory data quality, provenance, and traceability requirements across a range of high-risk domains. Organizations operating in EU markets must meet these standards; those that cannot will face exclusion. This creates a powerful incentive structure, though one that currently benefits incumbents with mature data infrastructure over smaller organizations that may be technically sophisticated but governance-deficient.

China has elevated data to a sovereign strategic resource[30], systematically integrating collection and control across public and private sectors as deliberate grand strategy — not commercial byproduct. The result is a model of data-centric authoritarianism[31], built on population surveillance and algorithmic social control, that is actively being exported. Internationally, data amassed through nominally private Chinese firms creates asymmetric intelligence advantages and manipulation vectors that US and EU governance frameworks are only beginning to reckon with. The competition is ultimately architectural: infrastructure and standards choices made now will determine whose vision of data governance, and whose values, become the global default.

The United States, in contrast, lacks a federal data integrity framework. The resulting regulatory fragmentation, with sector-specific[32] regimes governing financial data, health data, and critical infrastructure, creates compliance complexity without strategic coherence. American organizations competing internationally operate at a governance disadvantage relative to EU peers who benefit from regulatory clarity, even when their underlying technical capabilities are equivalent.

DATA INTEGRITY AND DEMOCRATIC RESILIENCE

The Information Ecosystem as Critical Infrastructure

Democracy runs on shared facts. Electoral commissions certify vote counts, courts authenticate evidence, and regulators verify the data behind environmental, financial, and public health decisions. None of this works if the underlying information cannot be trusted.

That trust is now under pressure from three directions simultaneously: adversarial actors who manipulate records deliberately, institutions that govern data negligently, and AI systems evolving and dependent on that data are deployed faster than the frameworks needed to manage them. The result is a world where provenance cannot be established, records can be altered without detection, and synthetic content is indistinguishable from the real thing. The cost of that world is not just

[28] European Union’s Data Act of 2025. Retrieved from https://digital-strategy.ec.europa.eu/en/policies/data-act
[29] AI Act of 2024. Retrieved from https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
[30] China's Grand Data Strategy. Retrieved from https://www.hoover.org/research/chinas-grand-strategy-global-data-dominance.
[31] Retrieved from https://www.ned.org/data-centric-authoritarianism-how-chinas-development-of-frontier-technologies-could-globalize-repression-2/.
[32] Retrieved from https://www.dlapiperdataprotection.com/

CC BY-NC 4.0

economic. It is democratic. When citizens cannot trust the information their institutions produce, the institutions themselves lose legitimacy. That is not a technical problem, it is a constitutional one.

Data as the Foundation of Institutional Trust

Modern democratic institutions operate through data-mediated processes. Ballots become vote tallies, testimony becomes records, and policy decisions rest on quantified evidence. In each case, trust is not derived from authority alone, but from the ability to verify that data has been collected, preserved, and used with integrity.

When data integrity fails, trust collapses asymmetrically. Errors and manipulation disproportionately benefit those with power to exploit ambiguity, while ordinary citizens bear the consequences: miscounted votes, erroneous enforcement actions, and distorted markets. Over time, this asymmetry produces cynicism, disengagement, and susceptibility to disinformation[33]. Data integrity failures thus act as a leading indicator of democratic fragility[34].

Restoring trust requires more than counter-disinformation campaigns or content moderation. It requires decentralized infrastructure that makes falsification difficult, manipulation detectable, and correction routine. Provenance, auditability, and verifiability are democratic necessities, not technical luxuries.

Digital Dignity and the Rights Dimension

Data integrity is the precursor to digital dignity. Individuals cannot exercise autonomy, ownership, consent, or due process [35] in digital environments if the data describing them is inaccurate, inaccessible, or altered without their knowledge. Before questions of fairness, transparency, or accountability arise, there must be confidence that the underlying data is correct and contestable.

Data integrity carries a rights dimension that policy frameworks have been slow to recognize. Individuals who cannot access accurate records of their own data (e.g., their financial history, their health records, their legal status) are materially disadvantaged in their capacity to exercise agency in digital environments. This is not a theoretical concern. Documented cases of algorithmic error[36], data corruption, and unauthorized modification have resulted in denied loans, erroneous deportations, wrongful benefit terminations, and misdiagnoses [37] . The integrity of the data environment is, in these cases, a precondition for the exercise of fundamental rights. In each case, the harm flows not from data use, but from the inability to verify, correct, or challenge corrupted records.

A rights-based approach to data integrity would extend beyond error prevention to encompass meaningful consent, granular access control, and enforceable rights to correction and explanation.

---

[33] Retrieved from https://academic.oup.com/poq/article/88/SI/472/7712444.
[34] Retrieved from https://www.oecd.org/en/publications/facts-not-fakes-tackling-disinformation-strengthening-information-integrity_d909ff7a-en.html.
[35] Retrieved from https://openscholarship.wustl.edu/law_lawreview/vol85/iss6/2/.
[36] https://oecd.ai/en/catalogue/tools/algorithmic-accountability-policy-toolkit.
[37] Retrieved from https://www.hrw.org/report/2023/06/13/automated-neglect/how-world-banks-push-allocate-cash-assistance-using-algorithms.

CC BY-NC 4.0

The technical infrastructure to support such rights, including cryptographic attestation, audit trails, verified provenance, already exists. The governance frameworks to mandate and operationalize it through the full data lifecycle do not yet meet the challenge. Digital dignity therefore depends not only on privacy law, but on enforceable integrity guarantees embedded into technical and institutional design.

THE TECHNICAL TOOLKIT: STATE OF THE ART AND GOVERNANCE GAPS

Data integrity does not emerge from a single technology or policy choice. It is built in layers, across the full lifecycle of data, from the moment it is created to the moment it is destroyed. This toolkit organizes those layers into nine functional domains, each outlining its purpose and function, as well as listing specific controls, technical features, and useful sources and references to give practical guidance on implementation.

This toolkit, and its listed "tools," are designed to help a variety of stakeholders, including technical practitioners and product developers who need specific standards and tools, and organizational leaders who need to understand what they are investing in and why.

Domain 1: Data Preparation, Minimization, and Digitization

Initial collection and ingestion are the origin points for data integrity. Data that enters a system without structure, without provenance documentation, and without validation is effectively ungoverned from the start. Downstream controls will find it difficult or impossible to fully compensate for failures to consider governance and provenance from the beginning of the lifecycle. The first domain of the toolkit addresses how data is created, captured, and initially managed.

A core principle in responsible governance and addressing concerns like privacy is minimization[38]: data that is not collected cannot be corrupted, stolen, misattributed, or misused. Organizations that treat data collection as inherently valuable (i.e., accumulating everything on the theory that it might be useful) systematically create integrity problems they then spend enormous resources trying to manage, and on data that may not even be of utility. A structured assessment of what data is genuinely necessary, what format it should be in, and how it will be governed from the point of creation is the single highest-leverage intervention available.

For analog-to-digital conversion workflows (digitizing paper records, historical archives, physical evidence), chain-of-custody documentation must begin at the moment of capture. Who digitized it, when, using what equipment, under what validation protocol, etc. Understanding and maintaining trust in provenance relies on these kinds of investments from the initial stages of data collection and use.

Domain 1 Tools – Data Preparation, Minimization, and Digitization

Control	Objective	Reference Standard / Tool

[38] Data Minimization Definition. Retrieved from https://epic.org/issues/consumer-privacy/data-minimization/.

CC BY-NC 4.0

Data Minimization	Reduce attack surface; collect only what is necessary	GDPR Article 5(1)(c); NIST Privacy Framework PR.DS-1
Schema Validation & Enforcement	Ensure structural integrity at ingestion	OWASP Data Validation Cheat Sheet; JSON Schema (IETF Draft)
Data Classification & Labeling	Enable risk-based controls and retention policies	NIST SP 800-60 Vol. 1 & 2; ISO/IEC 27001:2022 Annex A.5.12
Analog-to-Digital Chain of Custody	Document provenance from point of capture	FADGI Guidelines; ISO 19005 (PDF/A)
Real-Time Ingestion Validation	Detect errors and tampering at point of entry	OWASP Data Security Controls; NIST SP 800-53 SI-3, SI-10
Secure Data Destruction & Retention	Prevent unnecessary retention; support compliance	NIST SP 800-88 Rev. 1; GDPR Article 17

Domain 2: Cryptographic Integrity – Hashing, Signing, and Verification

Cryptography is the mathematical foundation of data integrity. A variety of methods and applications help to establish data integrity, such as in proving through these cryptographic means that data has not been changed (advertently or inadvertently) and supporting verification that the person on the other side of the internet is in fact whom they claim to be. For example, hash functions take any input (e.g., a document, a dataset, a software build, or a model weight file) and produce a fixed-length fingerprint. If you change a single character in the input, the fingerprint changes completely, making tampering detectable. Digital signatures extend this: they bind a hash to a verified identity, so you can establish not only that a piece of data has not changed, but who attested to its state at a given point in time.

An emerging challenge relates to the security of certain cryptographic methods in the wake of quantum computing. Current hash and signature standards are potentially vulnerable to sufficiently powerful quantum computers. The transition to post-quantum cryptographic standards is a planning imperative for any organization with long-horizon integrity requirements – such as government archives, health records, legal documentation, or financial records. NIST finalized its first post-quantum standards in 2024.[39]

Domain 2 Tools – Cryptographic Integrity: Hashing, Signing, and Verification

[39] NIST PWC, Retrieved from https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards.

CC BY-NC 4.0

Control	Objective	Reference Standard / Tool
Cryptographic Hashing	Tamper detection via file/dataset fingerprinting	FIPS 180-4 (SHA-2 family); FIPS 202 (SHA-3)
Digital Signatures	Bind integrity to verified identity	FIPS 186-5 (DSS); RFC 8017 (RSA-PKCS); RFC 8032 (EdDSA)
Cryptographic Key Management	Protect keys; prevent unauthorized use	NIST SP 800-57 Parts 1, 2, 3
Certificate Transparency	Public, auditable log of issued TLS certificates; tamper detection	IETF RFC 9162 (Certificate Transparency 2.0)
Merkle Tree Verification	Scalable integrity proofs for large/distributed datasets	RFC 9162 Merkle structure; ISO/TC 307 DLT standards
GPU-Accelerated Hashing	High-throughput integrity verification at scale	NVIDIA cuHash; NVIDIA cuPQC SDK
Post-Quantum Cryptography	Long-horizon integrity against quantum threats	NIST FIPS 203 (ML-KEM); FIPS 204 (ML-DSA); FIPS 205 (SLH-DSA) — finalized 2024
Open-Source Cryptographic Libraries	Production-grade implementation of above standards	OpenSSL; BouncyCastle; libsodium
Hash Verification Utilities	Operational hash management and comparison	HashCalc; QuickHash GUI; sha256sum (native CLI)

Domain 3: Provenance, Metadata, and Digital Attestation

While hashing can tell you a file has not been changed, provenance tells you where it came from, who made it, what tools were used, and what has happened to it since. In a world of AI-generated content, deepfakes, and synthetic media, provenance is the difference between information that can be trusted and information that merely appears trustworthy.

The Coalition for Content Provenance and Authenticity (C2PA[40]) is the leading industry standard for embedding cryptographically signed provenance metadata directly into media files. It has been adopted by Adobe, Microsoft, Google, the BBC, Reuters, and the Associated Press. A C2PA manifest travels with the file and records the complete editorial history, including whether AI tools were used in creation or modification, verified by a chain of cryptographic signatures that cannot be forged without detection.

Digital attestation extends this to institutional assertions, essentially as a signed, time-stamped statement by a verified party that certain facts about a data asset were true at a specific moment. This is the technical foundation of document notarization, software supply chain integrity, and evidence of authentication in legal proceedings.

[40] C2PA Homepage. Retrieved from https://c2pa.org/

CC BY-NC 4.0

Domain 3 Tools – Provenance, Metadata, and Digital Attestation

Control	Objective	Reference Standard / Tool
Content Provenance & Authenticity	Embed and verify cryptographic provenance in media and documents	C2PA Specification v2.0
Metadata Embedding	Attach structured provenance metadata to files at creation	ExifTool; C2PA; XMP (Adobe/ISO 16684)
Sidecar Metadata	Maintain provenance when embedding is not feasible	Dublin Core Metadata Initiative; PREMIS (Library of Congress)
Digital Attestation & Timestamping	Tamper-evident signed assertions about data state at a point in time	RFC 3161 (Trusted Timestamping); PGP/GPG; DocuSign eSignature API
Blockchain Notarization	Immutable proof-of-existence without trusted third party	ISO/TC 307 DLT Standards; OpenTimestamps
Decentralized Identifiers (DIDs)	Portable, cryptographically verifiable identity for data producers	W3C DID Core Specification
Verifiable Credentials	Machine-verifiable signed claims about subjects and datasets	W3C Verifiable Credentials Data Model 2.0
PGP / GPG Signing	Decentralized attestation; document and code signing	OpenPGP RFC 4880; GnuPG
AI Content Labeling	Trace AI-generated or AI-modified content	C2PA AI training and inference manifest; IPTC Photo Metadata Standard
Human Rights & Conflict Documentation	Cryptographic preservation of sensitive evidentiary records	Starling Lab Framework; Hala Systems methodology

Domain 4: Data Lineage, Versioning, and Lifecycle Management

Provenance documents the origin of data. Lineage documents what has happened to it since, including every transformation, join, aggregation, filter, and model inference applied to it from source to current state. Without supports data auditing, reproduction of results, and assigning accountability for errors.

This matters to organizations practically in three main contexts. In regulated industries, lineage is often a compliance requirement: financial regulators [41] and healthcare [42] oversight bodies increasingly require organizations to demonstrate that the data underlying a consequential decision can be traced back to its source. In AI development, lineage of training data[43] is the precondition for understanding model behavior, detecting bias, and auditing outputs. And in litigation [44] (civil or criminal [45]) and investigations, lineage underpins the chain of custody that determines whether evidence is admissible.

[41]Retrieved from https://www.bis.org/publ/bcbs_nl36.htm.
[42] Retrieved from https://pmc.ncbi.nlm.nih.gov/articles/PMC10384601/.
[43]Retrieved from https://airc.nist.gov/airmf-resources/playbook/.
[44] Retrieved from https://law.baylor.edu/sites/g/files/ecbvkj1546/files/2023-11/7_grimm_capra_joseph.pdf.
[45] Civil or Criminal Investigation Rules. Retrieved from https://www.law.cornell.edu/rules/fre/rule_902.

CC BY-NC 4.0

Versioning is lineage applied to the data asset itself - a record of every state the data has occupied over time, enabling rollback, comparison, and audit. Git-based version control, standard in software development, is increasingly being applied to datasets and machine learning models through tools like data version control (DVC).

Domain 4 Tools – Lineage, Versioning, and Lifecycle Management

Control	Objective	Reference Standard / Tool
Data Lineage Tracking	End-to-end traceability of data origin and transformation	OpenLineage (Linux Foundation); Apache Atlas
Provenance Data Model	Formal vocabulary for recording agents, activities, and entities	W3C PROV Data Model (PROV-DM)
Dataset Version Control	Track dataset states over time; enable rollback and audit	DVC (Data Version Control); Delta Lake (Linux Foundation)
Data Lifecycle Governance	Accountability and stewardship across full data lifecycle	ISO/IEC 38505-1; DAMA-DMBOK
Schema Registry & Enforcement	Enforce data structure consistency across pipelines	Apache Kafka Schema Registry; Confluent Schema Registry
Chain-of-Custody Controls	Legal traceability for evidentiary and forensic purposes	NIST SP 800-86; SWGDE Best Practices
Secure Audit Trails	Preserve tamper-evident log of all data access and modification	ISO/IEC 27001:2022 Annex A.8.15; W3C Linked Data Proofs
Immutable Logging	Hash-chained logs that detect any retrospective tampering	W3C RDF Dataset Canonicalization (RDFC-1.0, 2024)
Synthetic Data Governance	Label, document, and risk-manage AI-generated datasets	DARPA D3M; NIST AI RMF Appendix; IEEE 7003 (Draft)

Domain 5: Access Control, Authentication, and Secure Delegation

Integrity is not only about preventing external attackers from modifying data, but requires controlling who among legitimate users can read, write, modify, and certify data, as well as ensuring that every access event is logged and attributable. Most data integrity failures in practice are not the result of sophisticated external attacks. They are the result of excessive internal access, poorly managed credentials, and the absence of audit trails that would make unauthorized modification detectable. This vector is absolutely critical to secure – identity and access controls are one of the most common vulnerability and point of entry[46] in most cyber attacks.

Zero Trust Architecture[47] represents the current state of the art for access governance. The traditional model where organizations trust anyone inside the network perimeter while verifying those outside it has been rendered obsolete by continuous compromise of defenses and lateral movement across cloud computing environments, remote work increasing dependency on digital accesses and services, and supply chain compromises that bridge attackers inside the perimeter. Zero Trust inverts

[46] Most Common Entry Points, Retrieved from https://spycloud.com/blog/verizon-2025-data-breach-report-insights/.
[47] Zero Trust Architecture. Retrieved from https://doi.org/10.6028/NIST.SP.800-207.

CC BY-NC 4.0

the assumption: no access is trusted by default, regardless of where the request originates. Every access request is explicitly verified, every session is authorized, and every action is logged.

Domain 5 Tools – Access Control, Authentication, and Secure Delegation

Control	Objective	Reference Standard / Tool
Role-Based Access Control (RBAC)	Assign access rights to defined roles; prevent privilege creep	NIST SP 800-207; NIST RBAC Model (INCITS 359)
Attribute-Based Access Control (ABAC)	Fine-grained, context-sensitive access decisions	NIST SP 800-162; XACML 3.0 (OASIS)
Zero Trust Architecture	Eliminate implicit trust; verify all access explicitly	NIST SP 800-207; CISA Zero Trust Maturity Model v2
Multi-Factor Authentication (MFA)	Prevent unauthorized access via credential compromise	NIST SP 800-63B; FIDO2/WebAuthn
Secure Delegation & Revocation	Controlled authority transfer; timely revocation of access	OAuth 2.0 (RFC 6749); User-Managed Access (UMA 2.0)
Identity & Access Management (IAM)	Centralized identity governance across systems	NIST SP 800-63 Digital Identity Guidelines series
API Security	Protect data accessed via APIs; prevent unauthorized use	OWASP API Security Top 10; NIST SP 800-204
Privileged Access Management (PAM)	Control and monitor high-privilege accounts	CIS Control 5 (Account Management); NIST SP 800-53 AC-2, AC-6
Container & Cloud Integrity	Integrity controls for cloud-native and containerized environments	NIST SP 800-190; CIS Benchmarks (cloud-specific)

Domain 6: Privacy-Preserving Verification and Consent

A central tension in data integrity is that verifying data often requires disclosing it. Auditors need to see records, regulators need to verify compliance, partners need to confirm dataset properties before relying on them, etc. In each case, the act of verification creates an exposure where the data being verified can be seen, copied, or misused by the verifying party.

Privacy-preserving cryptographic techniques[48] can resolve this tension. Zero-knowledge proofs allow one party to prove to another that a statement about data is true without revealing the data itself. Homomorphic encryption allows computation to be performed on encrypted data so that the analyst never sees the underlying records. Secure multi-party computation allows multiple parties to jointly derive insights from their combined data without any party revealing its inputs to the others.

These are not theoretical capabilities, they are deployed in production at scale in financial services in areas like fraud detection across competing institutions, in healthcare such as for clinical research across hospital networks, and national security for intelligence sharing without source disclosure.

[48] Privacy-preserving cryptographic techniques. Retrieved from https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl-understanding-pets-and-ppts-dec2023.pdf.

CC BY-NC 4.0

The barrier to broader adoption is primarily awareness and governance clarity[49], not technical readiness.

Domain 6 Tools – Privacy-Preserving Verification and Consent

Control	Objective	Reference Standard / Tool
Zero-Knowledge Proofs (ZKPs)	Prove data properties without revealing underlying data	zkSNARKs / zkSTARKs; ZKProof Community Reference
Homomorphic Encryption	Compute on encrypted data without decryption	Microsoft SEAL; OpenFHE; ISO/IEC 18033-6 (in development)
Secure Multi-Party Computation (SMPC)	Collaborative analytics without data exposure	ISO/IEC 19592-1; ITIF PETs report; MP-SPDZ library
Federated Learning	Train AI models across distributed data without centralizing records	TensorFlow Federated; PySyft (OpenMined); Google FL framework
Differential Privacy	Provable privacy guarantees for analytics outputs	NIST SP 800-188; Apple DP framework; Google DP library
Data Anonymization & Pseudonymization	GDPR/HIPAA compliance; reduce re-identification risk	ISO/IEC 20889; ISO/IEC 27559
Consent Management	Automated, auditable, revisable user consent and rights	OneTrust; Secure Privacy; W3C Data Privacy Vocabulary (DPV)
Privacy Framework	Risk-based organizational privacy governance	NIST Privacy Framework v1.0
Cross-Border Data Transfer Controls	Manage legal and technical risks of international data flows	GDPR SCCs; APEC CBPR; EU-US Data Privacy Framework

Domain 7: Data Quality, Anomaly Detection, and AI Integrity

A data asset can be cryptographically intact, unmodified since it was created, and still be wrong. Data quality controls address accuracy, completeness, consistency, and fitness for purpose. Anomaly detection addresses drift in the quality and utility of the data for its intended purpose. This can derive from many causes, such as the gradual degradation of data quality over time as the world changes but the dataset does not, or the sudden introduction of anomalous records that may indicate manipulation or system failure.

For AI systems specifically, data quality is a cornerstone determinant of model behavior. A model trained on biased, incomplete, or manipulated data will produce biased, incomplete, or manipulated outputs, potentially with no visible indication that anything is wrong. The NIST AI Risk Management Framework[50] and IEEE 7003[51] provide the emerging governance standards for AI data quality; and automated pipeline validation tools can provide the operational implementation.

Domain 7 Tools – Data Quality, Anomaly Detection, and AI Integrity

[49]OECD. (2023). Emerging Privacy-Enhancing Technologies: Current Regulatory and Policy Approaches. OECD Digital Economy Papers No. 351. https://doi.org/10.1787/bf121be4-en

[50] NIST AI Risk Management Framework. Retrieved from https://www.nist.gov/itl/ai-risk-management-framework

[51] IEEE 7003. Retrieved from https://sagroups.ieee.org/7003/

CC BY-NC 4.0

Control	Objective	Reference Standard / Tool
Data Quality Metrics	Quantify accuracy, completeness, consistency, and timeliness	ISO/IEC 25012 (Data Quality Model); DAMA Data Quality Framework
Automated Pipeline Validation	Real-time data quality checks at every pipeline stage	Great Expectations; Apache Griffin; dbt tests
Anomaly & Drift Detection	Identify inconsistent, degraded, or manipulated data	NIST AI RMF (Govern 1.2, Map 2.3); Evidently AI; Whylogs
AI Training Data Governance	Document, version, and audit AI training corpora	NIST AI RMF; Hugging Face Dataset Cards; Croissant metadata standard
Bias Detection & Measurement	Identify and quantify bias in datasets and model outputs	IEEE 7003 (Draft); IBM AI Fairness 360; Google What-If Tool
Model Cards & Datasheets	Standardized documentation of model and dataset properties	Google Model Cards; Datasheets for Datasets
Secure Data Ingestion (Anti-Poisoning)	Prevent data poisoning attacks on AI training pipelines	OWASP Machine Learning Security Top 10; NIST SP 800-53 SI-10
Synthetic Data Labeling	Distinguish AI-generated data from empirically derived records	C2PA AI manifest; DARPA D3M; IEEE 7003

Domain 8: Archiving, Preservation, and Immutable Records

Most organizations think about data integrity in operational terms around the integrity of data in active use. Long-term preservation raises a different and harder set of questions. Cryptographic standards evolve and must be migrated before they are broken. Storage media degrades and file formats become unreadable. Organizations dissolve, and the custodial chain that maintains integrity guarantees breaks with them.

Institutions that have thought most seriously about this problem, such as national libraries, archival institutions, human rights organizations, have developed frameworks worth studying. The USC Shoah Foundation's collaboration with Starling Lab[52] represents the current state of the art for high-stakes decentralized[53] preservation: Holocaust survivor testimony secured through cryptographic attestation across a distributed network, designed to remain verifiable and accessible regardless of what happens to any single custodial institution. Additional case studies are included in the Appendix.

For organizations with regulatory, legal, or historical obligations that extend long-term, this domain deserves dedicated attention and consideration for requirements around integrity, privacy, and discoverability of the data that is stored.

Domain 8 Tools – Archiving, Preservation, and Immutable Records

Control	Objective	Reference Standard / Tool
Long-Term Digital Preservation	Format-resilient, institutionally sustainable archival storage	NDSA Levels of Preservation; ISO 14721 (OAIS Reference Model)

[52] USC Shoah Foundation's collaboration with Starling Lab. Retrieved from https://vha.usc.edu/home
[53] Decentralized Human Rights Verification. Retrieved from https://venturebeat.com/consumer/usc-and-stanford-launch-starling-lab-to-protect-human-rights-with-decentralization.

CC BY-NC 4.0

Content-Addressed Storage	Hash-based storage that makes integrity self-verifying	IPFS (InterPlanetary File System); Filecoin
Decentralized Archival	Cryptographically secured preservation without single custodian	Starling Lab Framework; Filecoin + IPFS integration
Immutable Records via DLT	Tamper-evident, independently verifiable records at scale	ISO/TC 307 DLT Standards; Hyperledger Fabric; OpenTimestamps
Long-Term Institutional Memory	10,000-year planning horizon for format and media resilience	Long Now Foundation; Internet Archive; Library of Congress
Archival Metadata Standards	Ensure long-term discoverability and context preservation	PREMIS; EAD (Encoded Archival Description); BagIt (RFC 8493)
Secure Backup & Redundancy	Protect against ransomware, media failure, and site loss	NIST SP 800-34 Rev. 1 (Contingency Planning); 3-2-1 backup principle
PDF/A Archival Format	Long-term readable document format with embedded metadata	ISO 19005 (PDF/A); PDF/A-3 for embedded attachments

Domain 9: Regulatory Compliance and Governance Frameworks

Technical controls without governance function as infrastructure without management. Governance frameworks provide the organizational structures, accountability mechanisms, and documented policies through which technical controls are deployed and maintained. Regulatory frameworks provide the legal mandates that require organizations to invest in data integrity in the first place, and the enforcement mechanisms that make those mandates real.

The current regulatory landscape is fragmenting in ways that create both complexity and opportunity. The EU Data Act and AI Act represent the most comprehensive attempt to date to codify data integrity requirements across an entire economy. The United States, by contrast, has sector-specific regimes without a unifying cross-sector framework[54]. Organizations operating internationally must navigate this fragmented terrain, even as policymakers consider means of enabling more streamlined portability and interoperability across frameworks and jurisdictions.

Domain 9 Tools – Regulatory Compliance and Governance Frameworks

Control	Objective	Reference Standard / Tool
GDPR	Data accuracy, subject rights, breach notification, cross-border controls	EU General Data Protection Regulation (2016/679)
EU Data Act	Data sharing obligations, interoperability, portability, quality standards	EU Data Act (2023/2854)
EU AI Act	Training data quality and provenance documentation for high-risk AI	EU AI Act (2024/1689)
CCPA/CPRA	Consumer data rights; accuracy; opt-out of sharing	California Consumer Privacy Act/Privacy Rights Act
HIPAA Security Rule	Integrity controls for protected health information	45 CFR Parts 160 and 164

[54]Retrieved from https://www.dlapiperdataprotection.com/

CC BY-NC 4.0

Sarbanes-Oxley (SOX)	Financial record integrity; audit trail requirements	SOX Sections 302, 404, 802
FedRAMP	Security authorization for federal cloud services including data integrity	FedRAMP Authorization Framework
ISO/IEC 27001:2022	Information security management system; comprehensive control framework	ISO/IEC 27001:2022
ISO/IEC 27701	Privacy information management; extension of 27001 to personal data	ISO/IEC 27701:2019
ISO/IEC 38505	Data governance accountability and stewardship	ISO/IEC 38505-1:2017
NIST Cybersecurity Framework 2.0	Identify, Protect, Detect, Respond, Recover – updated 2024	NIST CSF 2.0
NIST AI RMF	Risk management for AI systems including training data quality	NIST AI RMF 1.0 (2023)
DISA STIGs	Security Technical Implementation Guides for federal systems	DISA STIG Library
GLBA/FTC Safeguards Rule	Financial data confidentiality and integrity; written security program; breach notification for 500+ consumer incidents	15 U.S.C. §§ 6801-6827; FTC Safeguards Rule, 16 CFR Part 314
RFPA – Right to Financial Privacy Act	Restricts federal agency access to bank records; requires legal process, customer notice, and audit trail of all government disclosures	12 U.S.C. §§ 3401–3422
FERPA	Student record accuracy and confidentiality; right to inspect, correct, and control disclosure; applies to federally funded institutions	20 U.S.C. § 1232g; 34 CFR Part 99
Telecom Act §222 / FCC CPNI Rules	Carrier protection of call records, usage, and location data; opt-in consent for sharing; breach notification to FCC/FBI within 7 days (revised 2023)	47 U.S.C. § 222; FCC CPNI Rules, 47 CFR Part 64 Subpart U
CALEA	Lawful intercept capability requirements for carriers; System Security and Integrity Plans; secure records of all interceptions	47 U.S.C. §§ 1001-1010; 47 CFR Part 1 Subpart Z
ECPA	Integrity and confidentiality of electronic communications in transit and storage; governs wiretapping, stored data access, and metadata collection	18 U.S.C. §§ 2510-2523; §§ 2701-2712; §§ 3121-3127
FCRA – Fair Credit Reporting Act	Accuracy and integrity of consumer credit records; mandatory dispute and correction mechanisms; permissible use controls; FTC/CFPB enforcement	15 U.S.C. § 1681 et seq.
COPPA	Security and integrity of children's personal data; verifiable parental consent; right to review, correct, and delete; updated January 2025	15 U.S.C. §§ 6501-6506; COPPA Rule, 16 CFR Part 312
FISMA	Federal agency information security programs; mandates NIST SP 800-53 and FIPS 199/200 compliance; continuous monitoring and incident reporting to CISA	44 U.S.C. §§ 3551-3558; OMB Circular A-130

CC BY-NC 4.0

CONCLUSION

In a world increasingly defined by digital interactions, data integrity is a societal necessity. In a world where digital spaces are increasingly navigated by AI agents and their generated data, it cannot wait any longer. Data integrity underpins trust in institutions, fuels economic growth, and safeguards democratic freedoms. As generative AI, agentic systems, and synthetic data reshape the digital landscape, the need for robust, scalable, and interoperable data integrity frameworks has never been more urgent. This toolkit provides a practical, actionable set of tools for organizations of all sizes to implement data integrity controls, ensuring that their data remains accurate, reliable, and secure across its lifecycle.

By adopting the principles and tools outlined here, organizations can not only mitigate risks but also unlock new opportunities for innovation, efficiency, and collaboration. From cryptographic hashing and provenance tracking to privacy-preserving analytics and regulatory compliance, the toolkit equips stakeholders with the resources needed to navigate the complexities of the modern data ecosystem. Whether addressing national security concerns, combating misinformation, or enabling agentic commerce, data integrity is the quiet infrastructure that makes it all possible.

The stakes are high, but so is the potential for impact. By investing in data integrity, we can build a digital future that is resilient and efficient while also just and equitable. This is not just about protecting data; it is about protecting the trust, dignity, and agency of individuals and institutions alike. The time to act is now. Together, we can ensure that the digital foundations of our economies, democracies, and societies remain strong for decades to come.

CC BY-NC 4.0

Appendix: Data Integrity Tool Example Implementations and Use Cases

Domain	Tool / Standard	Example Implementations	Example Use Case
Domain 1: Data Preparation, Minimization & Digitization	NIST SP 800-60 (Data Classification)	Microsoft Purview Information Protection; Varonis; Titus Classification	• A defense contractor deploys a classification tool (Microsoft Purview) that asks employees to label documents when created. The system automatically applies protections: lower-sensitivity documents flow freely, while higher-sensitivity ones are encrypted, restricted to named individuals, and blocked from external email. • A government agency requires every document to carry a visible classification marking. A prompting tool (Titus) asks the author to select a classification each time they save a file. That classification travels with the document as metadata, allowing downstream systems to enforce access controls automatically.
	NIST SP 800-88 (Media Sanitization)	Blancco; DBAN; AWS S3 Object Expiry + Glacier Vault Lock; BitRaser	• A company disposing of several hundred old laptops uses specialized software (Blancco) to overwrite every drive multiple times in a pattern that meets NIST 800-88 standards, receiving a certificate for each drive confirming sanitization was completed. • Under GDPR, a cloud-based company meets deletion obligations by deleting the encryption key rather than hunting down every copy of the data. Once the key is gone, the data is mathematically unrecoverable; an approach NIST SP 800-88 explicitly recognizes as valid.

CC BY-NC 4.0

Domain 2: Cryptographic Integrity – Hashing, Signing & Verification	NIST SP 800-57 (Key Management)	AWS KMS; HashiCorp Vault; Azure Key Vault; PKCS#11-compliant HSMs (Thales, nCipher)	• A hospital uses HashiCorp Vault to centrally generate, rotate, and expire encryption passwords on a schedule — following NIST 800-57 guidelines for how long a given type of key should remain in use before being replaced. • A federal contractor stores signing credentials inside a dedicated HSM that is physically designed to prevent credentials from ever being copied or extracted, in line with NIST 800-57 Part 2 rules for high-assurance key storage.
	FIPS 180-4 (SHA Hash Standard)	OpenSSL; Python hashlib; Java MessageDigest; BouncyCastle	• A software company publishes a SHA-256 hash "fingerprint" alongside each product release. Anyone who downloads the file can run the same algorithm and compare fingerprints — if they match, the file is genuine and unaltered. • A digital forensics investigator immediately generates a hash of a disk image now of collection. If anyone later questions whether the evidence was modified, the investigator can re-hash the image and show the fingerprint is identical.
	RFC 8017 (PKCS#1 / RSA)	OpenSSL; Java JCE; BouncyCastle; cryptography (Python); Go crypto/rsa	• A contract management system uses RSA cryptography, so each employee signature produces a unique mathematical stamp tied to their identity. Anyone reviewing the document later can verify the stamp is authentic, and the text has not changed. • An organization's internal web services use RSA for a secure "handshake" at the start of each connection: each service proves its identity by demonstrating it holds a private credential that only it could possess.

CC BY-NC 4.0

	IETF RFC 9162 (Certificate Transparency 2.0)	Google Trillian; ctclient (Go); Let's Encrypt; Cloudflare Nimbus	• When a certificate authority issues a TLS certificate, it is required to record it in a public, append-only log. Anyone can audit the log to check whether a certificate was issued for a domain without the owner's knowledge. • A large company monitors public certificate transparency logs for entries relating to its own domain names. If an unauthorized certificate appears, an automated alert fires immediately, allowing investigation and revocation before interception.
	PGP / OpenPGP (RFC 4880)	GnuPG (gpg); Proton Mail; Sequoia-PGP (Rust); OpenKeychain (Android)	• A cybersecurity team digitally signs threat intelligence reports using PGP before sending, so recipients can verify the report genuinely came from this team and has not been altered in transit. • A journalist publishes their PGP public key on their website so that sources can encrypt messages that only the journalist holding the corresponding private key can decrypt.
Domain 3: Provenance, Metadata & Digital Attestation	W3C Linked Data Proofs & RDFC-1.0	jsonld-signatures (Digital Bazaar); rdf-canonize (npm); jsonld.js; DIF did-jwt-vc	• A university issues digital diplomas that graduates can store in a digital wallet on their phone. The diploma contains a cryptographic seal applied at the time of issue. When a graduate shares it with an employer, the employer's system can verify the seal is intact and the document is genuine, without needing to call the university's registrar.

CC BY-NC 4.0

Domain 4: Data Lineage, Versioning & Lifecycle Management	W3C PROV / OpenLineage	Apache Atlas; Marquez; dbt; Apache Airflow with OpenLineage plugin; Great Expectations	• A bank produces a daily regulatory report. When a regulator questions a number, the team uses a lineage tracking tool to trace it backwards through every transformation to identify exactly which source data table it came from, and which processing step introduced an error. • A data science team records provenance using W3C PROV-O so that when an auditor asks, "how was this model built and on what data?", the answer is structured, consistent, and machine-readable rather than buried in emails.
	Blockchain / DLT (ISO/TC 307)	Hyperledger Fabric; Ethereum; R3 Corda; DAML	• Pharmaceutical manufacturers, distributors, and pharmacies use a shared digital ledger to record every time a batch of medication changes hands. It becomes straightforward to trace a recalled drug batch back through the entire supply chain and identify where a counterfeit product entered. • In international trade, a shared digital ledger lets all parties see the same document in real time. When the exporter presents shipping documents, a smart contract automatically triggers payment, reducing settlement time from days to hours.
	ISO/IEC 27001:2022 Annex A.8.15 (Logging)	Splunk; Elastic SIEM; Graylog; Azure Monitor; AWS CloudTrail	• A fintech company feeds activity records into Splunk, configured to retain records for 12 months and flag gaps or anomalies. When auditors arrive, the company pulls a report demonstrating continuous logging coverage aligned to ISO 27001 Annex A.8.15. • A security services firm uses Graylog to collect and store activity logs from all client environments, applying tamper-detection

CC BY-NC 4.0

			controls so that if any log records are altered after the fact, an alert fires.
	ISO/IEC 38505 (Data Governance)	Collibra; Alation; Informatica Data Governance; Apache Atlas	• A multinational bank uses Collibra to build a central catalogue of all data assets, assigning a named owner to each one who is accountable for accuracy and appropriate use, following the ISO 38505 principle that governing boards should evaluate, direct, and monitor data management. • A hospital network applies Apache Atlas to scan datasets and flag any that contain health information, enforcing access restrictions mandated by the governance council before any analyst can query those records.
	E-Discovery & Forensics (EDRM, NIST SP 800-101, RFC 3227)	Relativity; Nuix; Cellebrite; Autopsy (OSS); FTK (AccessData); EnCase	• A law firm managing a securities lawsuit uses Relativity to learn from human-reviewed examples and automatically prioritize millions of documents, following the EDRM workflow, a standard, legally defensible framework for electronic evidence. • An incident response team takes a forensic copy of a departing employee's devices, capturing the most volatile evidence first (memory) before moving to less volatile storage, documented to meet admissibility standards for tribunal or criminal proceedings.

CC BY-NC 4.0

Domain 5: Access Control, Authentication & Secure Delegation	OAuth 2.0 / UMA 2.0 (User-Managed Access)	Keycloak; Auth0; Okta; Ping Identity; ORY Hydra; Spring Security OAuth	• A patient health portal uses UMA 2.0 to let patients decide exactly who can see their medical records and for how long, and revoke access at any time without calling IT. The patient remains in control. • When you click "Sign in with Google," Google issues a temporary, limited-permission token that lets the app do only what you explicitly approved (e.g., reading your calendar), and nothing else.
	NIST RBAC / ABAC Models	Casbin; Open Policy Agent (OPA); AWS IAM; XACML (OASIS); Apache Ranger	• A hospital assigns every employee a job role in the HR system. The access control system grants permissions based on those roles — a billing clerk has access to billing records but not patient clinical notes, simply because of the role they hold. • A large data platform uses a policy engine that evaluates multiple conditions together in real time (department, sensitivity level, time of day) each time someone requests access, rather than maintaining thousands of individual rules.
	OWASP Data Security Controls	OWASP Dependency-Check; OWASP ZAP; ModSecurity WAF; SonarQube	• A software development team automatically checks every new release for known security weaknesses. If a critical vulnerability is detected, the release is automatically blocked until the issue is resolved. • An online retailer uses the OWASP ASVS checklist for its annual penetration test and runs a WAF in front of the checkout page that automatically blocks common attack patterns such as code injection into form fields.

CC BY-NC 4.0

	OWASP API Security / NIST SP 800-204	Kong Gateway; AWS API Gateway + WAF; 42Crunch; Apigee	• A healthcare interoperability company reviews API design specs against the OWASP API Security Top 10 before any code is written. Catching problems at the design stage — such as an endpoint returning more patient data than needed — is far cheaper than fixing them after deployment. • Following NIST SP 800-204, a company requires every service-to-service connection to be mutually authenticated and encrypted, with access decisions enforced by a central policy engine — so a compromised service cannot impersonate another.
	NIST SP 800-190 / CIS Benchmarks (Container Security)	Aqua Security; Prisma Cloud; Trivy; Falco; kube-bench	• A platform team runs automated checks against the CIS Kubernetes Benchmark to verify container infrastructure is configured safely, tracking pass/fail on a dashboard and assigning remediation tasks. • A security team scans every container image for vulnerabilities before production deployment and runs a monitoring tool that watches for suspicious runtime behavior, such as a container unexpectedly opening a command-line shell.

CC BY-NC 4.0

Domain 6: Privacy-Preserving Verification & Consent	Zero-Knowledge Proofs (ZKPs)	libsnark; bellman (Zcash); circom/snarkjs; gnark (Go); StarkWare Cairo VM	• A user proves they are over 18 to access an age-restricted website without handing over their date of birth or any personal detail. Their digital identity wallet generates a mathematical proof — the website receives only the yes/no answer. • A bank proves to a regulator that its total assets exceed its liabilities — solvency — without disclosing individual account balances, customer names, or counterparties. Anyone can verify the proof is correct while financial details remain private.
	Homomorphic Encryption (HE)	Microsoft SEAL; OpenFHE; PALISADE; Zama's concrete	• Five hospitals collaborate on cancer research by sending encrypted patient data to a shared analysis environment. The analysis produces results (e.g., average treatment outcomes by age group) without the system ever seeing any individual patient record. • A company runs queries against encrypted financial data stored with a cloud provider. The query runs on encrypted data and returns an encrypted result. The cloud provider then performs useful computation while remaining blind to the underlying values.
	Multi-Party Computation (MPC)	MP-SPDZ; SCALE-MAMBA; Sharemind;	• Three competing supermarket chains use MPC to jointly compute industry-wide fraud statistics. No retailer ever sees a competitor's raw transaction records; they each receive only the

CC BY-NC 4.0

		Inpher; Duality Technologies	final aggregated insights. • A government statistics office publishes a national gender pay gap analysis drawing payroll data from hundreds of private employers via MPC. Each employer data stays private; the government receives only the final statistical output.
	ISO/IEC 18033-6 (HE) & ISO/IEC 19592-1 (Secret Sharing)	Shamir Secret Sharing (ssss CLI, tss-lib in Go); OpenFHE	• A cryptocurrency exchange splits a master key into five pieces held by custodians in different countries using secret sharing. No single custodian can act alone: any three of the five must come together to reconstruct the key and authorize a withdrawal. • A cloud storage provider offers encrypted search: the query arrives already encrypted, the server finds matching results while data remains locked, and only the customer can decrypt what comes back.
	ISO/IEC 20889 & ISO/IEC 27559 (PETs Framework)	ISO 27559 evaluation criteria for PETs vendor assessments	• A procurement team evaluating three analytics vendors that all claim to be “privacy-preserving” uses ISO 27559 criteria to score each vendor consistently and justify their selection decision to a governance committee.
	NIST Privacy Framework	OneTrust; TrustArc; Privitar	• A consumer tech company uses the NIST Privacy Framework to build a comprehensive inventory of all the personal data it collects, feeding it into both GDPR privacy notices and CCPA consumer rights processes from a single underlying record. • A state government agency uses the framework as a common language across departments handling benefits, taxation, health, and licensing data — so each team evaluates privacy practices in a comparable, rollup-ready way.

CC BY-NC 4.0

	ISO/IEC 27701 (PIMS)	OneTrust; SAP GRC; IBM OpenPages	• A cloud service provider extends its ISO 27001 certification to include ISO 27701, producing a single audit that covers both security and privacy — satisfying enterprise customers' due diligence requirements for both topics. • A payroll processor uses ISO 27701's separate obligation sets for data controllers (own employees) and data processors (clients' employees), managing both independently with different control owners, audit trails, and evidence packs.
	GDPR / CCPA / SCCs / BCRs	OneTrust; Osano; Transcend; Didomi	• A US-based software company with EU customers uses Standard Contractual Clauses as the legal mechanism for cross-border data transfer, tracking which customers have signed them and maintaining supporting risk assessments. • A global corporation obtains Binding Corporate Rules: a single set of internal privacy commitments recognized by European regulators as a valid legal basis for employee data transfers across its offices in 30 countries.
Domain 7: Data Quality, Anomaly Detection & AI Integrity	IEEE 7003 (Algorithmic Bias / Data Quality) – Draft	IBM AI Fairness 360; Google What-If Tool; Aequitas	• A bank tests whether its loan approval algorithm approves or denies applicants at different rates based on race or gender. If disparities are found, the model is adjusted before deployment and test results are shared with regulators as evidence of due diligence. • A company selling AI-assisted hiring tools runs regular bias audits and provides plain-language reports to each customer allowing HR departments to make informed decisions about whether the tool meets their diversity standards.

CC BY-NC 4.0

	NIST AI Risk Management Framework (AI RMF)	NIST AI RMF Playbook; IBM OpenPages; ServiceNow GRC; Microsoft Responsible AI Toolbox	• Program managers use the AI RMF, adapted to their operational context, to evaluate AI vendors and use cases with a clear-eyed view of the risks before any technology is procured.
	DARPA D3M / Synthetic Data Standards	Gretel.ai; Mostly AI; Synthetic Data Vault (sdv); CTGAN; ydata-synthetic	• A bank creates thousands of realistic but entirely artificial fraudulent transactions using synthetic data generation, giving a fraud detection model enough examples to learn from without exposing actual customer records. • A government agency generates a synthetic version of a restricted population survey preserves statistical relationships but contains no real individual's personal health information and makes it publicly available for academic research.
Domain 8: Archiving, Preservation & Immutable Records	Blockchain / DLT (ISO/TC 307) – Archival Applications	Hyperledger Fabric; Ethereum; R3 Corda; DAML	• Pharmaceutical supply chain participants use a shared ledger so that no single participant can alter records unilaterally, enabling precise tracing of recalled drug batches back through the entire chain. • Trade finance parties use a shared ledger with smart contracts to automate payment on presentation of shipping documents, creating an immutable record of the authoritative version of each document.
Domain 9: Regulatory Compliance & Governance Frameworks	GDPR / CCPA / SCCs / BCRs – Compliance Operations	OneTrust; Osano; Transcend; Didomi	• Standard Contractual Clauses provide the legal mechanism for a US company processing EU customer data on US servers, with a compliance platform tracking signatures and maintaining risk assessments. • Binding Corporate Rules allow a global

CC BY-NC 4.0

			corporation to move employee data between 30 countries under a single set of internal privacy commitments recognized by European regulators.
	ISO/IEC 27701 (PIMS) – Audit & Certification	OneTrust; SAP GRC; IBM OpenPages	• Extending ISO 27001 to ISO 27701 produces a single audit covering both security and privacy, reducing assessment burden for enterprise customers. • A payroll processor manages separate controller and processor obligations with distinct control owners and evidence packs for each role.
	ISO/IEC 38505 (Data Governance) – Board-Level Accountability	Collibra; Alation; Informatica Data Governance; Apache Atlas	• A multinational bank uses Collibra to catalogue all data assets and assign named owners accountable for accuracy, following the ISO 38505 principle that governing boards should evaluate, direct, and monitor data management. • A hospital network uses Apache Atlas to scan and flag datasets containing health information, enforcing governance council restrictions before any analyst can query the records.

References:

1. "AI Now Writes as Many Online Articles as Humans Do," Graphite Research, May 15, 2026, https://graphite.io/five-percent/ai-now-writes-as-many-online-articles-as-humans-do
2. "Bot vs. Human Web Traffic," Cloudflare, April 2026, https://radar.cloudflare.com/traffic#bot-vs-human.

CC BY-NC 4.0

3. Gartner. (n.d.). *Data quality: Why it matters and how to achieve it* (citing Gartner research from 2020). https://www.gartner.com/en/data-analytics/topics/data-quality
4. e-Estonia / Estonian Information System Authority (RIA). (n.d.). *Interoperability services: X-Road* (self-reported estimate of 1,345 working years saved annually). https://e-estonia.com/solutions/interoperability-services/x-road/
5. Forrester Consulting. (2025). *The Total Economic Impact of Microsoft Purview* (commissioned study; 355% composite ROI over three years across data security, governance, and compliance). https://tei.forrester.com/go/microsoft/purview/
6. International Organization for Standardization. (n.d.). *ISO/IEC JTC 1/SC 42: Artificial intelligence*. https://www.iso.org/committee/6794475.html
7. Europol Innovation Lab. (2022). *Facing reality? Law enforcement and the challenge of deepfakes*. Europol. https://www.europol.europa.eu/cms/sites/default/files/documents/Europol_Innovation_Lab_Facing_Reality_Law_Enforcement_And_The_Challenge_Of_Deepfakes.pdf
8. National Center for State Courts. (n.d.). *AI-generated evidence: Threat to public trust in courts*. https://www.ncsc.org/resources-courts/ai-generated-evidence-threat-public-trust-courts
9. Commodity Futures Trading Commission. (2024). *Technology Advisory Committee report on artificial intelligence*. https://www.cftc.gov/media/10626/TAC_AIReport050224/download
10. Atlantic Council GeoTech Center. (n.d.). *Why data governance matters*. Atlantic Council. https://www.atlanticcouncil.org/blogs/geotech-cues/why-data-governance-matters/
11. SAP Community. (n.d.). *Bad data costs the U.S. $3 trillion per year*. https://community.sap.com/t5/technology-blog-posts-by-sap/bad-data-costs-the-u-s-3-trillion-per-year/ba-p/13575387
12. SSRN. (2025). *Cognitive surrender*. Social Science Research Network. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6097646
13. e-Estonia. (n.d.). *X-Road: Interoperability services*. https://e-estonia.com/solutions/interoperability-services/x-road/
14. Forrester Consulting. (n.d.). *The Total Economic Impact™ of Snowflake AI Data Cloud*. Commissioned by Snowflake. https://tei.forrester.com/go/Snowflake/AIDataCloud/?lang=en-us
15. Taylor & Francis. (2025). *Information Systems Management*. https://www.tandfonline.com/doi/full/10.1080/10580530.2025.2477459

CC BY-NC 4.0

16. ITIC. (2024). *2024 hourly cost of downtime survey*. Information Technology Intelligence Consulting. https://itic-corp.com/itic-2024-hourly-cost-of-downtime-report/
17. Organisation for Economic Co-operation and Development. (2019). *Economic and social benefits of data access and sharing*. In *Enhancing access to and sharing of data*. OECD Publishing. https://www.oecd.org/en/publications/2019/11/enhancing-access-to-and-sharing-of-data_070835df/full-report/economic-and-social-benefits-of-data-access-and-sharing_836734cb.html
18. Amazon Web Services & Forrester Consulting. (n.d.). *Forrester study examines data integration ROI for public sector organizations*. AWS Public Sector Blog. https://aws.amazon.com/blogs/publicsector/forrester-study-examines-data-integration-roi-for-public-sector-organizations/
19. European Commission. (2025). *European Data Act*. Digital Strategy. https://digital-strategy.ec.europa.eu/en/policies/data-act
20. European Parliament & Council of the European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act)*. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
21. Hoover Institution. (n.d.). *China's grand strategy for global data dominance*. Stanford University. https://www.hoover.org/research/chinas-grand-strategy-global-data-dominance
22. National Endowment for Democracy. (n.d.). *Data-centric authoritarianism: How China's development of frontier technologies could globalize repression*. https://www.ned.org/data-centric-authoritarianism-how-chinas-development-of-frontier-technologies-could-globalize-repression-2/
23. DLA Piper. (n.d.). *Data protection laws of the world*. https://www.dlapiperdataprotection.com/?c=US
24. Oxford Academic. (2024). Disinformation and public opinion. *Public Opinion Quarterly, 88*(SI), 472. https://academic.oup.com/poq/article/88/SI/472/7712444
25. Organisation for Economic Co-operation and Development. (n.d.). *Facts not fakes: Tackling disinformation, strengthening information integrity*. OECD Publishing. https://www.oecd.org/en/publications/facts-not-fakes-tackling-disinformation-strengthening-information-integrity_d909ff7a-en.html
26. Washington University in St. Louis. (n.d.). Due process implications. *Washington University Law Review, 85*(6). https://openscholarship.wustl.edu/law_lawreview/vol85/iss6/2/
27. OECD AI Policy Observatory. (n.d.). *Algorithmic accountability policy toolkit*. https://oecd.ai/en/catalogue/tools/algorithmic-accountability-policy-toolkit

CC BY-NC 4.0

28. Human Rights Watch. (2023, June 13). *Automated neglect: How the World Bank's push to allocate cash assistance using algorithms harms the poor*. https://www.hrw.org/report/2023/06/13/automated-neglect/how-world-banks-push-allocate-cash-assistance-using-algorithms
29. Electronic Privacy Information Center. (n.d.). *Data minimization*. EPIC. https://epic.org/issues/consumer-privacy/data-minimization/
30. National Institute of Standards and Technology. (2024, August). *NIST releases first 3 finalized post-quantum encryption standards*. U.S. Department of Commerce. https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards
31. Coalition for Content Provenance and Authenticity. (n.d.). *C2PA specification*. https://c2pa.org/
32. Bank for International Settlements. (n.d.). *Basel Committee newsletter no. 36*. https://www.bis.org/publ/bcbs_nl36.htm
33. PubMed Central. (2023). *Data integrity in healthcare*. National Library of Medicine. https://pmc.ncbi.nlm.nih.gov/articles/PMC10384601/
34. National Institute of Standards and Technology. (n.d.). *AI RMF playbook resources*. U.S. Department of Commerce. https://airc.nist.gov/airmf-resources/playbook/
35. Grimm, P. W., Capra, D. J., & Joseph, G. P. (2023). *Authenticating digital evidence*. Baylor Law Review. https://law.baylor.edu/sites/g/files/ecbvkj1546/files/2023-11/7_grimm_capra_joseph.pdf
36. Legal Information Institute. (n.d.). *Federal Rules of Evidence, Rule 902: Evidence that is self-authenticating*. Cornell Law School. https://www.law.cornell.edu/rules/fre/rule_902
37. SpyCloud. (2025). *Verizon 2025 data breach investigations report: Key insights*. https://spycloud.com/blog/verizon-2025-data-breach-report-insights/
38. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture* (NIST SP 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207
39. Centre for Information Policy Leadership. (2023, December). *Understanding privacy-enhancing and privacy-preserving technologies*. CIPL. https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl-understanding-pets-and-ppts-dec2023.pdf
40. Organisation for Economic Co-operation and Development. (n.d.). *Emerging privacy-enhancing technologies: Current regulatory and policy approaches*. OECD Publishing. https://www.oecd.org/en/publications/emerging-privacy-enhancing-technologies_bf121be4-en.html

CC BY-NC 4.0

41. National Institute of Standards and Technology. (n.d.). *AI risk management framework*. U.S. Department of Commerce. https://www.nist.gov/itl/ai-risk-management-framework
42. IEEE Standards Association. (n.d.). *IEEE P7003: Algorithmic bias considerations*. https://sagroups.ieee.org/7003/
43. USC Shoah Foundation. (n.d.). *Visual History Archive*. University of Southern California. https://vha.usc.edu/home
44. VentureBeat. (n.d.). *USC and Stanford launch Starling Lab to protect human rights with decentralization*. https://venturebeat.com/consumer/usc-and-stanford-launch-starling-lab-to-protect-human-rights-with-decentralization
45. DLA Piper. (n.d.). *Data protection laws of the world*. https://www.dlapiperdataprotection.com/?c=US

CC BY-NC 4.0