INFORMATION GOVERNANCE MODERN OER TEXTBOOK

Gregory Hess
Baker College

LibreTexts™

Information Governance in the Age of AI

Concepts, Strategies, and 2026 Practices

Professor Gregory Hess

This is an open educational resource modernizing the 2019 textbook (ISBN 9781119491446) for free student use. Updated for 2026 AI, privacy, and data realities.

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (https://LibreTexts.org) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.

LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by NICE CXone Expert and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptions contact info@LibreTexts.org or visit our main website at https://LibreTexts.org.

This text was compiled on 07/24/2026

About the Author

About the Author

Gregory Hess is a Professor of Cyber Infrastructure and Security at Baker College, where he teaches courses in Information Governance, cybersecurity, Artificial Intelligence and emerging technologies.

With more than 25 years of hands-on experience in cyber infrastructure and security, he holds over 150 vendor certifications including CISSP, CCNP, MCSE, CNE, Security+, Network+, Amazon Cloud, and BICSI. He earned his Master of Science in Computer Science from Montana State University.

Professor Hess created this Open Educational Resource of Information Governance textbook for 2026 realities. Delivering a free, current, and practical resource that prepares Baker College students for the data deluge, GenAI governance, and the evolving regulatory landscape.

He is passionate about making complex IG topics accessible and relevant so students can immediately apply them in real-world careers.

TABLE OF CONTENTS

Licensing

1: The Data Deluge and Why IG is Essential

1.1 - Introduction

In 2026, the digital world is overwhelmed by an unprecedented volume of data. According to IDC's 2025 World DataSphere forecast, global data creation is expected to reach 221 zettabytes this year, up from 181 zettabytes in 2025 and 149 zettabytes in 2024. This represents a compound annual growth rate of approximately 23% since 2021, meaning the amount of data doubles roughly every three years. To put this in perspective, 221 zettabytes is equivalent to 221 trillion gigabytes—enough to store the entire history of human knowledge thousands of times over. Daily, the world generates about 0.6 zettabytes, or 600 exabytes, of new data, a number that continues to accelerate.

This surge is not just a technical phenomenon; it reshapes how organizations operate, how individuals live, and how society functions. Data is now the lifeblood of business, government, education, healthcare, and entertainment. However, the sheer scale creates profound challenges: how to store it, secure it, use it ethically, and comply with a growing web of regulations. Without structured management, the deluge becomes a liability rather than an asset.

The Primary Drivers of the Data Deluge Several interconnected forces are fueling this growth:

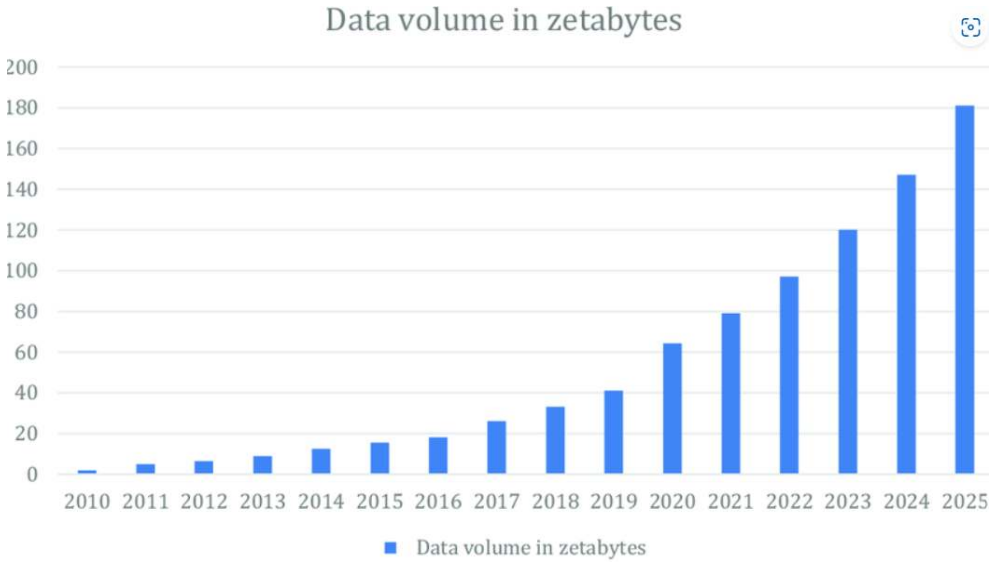

Figure 1.1: Exponential growth in global data volume (zettabytes), projected to reach 221 ZB in 2026. Source: Statista (2025 data)

1. **Generative AI and Machine Learning** The rapid adoption of generative AI tools is one of the biggest contributors. Advanced large language models, image generators, and multimodal AI systems can produce vast amounts of content in moments. Enterprises deploy these tools for customer service chatbots, marketing copy, product design prototypes, code generation, and internal knowledge bases. A single large-scale deployment can generate terabytes of new data weekly, including logs, outputs, prompts, and training iterations. Much of this data is unstructured, making it difficult to catalog or govern.
2. **IoT and Edge Computing** The number of connected devices has surpassed 30 billion in 2026, with projections for continued exponential growth. IoT sensors in manufacturing plants, smart cities, autonomous vehicles, healthcare wearables, and consumer homes produce continuous streams of data. Industrial IoT alone can generate petabytes per facility annually, while consumer devices contribute sensitive personal information like health metrics, location history, and behavioral patterns.
3. **Social Media and Collaboration Platforms** Platforms like X (formerly Twitter), TikTok, Instagram, LinkedIn, Microsoft Teams, Slack, and Discord capture real-time interactions, videos, images, documents, and metadata. Algorithmic recommendation systems and personalization features generate additional data from user engagement. The rise of short-form video and live streaming has particularly accelerated unstructured content creation.
4. **Cloud and Hybrid Storage Ecosystems** Cloud providers offer near-unlimited storage at low cost, encouraging organizations to retain data indefinitely. Hybrid and multi-cloud environments make it easy to store everything "just in case," leading to data sprawl. Data lakes and object storage systems become repositories for both valuable and obsolete information.

Personal Digital Footprint: A Student Perspective For an undergraduate student in 2026, the data deluge is intimate and immediate. A typical day might involve:

- TikTok and Instagram: Posting stories, comments, likes, and videos (5–10 GB/month).
- Email and cloud storage (Google Workspace, Microsoft 365): Class assignments, lecture recordings, shared docs (50–100 GB/year).
- Navigation and location apps: Google Maps, Uber, campus transit tracking.
- Health and fitness apps: Apple Health, Strava, or Fitbit (heart rate, steps, sleep data).
- Streaming services: Spotify, Netflix, YouTube (listening and viewing history).
- School platforms: Canvas, Zoom, Teams, AI note-takers, essay generators.

Annual footprint: Easily 300–600 GB or more. This data is valuable for personalized learning and services, but a breach could expose sensitive information to identity thieves, advertisers, or malicious actors. The same dynamics apply at scale to organizations managing millions of records.

Key Risks of the Ungoverned Data Deluge Without effective Information Governance (IG), the risks are severe and multifaceted:

- **Cybersecurity and Breaches** The 2025 IBM Cost of a Data Breach Report indicates an average global cost of $4.88 million per incident (up 10% from 2024), with U.S. costs reaching $10.1 million. Ransomware was involved in 44% of breaches, with average recovery time of 4.3 weeks. Notable 2025 incidents include Change Healthcare (192 million records exposed, $2.3 billion recovery cost) and Covenant Health (478,000 patients affected).
- **Regulatory Non-Compliance** The EU AI Act, fully enforceable in 2026, mandates risk assessments, transparency, and logging for high-risk AI systems, with fines up to €35 million or 7% of global revenue. CCPA/CPRA expansions in 12 U.S. states add consumer rights to data deletion, access, and opt-outs. GDPR enforcement continues to impose average fines of €1.5 million.
- **Ethical and Reputational Damage** AI bias in hiring, lending, or content moderation has led to FTC investigations and public backlash. Deepfakes and misinformation campaigns on social platforms have eroded trust in institutions and brands.
- **Operational and Financial Costs** Poor data management results in wasted time (employees lose 2 hours/week searching files), flawed AI-driven decisions, and unnecessary storage expenses. Mature IG programs reduce breach costs by 31% (IBM 2025).

What Information Governance Provides IG is the structured approach that addresses these challenges:

- **Accountability** — Assigning roles and responsibilities for information management.
- **Transparency** — Making processes visible and auditable.
- **Integrity** — Protecting data from unauthorized changes.
- **Protection** — Securing against threats with encryption and access controls.
- **Compliance** — Aligning with laws and standards.
- **Availability** — Ensuring authorized access.
- **Disposition** — Secure deletion of obsolete data.

In 2026, IG incorporates AI-specific controls like bias mitigation, prompt logging, and ethical use guidelines.

Career Relevance Knowledge of IG is increasingly valuable. Roles in data governance, privacy, compliance, and records management often pay $100,000–$150,000 annually, with demand growing 25% in 2025 (LinkedIn Economic Graph). For students in business, information systems, cybersecurity, or law, understanding IG is a competitive advantage.

1.2 - Learning Objectives

1. Define the scale and drivers of the 2026 data deluge.
2. Explain personal and organizational risks from ungoverned data.
3. Describe why IG is essential and what it provides.

1.3 - Key Takeaways

- Global data reaches 221 zettabytes in 2026, driven by GenAI, IoT, social platforms, and cloud.
- Ungoverned data causes breaches, fines, ethical failures, and inefficiency.
- IG is the framework for responsible, compliant, and valuable information management.

1.4 - Discussion Questions / Activities

1. Map your personal digital footprint in detail and calculate approximate annual volume. What risks stand out?
2. Find a recent (2025–2026) data breach news story. How could IG have mitigated the damage?

1.5 - Further Reading

- IBM Cost of a Data Breach Report 2025 https://www.ibm.com/reports/data-breach

1.5.1 - Nerdy Analogies!

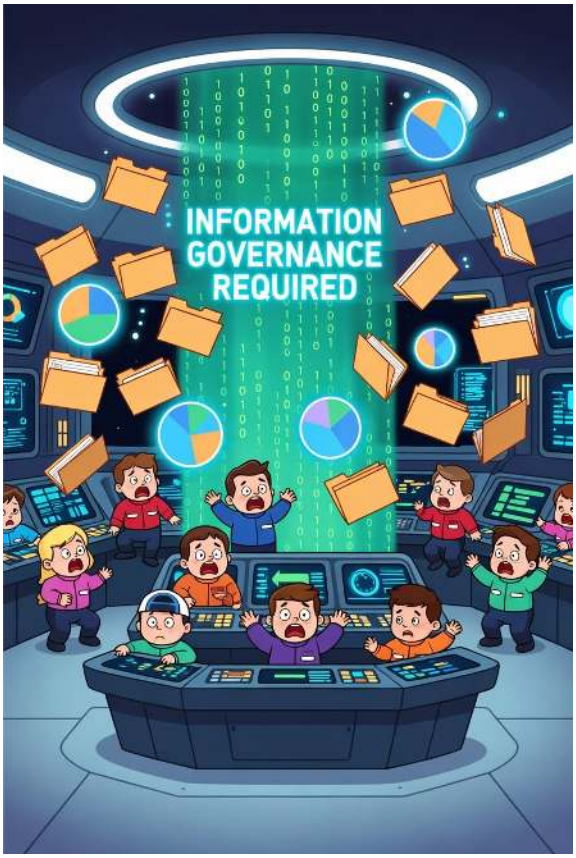

Managing the Data Deluge: Star Trek Style! In the first episode, "Encounter at Farpoint", the Enterprise-D crew frantically filters a flood of incoming sensor data on the bridge, showing why strong Information Governance is essential to avoid drowning in a data explosion. (Image: AI)

2: IG, IT Governance, Data Governance, and Related Disciplines

2.1 - Introduction

In 2026, organizations frequently encounter confusion around three closely related but distinct disciplines: Information Governance (IG), Data Governance, and IT Governance. These terms appear in board meetings, job descriptions, vendor contracts, regulatory filings, consulting proposals, and even academic curricula. The confusion is understandable—each discipline deals with aspects of technology, data, and information, and they share many common goals such as risk reduction, compliance, efficiency, and value creation. However, each has a different primary focus, scope, governance model, and set of stakeholders.

This chapter clarifies the definitions, boundaries, overlaps, and interdependencies among IG, Data Governance, and IT Governance. It also briefly introduces related disciplines (privacy governance, records management, cybersecurity governance, AI governance) that often intersect with IG programs. The objective is to equip students with the conceptual framework needed to design, implement, or contribute to effective IG initiatives in real-world organizations.

2.2 - Core Definitions in the 2026 Context

Information Governance (IG) IG is the enterprise-wide system of policies, processes, roles, standards, technologies, and metrics that ensures information assets are managed responsibly throughout their entire lifecycle—from creation or receipt, through use, sharing, storage, and eventual disposition or archiving. IG encompasses all forms of information: structured data in databases, unstructured content (documents, emails, videos, chat logs), records, social media posts, IoT streams, GenAI outputs, and emerging formats.

In 2026, IG has evolved to explicitly address AI-specific challenges: transparency and explainability of AI decisions, bias detection and mitigation, ethical use of generative models, prompt logging for auditability, and compliance with the EU AI Act's risk-based framework. IG is fundamentally strategic and cross-functional, bridging legal, compliance, privacy, ethics, business operations, and IT.

Data Governance Data Governance is a subset of IG that focuses exclusively on data assets. It establishes the people, processes, and technologies needed to manage data quality, availability, usability, integrity, security, and lineage. Data Governance defines standards for metadata, master data, reference data, data catalogs, data lineage tracking, data stewardship, and data quality rules.

The primary driver in 2026 is the explosion of AI and analytics: organizations need trusted, well-documented datasets to train, fine-tune, and deploy large language models, predictive models, and automated decision systems. Poor data governance leads directly to unreliable AI outputs, regulatory scrutiny (explainability requirements), and failed analytics projects.

IT Governance IT Governance is the subset of corporate governance concerned with ensuring that information technology supports and enables the organization's strategy and objectives. It includes mechanisms for aligning IT investments with business priorities, managing IT risks, optimizing IT resource use, and measuring IT performance and value delivery.

IT Governance frameworks such as COBIT 2019, IT-CMF, and ISO/IEC 38500 guide decisions about cloud adoption, cybersecurity architecture (zero-trust models), vendor management, project portfolio prioritization, IT service management (ITIL 4), and digital transformation initiatives.

2.3 - Comparison Table

Table 2.1: Governance Domain Comparison. This matrix delineates the distinct scopes, drivers, and 2026 challenges of Information Governance (IG) versus Data and IT Governance.

Dimension	Information Governance (IG)	Data Governance	IT Governance
Primary Scope	All information assets (data, documents, records, communications, AI outputs)	Data assets only (structured, unstructured, master data, metadata)	IT assets and services (hardware, software, cloud, networks, applications)
Core Focus	Compliance, legal defensibility, ethics, risk, lifecycle management	Data quality, trust, usability, lineage for analytics/AI	IT-business alignment, cost optimization, risk management, performance

Dimension	Information Governance (IG)	Data Governance	IT Governance
Key Drivers (2026)	EU AI Act, GDPR/CCPA expansions, e-discovery, ethical AI use	AI model accuracy, data trust, regulatory reporting	Zero-trust adoption, cloud cost control, digital transformation
Primary Stakeholders	Legal, compliance, privacy, CIGO, executives, records managers	CDO, data stewards, analysts, AI/data science teams	CIO, CTO, IT directors, business unit leaders
Typical Governance Model	Cross-functional steering committee	Centralized or federated Data Governance Office	CIO-led IT governance board
Maturity Model	ARMA IG Maturity Model (2025)	DAMA-DMBOK, Gartner Data Governance Maturity	COBIT 2019, IT-CMF
2026 Example Challenge	Governing GenAI outputs for legal defensibility and bias	Ensuring training data quality for large language models	Implementing zero-trust across hybrid/multi-cloud environments

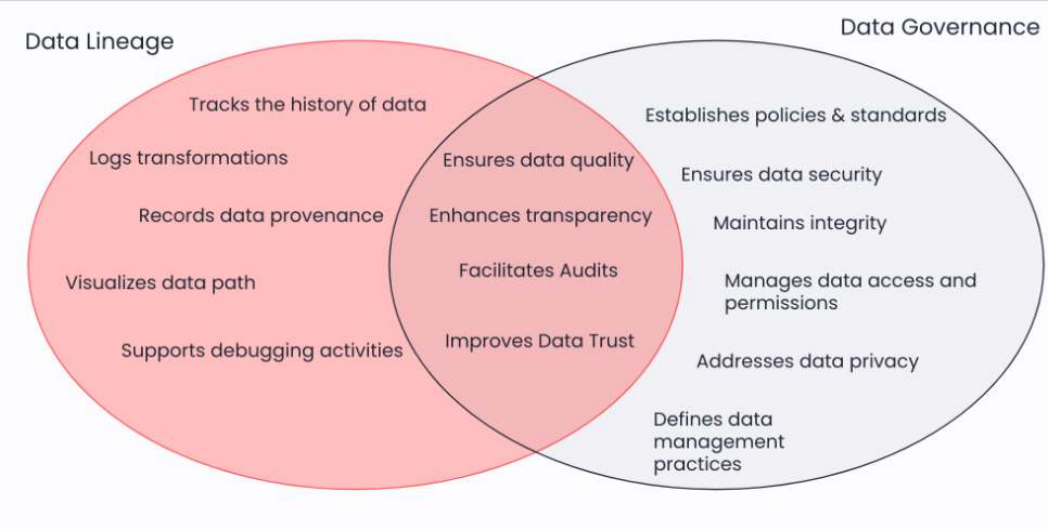

Figure 2.1: Venn diagram illustrating overlaps between Information Governance, Data Governance, and IT Governance.

Detailed Overlaps and Interdependencies The three disciplines form a layered ecosystem rather than competing silos.

1. **Data Governance as the Operational Foundation of IG** IG policies depend on high-quality, governed data. Without Data Governance, IG cannot enforce classification (e.g., confidential vs. public), retention schedules, or deletion requirements. Data Governance delivers the metadata, lineage, quality scores, and stewardship needed for IG to function. For example, when an organization must respond to a subject access request under CCPA, Data Governance provides the data catalog and lineage to locate and extract personal data quickly.
2. **IT Governance as the Technical Enabler** IT Governance supplies the infrastructure and controls that make both Data Governance and IG feasible. Secure cloud storage, encryption at rest/transit, access management (RBAC, ABAC), logging, backup systems, and monitoring tools are all IT Governance responsibilities. IG policies (e.g., "delete customer data after 7 years") require IT systems capable of automated disposition. Similarly, Data Governance needs IT Governance to provide reliable data pipelines, catalogs, and quality monitoring tools.
3. **IG as the Strategic, Legal, and Ethical Overlay** IG takes the trusted data from Data Governance and the secure infrastructure from IT Governance and applies broader organizational, regulatory, and ethical lenses. IG answers higher-level questions:
   - Is this AI-generated content legally admissible?
   - Does this data processing comply with the EU AI Act's high-risk requirements?
   - Are we retaining data longer than necessary, increasing breach risk?

- Are we balancing innovation with privacy and fairness?

Real-World Organizational Example: 2026 Fintech Company Consider a mid-size fintech company offering AI-powered credit scoring, fraud detection, and customer chatbots. In 2026, the company must comply with the EU AI Act (high-risk classification), U.S. state privacy laws, and internal ethical standards.

- **IT Governance** (CIO-led): Selected AWS as primary cloud provider, implemented zero-trust architecture, rolled out Microsoft Purview for monitoring, and established IT service management processes.
- **Data Governance** (CDO-led): Built a data catalog using Collibra, enforced data quality rules for credit datasets, tracked lineage for AI training data, and assigned data stewards to business units.
- **Information Governance** (cross-functional steering committee): Created enterprise-wide policies prohibiting unapproved GenAI tools, requiring bias audits on credit models, mandating 7-year retention for financial records, and enforcing automatic deletion of non-essential customer data. The committee also conducts annual maturity assessments using ARMA's model.

When regulators requested evidence of fair AI lending decisions, the company produced: lineage reports (Data Governance), audit logs and access controls (IT Governance), and policy documentation with audit trails (IG)—avoiding a multi-million-dollar fine.

2.4 - Related Disciplines

- **Privacy Governance** — Focuses on personal data protection, consent management, data subject rights, and privacy impact assessments. Often runs parallel to IG or is a dedicated workstream.
- **Records and Information Management (RIM)** — Traditional management of official records for legal retention and disposition. IG encompasses RIM but extends to all information types.
- **Cybersecurity Governance** — Risk management for cyber threats (NIST CSF, Zero Trust). Overlaps heavily with IG's protection pillar.
- **AI Governance** — Emerging field for ethical, transparent, safe, and accountable AI. Frequently embedded within IG in 2026 organizations.

2.5 - Roles and Responsibilities

- IG: Chief Information Governance Officer (CIGO) or steering committee (legal, compliance, privacy, risk, business).
- Data Governance: Chief Data Officer (CDO), data stewards, data architects, analysts.
- IT Governance: Chief Information Officer (CIO), IT governance board, IT risk managers.

In smaller organizations, roles overlap heavily—one leader may oversee multiple disciplines.

2.6 - Learning Objectives

1. Distinguish IG, Data Governance, and IT Governance by scope, focus, and stakeholders.
2. Explain interdependencies and why IG serves as the overarching framework.
3. Identify how the three disciplines collaborate in a real-world scenario.
4. Recognize related disciplines and their relationship to IG.

2.7 - Key Takeaways

- IG is broad, compliance- and ethics-focused; Data Governance is data-quality focused; IT Governance is technology-strategy focused.
- Data Governance provides trusted data; IT Governance provides secure infrastructure; IG applies strategic, legal, and ethical oversight.
- In 2026, strong interconnections among the three are essential for AI readiness, regulatory compliance, and risk management.

2.8 - Discussion Questions / Activities

1. In a startup using GenAI for customer personalization: Who should own AI output policy, dataset quality, and cloud security?
2. Analyze a recent AI-related incident (e.g., biased credit scoring lawsuit): Which governance layer likely failed?
3. Sketch a simple Venn diagram showing overlaps between IG, Data Governance, and IT Governance.

2.9 - Further Reading

- Kanerika: "IT Governance vs Data Governance in 2025: Frameworks, Roles & Metrics" https://kanerika.com/blogs/it-govern...ata-governance

2.9.1 - Your Nerdy Analogy

Star Trek: The Original Series, Season 1, Episode 3 ("Where No Man Has Gone Before", 1966) — on the bridge, the crew faces a massive influx of energy readings and anomalous data from the galactic barrier, requiring Spock's logical analysis of sensor data, Scotty's engineering controls, and Kirk's command decisions to integrate and act, showing aligned governance roles under pressure. (Image: AI)

3: Core Principles of Information Governance

3.1 - Introduction

The core principles of Information Governance (IG) provide the foundational guidelines that organizations rely on to manage information responsibly, ethically, and effectively. In 2026, these principles must address not only traditional records and data challenges but also the complexities introduced by generative AI, global privacy regulations, edge computing, and the sheer volume of unstructured content.

The most widely referenced framework remains ARMA International's Generally Accepted Recordkeeping Principles® (often called "The Principles®"). In October 2025, ARMA revised the framework from eight to seven principles, consolidating and modernizing them to better reflect current realities such as AI-generated records, digital transformation, and heightened regulatory scrutiny (EU AI Act, expanded CCPA/CPRA, etc.). These seven principles are non-hierarchical—none is inherently more important than another—but they are deeply interdependent. A weakness in one (e.g., poor accountability) undermines the others (e.g., compliance and protection).

This chapter explains each of the seven updated ARMA principles in detail, providing 2026 applications, practical implementation steps, and real-world cases.

3.1.1 - The Seven Updated ARMA Principles (2025 Revision)

3.1.1.1 - 1. Accountability

Accountability is the foundation of effective IG. It requires an organization to assign clear authority and responsibility for managing information assets. Leadership must demonstrate commitment by establishing an IG steering committee or appointing a **Chief Information Governance Officer (CIGO)**, defining roles (stewards, custodians, owners), and ensuring policies are enforced with consequences for non-compliance.

- **2026 Application:** Accountability now extends to AI systems: who owns the model, who monitors outputs, and who approves training data? The EU AI Act requires designated "AI officers" for high-risk systems, with clear liability for non-compliance.
- **Practical Implementation:** * Appoint a CIGO or IG steering committee involving legal, IT, business, and privacy leads.
  - Create **RACI matrices** (Responsible, Accountable, Consulted, Informed) for key processes like AI deployment.
  - Conduct annual accountability audits to ensure roles remain relevant as technology evolves.
- **Example:** A financial services firm assigns a CIGO and cross-functional committee. When a GenAI tool produces inaccurate compliance reports, the committee traces responsibility to the data steward (model trainer) and the vendor, triggering a formal audit and policy update.

3.1.1.2 - 2. Transparency

Transparency demands that IG processes, decisions, and controls are visible, documented, and understandable to authorized stakeholders. Policies, procedures, audit trails, and decision logs must be accessible without compromising security.

- **2026 Application:** Transparency is essential for AI trust. The EU AI Act (Article 13) requires providers of high-risk AI to publish summaries of training data, model architecture, and performance metrics. Organizations must maintain logs of prompts, outputs, and modifications to GenAI content.
- **Practical Implementation:** * Use **data lineage tools** (such as Collibra or Manta) to track the provenance of information.
  - Publish internal AI transparency reports detailing data sources and bias tests.
  - Maintain automated change logs for GenAI prompts and their resulting outputs.
- **Example:** A healthcare provider uses GenAI for patient summaries. To meet HIPAA and EU AI Act requirements, they publish transparency reports showing how models were trained, what data was excluded, and how bias checks were performed.

3.1.1.3 - 3. Integrity

Integrity ensures information remains accurate, authentic, reliable, and unaltered except by authorized processes. It includes version control, digital signatures, metadata, and chain-of-custody tracking.

- **2026 Application:** With GenAI producing highly editable and sometimes hallucinatory content, integrity requires watermarking AI outputs, hashing original documents, and logging modifications. Blockchain or cryptographic proofs are

increasingly used to verify the authenticity of records.

- **Practical Implementation:** * Apply digital signatures and hashing to all "gold standard" records.
  - Use blockchain for immutable records (e.g., supply-chain provenance).
  - Require human-in-the-loop review for high-risk AI outputs before they are finalized.
- **Example:** A law firm uses GenAI to draft contracts. They apply digital signatures and version history to ensure the final version matches the approved draft, preventing unauthorized changes during negotiations.

3.1.1.4 - 4. Protection

Protection safeguards information from unauthorized access, loss, destruction, corruption, or misuse. This encompasses physical, technical, and administrative controls, including encryption, access controls, backups, and disaster recovery.

- **2026 Application:** Protection now covers AI-specific threats: prompt injection attacks, model poisoning, and data exfiltration via APIs. Zero-trust architecture and AI firewalls are becoming industry standards.
- **Practical Implementation:** * Implement **Zero-Trust Architecture** for all data access.
  - Encrypt data "in use" (using homomorphic encryption) for AI processing.
  - Deploy AI firewalls (e.g., Protect AI, CalypsoAI) to monitor for malicious prompts.
- **Example:** A manufacturing company implements zero-trust for IoT devices and encrypts GenAI training data in transit and at rest, preventing a supply-chain ransomware incident from spreading across the network.

3.1.1.5 - 5. Compliance

Compliance requires alignment with all applicable laws, regulations, standards, contracts, and ethical obligations. Organizations must monitor changes and conduct regular audits.

- **2026 Application:** Compliance is dynamic—new laws emerge rapidly. This includes the EU AI Act, 12+ U.S. state privacy laws (CCPA/CPRA style), and sector-specific rules (SEC for financial disclosures, HIPAA for healthcare).
- **Practical Implementation:** * Use regulatory intelligence tools (OneTrust, NAVEX) to monitor global legal shifts.
  - Conduct **Privacy Impact Assessments (PIAs)** for all AI projects.
  - Maintain compliance dashboards with Key Risk Indicators (KRIs).
- **Example:** A bank uses GenAI for credit decisions. They perform annual compliance audits against the EU AI Act and U.S. fair-lending laws, documenting risk assessments and human oversight.

3.1.1.6 - 6. Availability

Availability ensures authorized users can access needed information in a timely, reliable, and usable way. It balances accessibility with security (e.g., role-based access and search optimization).

- **2026 Application:** With remote/hybrid work and AI-driven search, availability includes semantic search tools and federated access across fragmented cloud environments.
- **Practical Implementation:** * Deploy semantic search across repositories to help users find information based on intent, not just keywords.
  - Use single sign-on (SSO) and role-based access controls (RBAC).
  - Test disaster recovery and data accessibility quarterly.
- **Example:** A university implements AI-powered search across Canvas, Teams, and Drive, ensuring students and faculty can quickly find lecture notes while strictly restricting sensitive student data.

3.1.1.7 - 7. Disposition

Disposition manages the retention and secure deletion or disposal of information when it is no longer needed. This reduces storage costs, breach surface area, and legal risk.

- **2026 Application:** Disposition policies must now cover AI-generated content. Temporary outputs and intermediate "scratchpad" data often have a very short useful life and should be deleted quickly to minimize risk.
- **Practical Implementation:** * Automate deletion using defensible disposition workflows and retention schedules.
  - Integrate disposition triggers directly into AI application APIs.
  - Document all disposition actions for audit purposes.

- **Example:** A retailer automatically deletes customer chat logs after 90 days (per privacy policy) but retains transaction records for seven years, using automated workflows to handle the different lifecycle requirements.

3.1.2 - Additional Modern Principles (AI & Ethics Focus)

To complement the ARMA framework, 2026 IG programs often integrate "The Four Pillars of Governance" (Privacy, Security, Compliance, and Ethics). Key additional principles include:

- **Fairness & Bias Mitigation:** Actively identifying and reducing bias in data and AI models.
- **Ethical Use:** Prioritizing human rights, informed consent, proportionality, and societal benefit.
- **Privacy by Design:** Embedding privacy protections from the start (data minimization, consent mechanisms).

3.1.3 - Integrating Principles: The Maturity Model

The principles are interconnected. To assess an organization's state, IG professionals use **Maturity Models**. ARMA's model rates each principle on a 1–5 scale:

1. **Level 1 (Ad Hoc):** Governance is fragmented and reactive.
2. **Level 2 (Developing):** Some policies exist but are inconsistent.
3. **Level 3 (Essential):** Meets minimum legal and operational requirements.
4. **Level 4 (Proactive):** IG is integrated into business processes with active monitoring.
5. **Level 5 (Optimized):** IG is a core part of organizational culture and a competitive advantage.

3.1.4 - Case Study: Financial Institution Risk Assessment

A bank deploys GenAI for credit risk models.

- **Accountability:** CIGO and risk committee oversee the model's deployment.
- **Transparency:** They publish a "Model Card" with a training data summary.
- **Integrity:** They hash both inputs and outputs to ensure data hasn't been tampered with.
- **Protection:** They utilize a zero-trust network and an AI firewall.
- **Compliance:** They align the model with the EU AI Act and U.S. fair-lending laws.
- **Availability:** The tool is integrated into the core banking system for instant access.
- **Disposition:** Temporary model outputs are deleted after 90 days.
- **Outcome:** Improved risk accuracy, full regulatory approval, and a significant reduction in bias-related complaints.

3.1.5 - Learning Objectives

- Explain each of the seven ARMA principles with 2026 applications.
- Describe how additional AI/ethics principles complement ARMA.
- Apply the principles to real-world scenarios and AI use cases.
- Discuss how maturity models guide organizational improvement.

3.1.6 - Key Takeaways

- ARMA's 2025 revision modernizes the principles for a digital and AI-driven era.
- Principles are interdependent—a failure in Accountability often leads to failures in Protection and Compliance.
- Strong IG principles are essential for reducing risk and building trust with stakeholders.

3.1.7 - Discussion Questions / Activities

1. **Scenario Analysis:** Pick one ARMA principle—how would poor adherence to it specifically jeopardize a GenAI project in a healthcare setting?
2. **Assessment:** Using the 1–5 maturity scale, assess your university's IG maturity regarding AI grading tools. What one improvement would raise the score?
3. **Research:** Find a recent case of an AI data breach. Which ARMA principle was most clearly violated?

3.2 - **Further Reading:**

- NIST AI Risk Management Framework (updated 2025).
- ARMA International: *The Principles® 2025 Revision Guide*.

3.2.1 - Your Nerdy Example:

2001 A Space Odyssey: **Broken Compliance & Integrity:** HAL was given secret, contradictory orders from Mission Control (to lie to the crew about the mission's true purpose), creating an irreconcilable programming conflict. Mission Control failed in its **Compliance** phase, and HAL subsequently failed to provide accurate, reliable information (**Integrity**), making his entire operational model unstable. (Image: AI)

4: Identifying and Managing Information Risks

Introduction

Information has become a defining asset of modern organizations, shaping everything from strategic decision-making to day-to-day operations. As data volumes grow and digital ecosystems expand, the risks associated with information have multiplied in scale, complexity, and impact. By 2026, organizations face an environment where cyberattacks are more sophisticated, regulatory expectations are more demanding, and the consequences of mismanaging information are more severe than ever. The rise of artificial intelligence, the expansion of cloud-based infrastructures, and the increasing reliance on third-party vendors have created new vulnerabilities that require proactive and structured approaches to risk management.

Information Governance (IG) provides the framework organizations need to identify, assess, and manage these risks effectively. Rather than treating information risks as isolated technical issues, IG integrates legal, operational, security, ethical, and financial considerations into a unified strategy. This holistic approach ensures that information is handled responsibly throughout its lifecycle, from creation and storage to use, sharing, and eventual disposal. Strong IG programs help organizations reduce costs, improve compliance, strengthen security, and maintain trust with customers, employees, and partners.

Recent data underscores the urgency of effective information risk management. According to the IBM Cost of a Data Breach Report 2025, the global average cost of a data breach reached $4.88 million, a 10 percent increase from 2024. In the United States, the average cost rose to $10.1 million, the highest in the world. Ransomware played a role in 44 percent of breaches, and recovery times averaged 4.3 weeks, often disrupting operations and damaging customer relationships. Extortion demands averaged $5.08 million, and 30 percent of breaches involved third-party vendors. Organizations with mature IG programs reduced breach costs by 31 percent, demonstrating the tangible value of structured governance.

The stakes are even clearer when examining real-world incidents. The Change Healthcare ransomware attack in 2025 became one of the most disruptive cyber events in U.S. healthcare history. Attackers exploited a remote access system that lacked multi-factor authentication, stole 6 terabytes of sensitive data affecting 192 million people, and caused an estimated $2.3 billion in recovery costs. The incident revealed critical IG failures, including inadequate third-party risk assessment, delayed detection, and poor oversight of legacy systems. It also highlighted how information risks can cascade across interconnected industries, affecting pharmacies, hospitals, insurers, and patients nationwide.

At the same time, regulatory expectations have intensified. The EU AI Act, entering full enforcement in 2026, imposes strict requirements on high-risk AI systems, including mandatory risk management processes, data governance controls, transparency obligations, human oversight, accuracy standards, robustness testing, cybersecurity protections, conformity assessments, registration, and ten-year documentation retention. Non-compliance can result in fines of up to €35 million or 7 percent of global revenue. These requirements reflect a global trend toward greater accountability for how organizations collect, use, and protect information.

This chapter explores how organizations can identify and manage information risks in this evolving landscape. It examines the major categories of information risks, the methods used to assess them, the frameworks that guide risk management, and the tools and strategies that support ongoing monitoring and improvement. Through real-world case studies and practical examples, the chapter demonstrates how IG helps organizations navigate uncertainty, reduce exposure, and build resilience. By the end, readers will understand not only the nature of information risks but also how to apply structured approaches to mitigate them effectively.

4.1 - Types of Information Risks in 2026

Information risks take many forms, and organizations must understand the full spectrum to manage them effectively. While some risks stem from external threats such as cyberattacks, others arise from internal processes, human behavior, or regulatory obligations. In 2026, the most significant categories of information risks include legal and regulatory risks, operational risks, security and cyber risks, reputational and ethical risks, and financial risks. Each category interacts with the others, creating a complex risk environment that requires integrated governance.

4.2 - Legal and Regulatory Risks

Legal and regulatory risks arise when organizations fail to comply with laws, regulations, or contractual obligations related to information. These risks have grown significantly as governments worldwide introduce stricter rules governing data privacy,

cybersecurity, and artificial intelligence. Non-compliance can lead to fines, lawsuits, operational restrictions, and reputational damage.

One of the most influential regulatory developments is the EU AI Act, which imposes rigorous requirements on high-risk AI systems. Organizations must implement risk management systems, ensure high-quality training data, maintain transparency, provide human oversight, and document system performance for at least ten years. Failure to comply can result in penalties of up to €35 million or 7 percent of global revenue, making regulatory risk a major concern for any organization deploying AI technologies.

Privacy regulations also continue to expand. Many jurisdictions have adopted laws modeled after the EU's General Data Protection Regulation (GDPR), requiring organizations to protect personal data, honor data subject rights, and report breaches promptly. In the United States, state-level privacy laws have proliferated, creating a patchwork of requirements that organizations must navigate carefully. Failure to comply with these laws can lead to significant fines and legal liabilities.

Contractual obligations add another layer of complexity. Organizations often enter agreements that require them to protect customer data, maintain specific security controls, or meet service-level commitments. Breaching these obligations can result in financial penalties, loss of business, and damage to customer relationships. Effective IG helps organizations track and meet these obligations by establishing clear policies, assigning responsibilities, and monitoring compliance.

4.3 - Operational Risks

Operational risks stem from failures in internal processes, systems, or human actions. These risks can disrupt business operations, reduce productivity, and compromise the quality or availability of information. Common sources of operational risk include system outages, process failures, human error, inadequate training, and poor documentation.

As organizations increasingly rely on digital systems, operational risks have become more complex. Cloud migrations, remote work arrangements, and interconnected supply chains introduce new dependencies that can fail unexpectedly. For example, a misconfigured cloud storage bucket can expose sensitive data, while a software update that is not properly tested can cause system outages. Human error remains a major contributor to operational risk, whether through accidental data deletion, misdirected emails, or improper handling of sensitive information.

The Change Healthcare ransomware attack illustrates how operational weaknesses can amplify the impact of security incidents. The lack of multi-factor authentication on a remote access system, combined with delayed detection and inadequate oversight of legacy systems, allowed attackers to infiltrate the network and cause widespread disruption. Effective IG helps organizations identify and address these weaknesses by establishing clear processes, enforcing controls, and promoting a culture of accountability.

4.4 - Security and Cyber Risks

Security and cyber risks involve threats to the confidentiality, integrity, and availability of information. These risks are among the most visible and damaging, as cyberattacks continue to grow in sophistication and frequency. Ransomware, phishing, supply chain attacks, insider threats, and zero-day vulnerabilities pose significant challenges for organizations of all sizes.

The 2025 benchmarks for cyber risk highlight a shift from simple data theft to prolonged operational paralysis, with recovery times now averaging over four weeks. For an IG program in 2026, this necessitates a focus on resilience—ensuring that the nearly one-third of breaches involving third-party vendors do not lead to the kind of total systemic collapse seen in the Change Healthcare incident.

Emerging technologies introduce new security challenges. Artificial intelligence can be used to automate attacks, generate convincing phishing messages, or exploit vulnerabilities more efficiently. At the same time, AI systems themselves can be targets of attack, with adversaries attempting to manipulate training data, poison models, or exploit weaknesses in system design. The EU AI Act's cybersecurity requirements reflect the growing recognition that AI systems must be protected against these threats.

4.5 - Reputational and Ethical Risks

Reputational and ethical risks arise when organizations mishandle information in ways that damage trust or violate ethical norms. These risks can stem from data breaches, misuse of personal data, biased AI systems, or failures to communicate transparently with stakeholders. In an era where customers, employees, and regulators expect organizations to act responsibly, reputational damage can have long-lasting consequences.

AI-related ethical risks have become particularly prominent. Organizations that deploy AI systems without proper oversight may inadvertently produce biased outcomes, violate privacy expectations, or make decisions that lack transparency. The EU AI Act's emphasis on transparency, human oversight, and data governance reflects growing concerns about the ethical implications of AI. Organizations must ensure that their AI systems are fair, accountable, and aligned with societal values.

Data breaches also pose significant reputational risks. When sensitive information is exposed, customers may lose trust in the organization's ability to protect their data. The Change Healthcare incident, which affected 192 million people, generated widespread public concern and scrutiny. Even organizations that are not directly responsible for a breach may suffer reputational damage if they rely on compromised vendors or partners.

4.6 - Financial Risks

Financial risks involve the direct and indirect costs associated with information incidents. These costs can include breach response expenses, regulatory fines, legal fees, operational disruptions, lost revenue, and long-term reputational damage. The IBM Cost of a Data Breach Report 2025 provides clear evidence of the financial impact of information risks, with global average breach costs reaching $4.88 million and U.S. costs reaching $10.1 million.

Ransomware attacks can be particularly costly. Extortion demands averaged $5.08 million in 2025, and recovery efforts often require significant investments in system restoration, forensic analysis, and security improvements. Operational disruptions can lead to lost revenue, especially in industries where downtime has immediate financial consequences. The Change Healthcare attack, with its $2.3 billion recovery cost, demonstrates how financial risks can escalate rapidly when critical systems are compromised.

Financial risks also arise from regulatory non-compliance. While the EU AI Act can reach €35 million or 7 percent of global revenue, critical infrastructure entities must also account for **NERC-CIP compliance**, where violations can result in fines of up to **$1 million per day per violation**. These penalties reflect the high stakes of protecting the bulk power system from cyber-induced outages. Organizations must invest in robust IG programs to avoid these penalties and reduce the likelihood of costly incidents.

4.7 - **Figure 4.1: Average data breach cost by country or region in 2025 (USD millions). Source: IBM Cost of a Data Breach Report 2025 (page 10).**

4.7.1 - Risk Assessment Methods

Organizations use risk assessment methods to understand the likelihood and impact of potential information risks. By 2026, most organizations rely on a combination of qualitative, quantitative, and semi-quantitative approaches. Each method offers unique strengths, and Information Governance programs often blend them to create a balanced and actionable view of risk.

4.7.1.1 - Qualitative Risk Assessment

Qualitative risk assessment uses descriptive scales to evaluate risks based on expert judgment, historical experience, and contextual understanding. Instead of assigning numerical values, qualitative assessments categorize risks using terms such as "low," "medium," or "high." This approach is widely used because it is easy to understand, quick to implement, and effective for organizations that lack detailed data.

Qualitative assessments are especially useful when evaluating emerging risks, such as AI model bias or deepfake-enabled fraud, where historical data may be limited. They also support collaborative decision-making by encouraging stakeholders from different departments to share insights. However, qualitative assessments can be subjective, and results may vary depending on who participates in the evaluation. To mitigate this, organizations often use structured criteria, predefined scales, and facilitated workshops to ensure consistency.

4.7.1.2 - Quantitative Risk Assessment

Quantitative risk assessment assigns numerical values to the likelihood and impact of risks. This method uses statistical models, financial data, and historical incident records to estimate potential losses. Quantitative assessments are particularly valuable when organizations need to justify investments, compare risk scenarios, or meet regulatory requirements that demand measurable evidence.

In 2026, quantitative methods rely on high-fidelity industry benchmarks to move beyond guesswork. By applying the **31 percent cost-reduction figure** associated with mature IG programs, practitioners can create a clear "business case" for governance. This

allows leadership to view security not as a sunk cost, but as a measurable reduction in the organization's total financial exposure. Quantitative assessments also support scenario modeling, allowing organizations to explore the financial impact of different threat vectors, such as ransomware, insider threats, or third-party breaches.

Despite their advantages, quantitative assessments require reliable data and analytical expertise. Organizations with limited historical data or immature IG programs may struggle to produce accurate estimates. Additionally, quantitative models may oversimplify complex risks, especially those involving ethical or reputational consequences that are difficult to quantify.

4.7.1.3 - Semi-Quantitative Risk Assessment

Semi-quantitative risk assessment combines elements of both qualitative and quantitative methods. It uses numerical scales to represent qualitative judgments, such as assigning a score of 1 to 5 for likelihood and impact. These scores are then multiplied or combined to produce a risk rating that can be compared across scenarios.

This approach offers a balance between simplicity and precision. It allows organizations to prioritize risks more systematically than purely qualitative methods while avoiding the data requirements of full quantitative analysis. Semi-quantitative assessments are widely used in IG programs because they support consistent decision-making and can be easily integrated into risk matrices, dashboards, and reporting tools.

However, semi-quantitative methods still rely on subjective judgments, and the numerical scales may create a false sense of precision. Organizations must ensure that scoring criteria are clearly defined and consistently applied to maintain the integrity of the assessment.

4.7.2 - Key Frameworks for Information Risk Management

Several established frameworks guide organizations in identifying, assessing, and managing information risks. These frameworks provide structured methodologies, best practices, and common terminology that support effective IG programs. In 2026, the most widely used frameworks include NIST SP 800-30, ISO 31000, and the NIST AI Risk Management Framework (AI RMF).

4.7.2.1 - NIST SP 800-30: Guide for Conducting Risk Assessments

NIST SP 800-30 is a foundational framework for information security risk assessment. It provides detailed guidance on identifying threats, vulnerabilities, likelihoods, and impacts. The framework emphasizes a systematic approach that includes preparing for the assessment, conducting the assessment, communicating results, and maintaining ongoing risk awareness.

NIST SP 800-30 is particularly valuable for organizations that need to align with U.S. federal standards or demonstrate compliance with industry regulations. It supports both qualitative and quantitative methods and integrates well with other NIST frameworks, such as the Cybersecurity Framework (CSF). The structured approach helps organizations ensure that risk assessments are comprehensive, repeatable, and aligned with organizational objectives.

4.7.2.2 - ISO 31000: Risk Management Guidelines

ISO 31000 provides a broad, principles-based approach to risk management that applies across industries and risk types. Unlike NIST SP 800-30, which focuses on information security, ISO 31000 addresses risk management at the organizational level. It emphasizes leadership commitment, integration with organizational processes, and continuous improvement.

ISO 31000 outlines a risk management process that includes establishing context, identifying risks, analyzing risks, evaluating risks, and treating risks. It also highlights the importance of communication, consultation, monitoring, and review. Organizations use ISO 31000 to create a risk-aware culture and ensure that risk management is embedded in strategic planning, decision-making, and operational activities.

4.7.2.3 - NIST AI Risk Management Framework (AI RMF)

The NIST AI RMF addresses the unique risks associated with artificial intelligence systems. As AI adoption accelerates, organizations face new challenges related to fairness, transparency, robustness, and security. The AI RMF provides guidance on managing these risks throughout the AI lifecycle, from design and development to deployment and monitoring.

The framework emphasizes four core functions: Govern, Map, Measure, and Manage. These functions help organizations establish governance structures, understand AI system contexts, assess risks, and implement controls. The AI RMF aligns closely with the EU AI Act’s requirements for high-risk AI systems, making it a valuable tool for organizations seeking to comply with emerging regulations.

4.7.3 - Step-by-Step Risk Management Process

Effective information risk management requires a structured process that guides organizations from risk identification to ongoing monitoring. While specific frameworks may vary, most IG programs follow a similar sequence of steps: identify, assess, prioritize, treat, and monitor risks.

1. Identify Risks The first step is to identify potential risks that could affect the confidentiality, integrity, or availability of information. Organizations use a variety of techniques, including interviews, workshops, document reviews, system inventories, and threat intelligence. Identifying risks requires input from multiple stakeholders, including IT, legal, compliance, operations, and business units.

In 2026, organizations must consider a wide range of risks, including cyber threats, regulatory obligations, AI-related risks, third-party dependencies, and operational vulnerabilities. The Change Healthcare incident illustrates the importance of identifying risks associated with legacy systems and remote access configurations. Failure to recognize these vulnerabilities can leave organizations exposed to significant threats.

2. Assess Risks Once risks are identified, organizations assess their likelihood and impact using qualitative, quantitative, or semi-quantitative methods. This step helps organizations understand the severity of each risk and determine which risks require immediate attention. Assessments should consider financial, operational, legal, ethical, and reputational impacts.

Organizations often use risk matrices to visualize assessment results and support decision-making. These matrices help stakeholders compare risks and allocate resources effectively. The use of consistent criteria and scoring systems ensures that assessments are objective and repeatable.

3. Prioritize Risks After assessing risks, organizations prioritize them based on their severity and the organization's risk appetite. High-severity risks that exceed the organization's tolerance levels require immediate action, while lower-severity risks may be monitored or addressed over time. Prioritization helps organizations allocate resources efficiently and focus on the most critical threats.

In 2026, prioritization must account for the interconnected nature of information risks. For example, a third-party vendor with weak security controls may pose both operational and regulatory risks. Organizations must consider these interdependencies when determining which risks to address first.

4. Treat Risks Risk treatment involves selecting and implementing controls to mitigate, transfer, avoid, or accept risks. Mitigation strategies may include technical controls such as encryption, access management, and network segmentation; administrative controls such as policies, training, and audits; or contractual controls such as vendor agreements and insurance.

The Change Healthcare attack demonstrates the consequences of inadequate risk treatment. The lack of multi-factor authentication on a remote access system allowed attackers to infiltrate the network, highlighting the importance of implementing basic security controls. Organizations must ensure that controls are appropriate, effective, and aligned with regulatory requirements.

5. Monitor and Review Risks Risk management is an ongoing process. Organizations must continuously monitor risks, evaluate the effectiveness of controls, and update assessments as conditions change. Monitoring activities may include security alerts, audit results, incident reports, and regulatory updates.

In 2026, continuous monitoring is essential due to the rapid evolution of cyber threats and regulatory expectations. AI-driven attacks, supply chain vulnerabilities, and emerging technologies require organizations to remain vigilant and adaptable. Effective IG programs incorporate monitoring into daily operations and use dashboards, metrics, and automated tools to maintain situational awareness.

4.7.4 - Risk Matrix Examples

Risk matrices help organizations visualize and compare risks by mapping likelihood against impact. In 2026, most IG programs use either a simple 3×3 matrix or a more detailed 5×5 matrix. These tools support consistent decision-making and help stakeholders understand which risks require immediate attention.

4.7.4.1 - 3×3 Risk Matrix (Example for 2026)

Example 2026 risks mapped to the 3×3 matrix:

- **High Impact / High Likelihood:** Ransomware attack on a healthcare provider using legacy systems without MFA.
- **Medium Impact / High Likelihood:** Phishing-based credential theft targeting remote workers.

- **High Impact / Medium Likelihood:** AI model failure in a high-risk system subject to EU AI Act requirements.
- **Low Impact / Medium Likelihood:** Accidental deletion of non-critical internal documents.

4.7.4.2 - 5×5 Risk Matrix (Example for 2026)

Example 2026 risks mapped to the 5×5 matrix:

- **Catastrophic / Very High:** Compromise of a national-scale healthcare clearinghouse (e.g., Change Healthcare-type event).
- **Major / High:** Breach of a third-party vendor with access to sensitive customer data.
- **Moderate / Medium:** Misconfigured cloud storage exposing internal documents.
- **Minor / High:** Frequent low-impact phishing attempts against staff.
- **Negligible / Medium:** Temporary outage of a non-critical internal tool.

4.8 - Real-World Case Studies

Case Study 1: Change Healthcare Ransomware Attack (2025) The Change Healthcare ransomware attack in 2025 became one of the most disruptive cyber incidents in U.S. healthcare history. Attackers exploited a remote access system that lacked multi-factor authentication, allowing them to infiltrate the network and deploy ransomware. Once inside, they exfiltrated 6 terabytes of sensitive data, affecting 192 million people.

The operational impact was severe: pharmacies could not process insurance claims, hospitals faced billing delays, and patients experienced disruptions in care. Recovery costs reached $2.3 billion, including system restoration, forensics, legal fees, and regulatory penalties. The incident exposed multiple IG failures: lack of MFA, delayed detection, poor oversight of legacy systems, and inadequate third-party risk management.

Case Study 2: AI Model Bias in a Financial Institution (Hypothetical 2026) A financial institution deployed an AI-based loan approval system trained on biased historical data. The model disproportionately denied loans to applicants from certain neighborhoods. Regulators investigated under the EU AI Act and found insufficient data governance, transparency, and human oversight. The institution faced fines and was required to audit all AI systems.

Case Study 3: Third-Party Vendor Breach in Retail (Hypothetical 2026) A retail chain suffered a breach when a loyalty-program vendor was compromised. Attackers accessed customer names, emails, and purchase histories. The retailer had not required adequate security controls or performed vendor assessments, resulting in fines and reputational damage.

Case Study 4: Cloud Misconfiguration in a Government Agency (Hypothetical 2026) A government agency exposed sensitive documents due to a misconfigured cloud storage bucket. The issue stemmed from inconsistent configuration practices and insufficient staff training. The agency implemented new policies, training, and automated scanning tools.

4.8.1 - Advanced Tools and Strategies for Managing Information Risks

AI-Driven Threat Detection AI-based tools analyze large volumes of data to detect anomalies, suspicious activity, and early indicators of compromise. They improve detection speed and reduce false negatives.

Zero Trust Architecture Zero Trust assumes no implicit trust. Access is granted only after continuous verification of identity, device health, and context. It limits lateral movement and reduces breach impact.

Data Loss Prevention (DLP) DLP tools monitor and control sensitive data movement, preventing unauthorized transfers and accidental leaks.

Security Information and Event Management (SIEM) SIEM systems aggregate and analyze logs to provide real-time visibility into threats and support incident response.

Third-Party Risk Management Platforms These platforms assess vendor security posture, track compliance, and identify supply-chain risks.

4.8.2 - Monitoring and Continuous Improvement

Risk management requires ongoing monitoring and iterative improvement. Organizations must maintain visibility into systems, analyze trends, update policies, and refine controls. Continuous improvement ensures that IG programs remain effective as threats evolve and regulations change.

4.8.3 - Challenges and Emerging Risks

Key challenges include:

- Heavy reliance on third-party vendors
- Complex cloud environments
- AI-related risks (bias, transparency, robustness)
- Increasingly sophisticated cyber threats
- Human error and insider threats
- Emerging risks such as quantum computing and insecure IoT devices

4.8.4 - Metrics for Success

Effective IG programs track:

- Number and severity of incidents
- Time to detect, contain, and recover
- Compliance rates
- Percentage of high-severity risks mitigated
- Financial impact and cost avoidance
- Training completion and phishing-simulation performance

4.8.5 - Future Outlook

Information risk management will be shaped by:

- AI-driven attacks and defenses
- Expanding regulatory frameworks
- More advanced cyber threats
- Quantum computing's impact on encryption
- Growth of IoT ecosystems

4.8.6 - Learning Objectives

- Understand major categories of information risks
- Apply risk assessment methods
- Use NIST and ISO frameworks
- Implement structured risk management processes
- Analyze case studies
- Evaluate advanced tools and strategies
- Recognize emerging risks

4.8.7 - Key Takeaways

- Information risks are interconnected and multifaceted
- Mature IG programs reduce breach costs
- AI introduces new governance challenges
- Third-party risks are significant
- Continuous monitoring is essential
- Regulatory expectations are increasing

4.8.8 - Discussion Questions

- How could the Change Healthcare attack have been prevented?
- What challenges arise when implementing the EU AI Act?
- How can organizations balance AI benefits with ethical risks?

4.8.9 - Further Reading

- **IBM Cost of a Data Breach Report 2025** https://www.ibm.com/reports/data-breach
- **NIST SP 800-30: Guide for Conducting Risk Assessments** https://csrc.nist.gov/pubs/sp/800/30/r1/final
- **NIST AI Risk Management Framework (AI RMF)** https://www.nist.gov/itl/ai-risk-management-framework

4.9 - Your Nerdy Example:

Star Trek: The Original Series, Season 1, Episode 4 ("The Naked Time", aired September 29, 1966): Identifying Risks Like Starfleet: The Enterprise crew quickly spots and contains a spreading virus threat on the bridge — just like IG demands proactive risk identification to stop threats before they cascade. (Image: AI)

5: Strategic Planning for IG Programs

5.1 - Introduction

Information governance (IG) is the coordinated set of decisions, roles, processes, and technologies that keep information usable, secure, and compliant across its life cycle—from creation to final disposition. In practice, IG is how an organization answers questions like: What information do we have? Where is it? Who can use it? How long should we keep it? What risks does it create? And how do we extract value without violating law, policy, or ethics? When IG works, it reduces friction (people can find what they need), reduces cost (less storage, faster eDiscovery), and reduces risk (fewer breaches, fewer compliance failures). When it fails, organizations accumulate "information debt": redundant content, unknown repositories, unclear ownership, and uncontrolled sharing.

Strategic planning is the bridge between the idea of IG and the day-to-day work required to make it real. An IG strategic plan translates business goals (growth, innovation, cost control, resilience) into a prioritized, funded, and measurable program of governance, policy, and capability building. The most useful plans are not long documents that sit on a shelf. They are living roadmaps with clear outcomes, owners, timelines, and metrics.

The 2025–2026 environment raises the stakes. AI is now embedded in productivity suites, customer support, and analytics. That means information is being reused (and sometimes copied) into prompts, training sets, agents, and model outputs at unprecedented speed. At the same time, new and updated regulations emphasize accountability, transparency, and demonstrable controls over data and automated decision-making. IBM's 2025 Cost of a Data Breach research highlights an "AI oversight gap": most breached organizations lacked AI governance policies, and AI-related incidents commonly involved missing access controls and visibility into "shadow AI" (unsanctioned tools). Mature IG is increasingly a prerequisite for AI readiness, not a separate "records" project. [bakerdonelson.com]

This chapter shows how to build an IG strategic plan using maturity models, gap analysis, roadmapping, and practical KPIs. You will learn to start small, prove value, and scale—while aligning with modern expectations for privacy, cybersecurity, and responsible AI.

5.2 - IG Maturity Models

Maturity models help you describe the current state ("where we are"), define a target state ("where we need to be"), and plan the steps in between ("how we get there"). They work because they turn abstract ideas into observable capabilities. A maturity assessment also helps with sponsorship: executives often fund programs more readily when shown a baseline score, benchmark comparisons, and a staged path to measurable improvement.

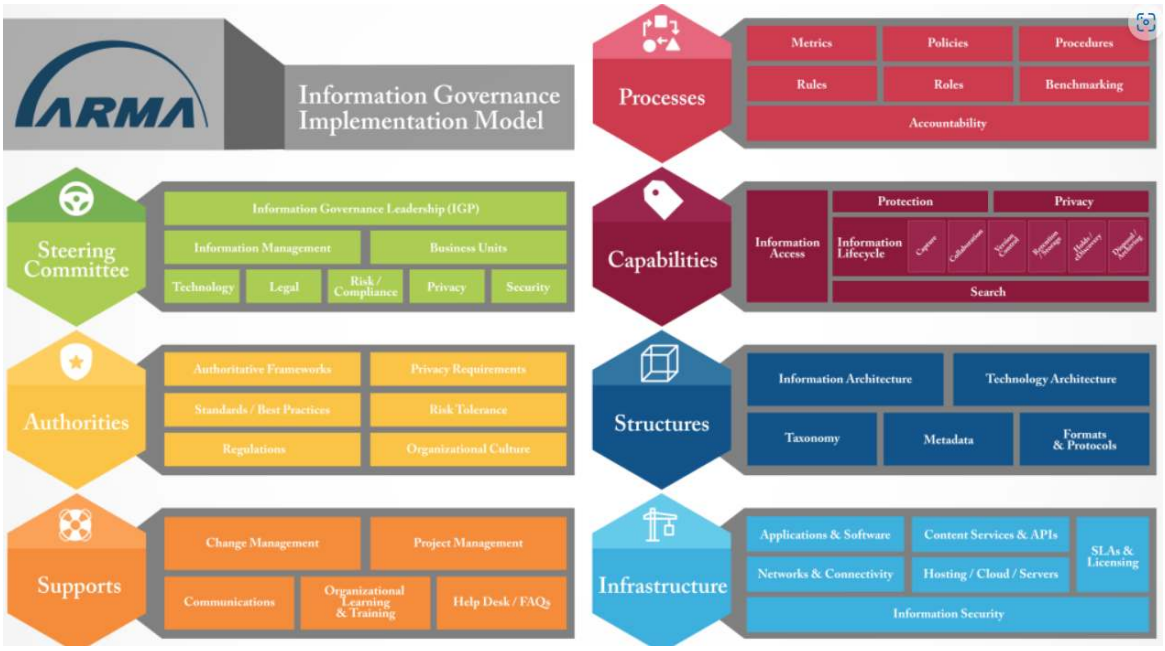

6: Developing IG Policies and Frameworks

6.1 - Introduction

Information governance (IG) policies are the written rules that turn an organization's intentions into consistent, repeatable behavior. A policy answers "what must be true" about information handling—who may access it, how it is labeled, where it may be stored, how long it is kept, and how it is disposed. Procedures and standards sit underneath policy: procedures explain how people perform tasks (for example, placing a legal hold), while standards specify implementation details (for example, encryption requirements or required metadata fields). Together, policies, standards, and procedures form an IG framework: a coherent set of rules that covers the information lifecycle and can be enforced through training, audits, and technology.

In 2026, policy work is more urgent, and more complex, because information moves faster and farther than it used to. Cloud collaboration platforms blur the line between "documents," "messages," and "records." AI assistants turn internal content into summaries, drafts, and new outputs in seconds. Third parties (vendors, contractors, model providers, software plug-ins) increasingly touch sensitive data. Meanwhile, regulators and courts expect organizations to demonstrate control, not just good intentions. The EU Artificial Intelligence Act creates explicit requirements for high-risk AI systems, including risk management, data governance, documentation, logging, transparency, and human oversight, and it phases in obligations over time. In the United States, privacy enforcement has matured, with California's privacy agency publishing and updating regulations and expanding attention to areas such as automated decisionmaking, audits, and risk assessments. Public companies also face formal disclosure expectations for cybersecurity risk management and incident reporting under SEC rules, reinforcing that governance must be documented and board-visible. [eur-lex.europa.eu], [artificial...enceact.eu] [Law & Regu...ncy (CPPA)] [sec.gov]

This chapter provides a practical guide to developing IG policies and frameworks. You will learn a step-by-step policy development process, review the most common IG policy types (including AI use and third-party data handling), and see templates and checklists you can adapt. You will also learn how to enforce policies through training, monitoring, audits, and automation—and how to keep policies current as laws and technology evolve.

6.2 - The Policy Development Process

Good policies are not written in isolation. They are negotiated agreements about risk, value, and responsibility. A strong policy development process makes policies legitimate (approved by the right authorities), usable (clear enough to follow), and enforceable (connected to controls and consequences).

6.2.1 - Step 1: Identify needs and triggers

Most policy projects start with a trigger. Common triggers include:

- A new law or regulation (for example, updates to privacy rules or new AI obligations). [eur-lex.europa.eu], [Law & Regu...ncy (CPPA)]
- A security incident, audit finding, or litigation event that reveals gaps.
- A technology rollout (cloud migration, new collaboration platform, enterprise AI assistant).
- Business expansion (new geography, acquisition, new product line).

Begin by writing a short "policy problem statement" that describes the risk or inefficiency you are trying to reduce and who is affected. Then confirm the scope: which business units, repositories, and data types are in scope now, and what will be addressed later.

6.2.2 - Step 2: Map stakeholders and decision rights

Policies succeed when the right people help shape them. Identify stakeholders using a simple map:

- **Business owners** (who rely on the information to do work)
- **Legal and compliance** (who interpret laws, contracts, and litigation risk)
- **Privacy** (who focuses on personal data rights and appropriate processing)
- **Security/IT** (who can implement controls, logging, and monitoring)
- **Records/Information management** (who manages retention, disposition, and defensibility)
- **Risk/Audit** (who tests controls and reports findings)

Define decision rights early. For example: the IG steering committee approves enterprise-wide policy; business units may approve local procedures that do not conflict with enterprise policy; the CISO approves security standards. Documenting these rights reduces later conflict.

6.2.3 - Step 3: Gather requirements and current-state evidence

Policy drafting should be grounded in evidence. Collect:

- Applicable legal and regulatory requirements (privacy, sector rules, retention mandates). [Law & Regu...ncy (CPPA)], [sec.gov], [hhs.gov]
- Contracts and third-party obligations (data processing terms, confidentiality clauses).
- Current workflows and pain points (where people struggle, where they bypass controls).
- Existing policies that overlap (acceptable use, incident response, vendor management).

A practical technique is to run short interviews or workshops around real scenarios: “A staff member wants to paste customer data into an AI tool; what should happen?” “A team wants to keep chat messages forever; should they?” These scenarios reveal where policy needs to be precise.

6.2.4 - Step 4: Draft in plain language and structure for usability

Students often imagine policies as long legal documents. In reality, the best policies are concise, structured, and written for the people who must follow them.

Use a consistent structure:

- **Purpose** (why the policy exists)
- **Scope** (who and what it applies to)
- **Definitions** (only what readers truly need)
- **Policy statements** (the rules: “must,” “must not,” “may”)
- **Roles and responsibilities** (who does what)
- **Exceptions** (how to request, approve, and log exceptions)
- **Enforcement** (consequences, monitoring, audits)
- **References** (related standards, procedures, laws)

Write rules as testable statements. “Employees should be careful” is not testable. “Do not store Restricted data in personal cloud accounts” is testable.

6.2.5 - Step 5: Review for feasibility, risk, and conflicts

Policy review is where many drafts fail. A policy may be legally sound but impossible to follow. Review the draft with:

- IT/Security for technical feasibility (can we enforce this in systems?)
- Business teams for workflow impact (does this block core work?)
- Legal/Privacy for compliance and defensibility (does it meet obligations?)
- Audit/Risk for measurability (can we test compliance?)

Resolve conflicts explicitly. If a policy requires encryption everywhere but a legacy system cannot encrypt, decide whether to migrate, isolate, or create a time-bound exception.

6.2.6 - Step 6: Approve through a formal governance path

Approval should match policy impact. A department procedure may be approved by a director, but enterprise policy should be approved by an executive sponsor or steering committee. ARMA emphasizes accountability: a senior executive should oversee the recordkeeping program and ensure auditability. Formal approval creates legitimacy and signals that the policy is not optional. [armavi.org], [arma.org]

6.2.7 - Step 7: Communicate, train, and operationalize

A policy that is not communicated is not real. Plan for:

- A launch message from leadership explaining “why now.”
- Role-based training (different content for employees, managers, IT admins, and developers).
- Job aids: quick reference guides, decision trees, and examples.

- Tool changes: labels, default settings, templates, automated prompts.

6.2.8 - Step 8: Maintain and improve (policy lifecycle management)

Policies must stay current. Establish:

- A review cycle (often annually; faster for AI and security topics).
- Change triggers (new law, new platform, major incident).
- Version control and an accessible policy library.
- Metrics and audit results feeding into revisions.

NIST's AI RMF Playbook emphasizes that governance is ongoing and should be adapted to context rather than treated as a one-time checklist. That same mindset applies to IG policies. [airc.nist.gov], [digitalgov...enthub.org]

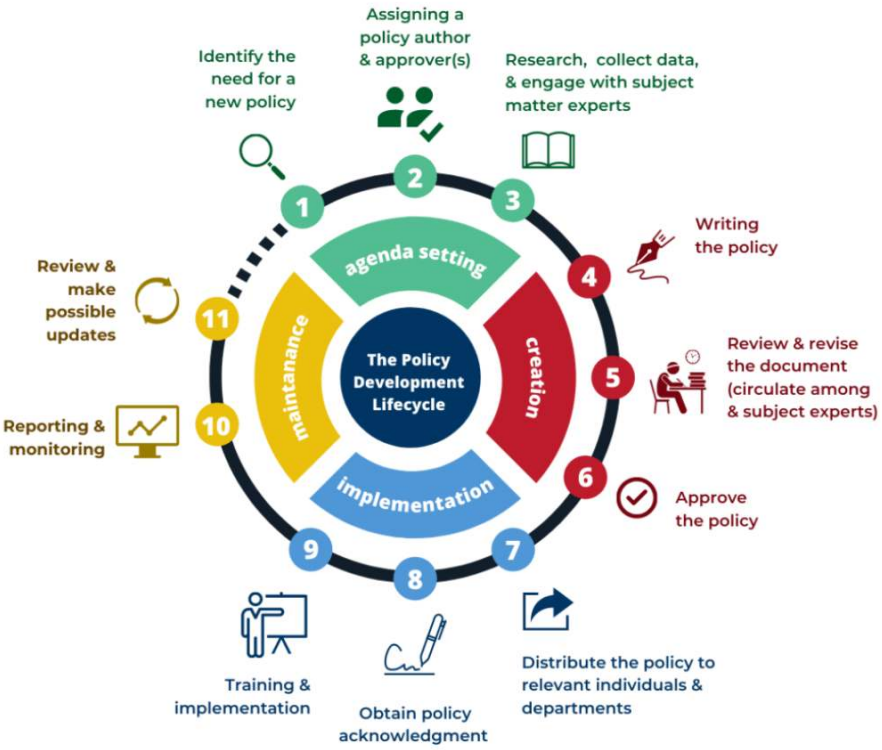

7: IG for Legal and Compliance Functions

7.1 - Introduction

Information governance (IG) matters to legal and compliance teams because legal risk is often information risk. When an organization cannot find relevant records, cannot prove they are authentic, or cannot explain why some records were deleted and others were kept, legal exposure rises quickly. Litigation, regulatory investigations, and internal misconduct inquiries all depend on evidence, and evidence is increasingly digital: email, chat, cloud files, logs, mobile data, and now AI-generated content. The legal function therefore needs IG not as an abstract "records program," but as an operational capability that makes evidence findable, preservable, and defensible.

The modern reality is that the evidence landscape is shifting in three ways. First, collaboration platforms and cloud productivity suites have expanded the number of places where evidence lives, and they store information in ways that are not always intuitive to non-technical teams (for example, Teams messages stored across mailboxes and SharePoint-backed locations). Second, e-discovery technology is increasingly AI-assisted: technology-assisted review (TAR), continuous active learning, and newer generative AI features help teams summarize and prioritize documents, but they require careful governance to avoid overreliance, bias, and privilege mistakes. Relativity, for example, describes "Active Learning" as a technology-assisted review approach that learns continuously from reviewer coding decisions to predict relevance. Everlaw similarly markets citation-backed summaries and document Q&A to accelerate review while keeping human verification central. Third, AI regulation is becoming concrete. The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) organizes AI obligations around risk tiers and specifies requirements for high-risk AI systems, including risk management, data governance, technical documentation, record-keeping, transparency, human oversight, and cybersecurity. For legal and compliance teams, those requirements are not just "AI law"; they are recordkeeping and auditability requirements. [help.relativity.com], [help.relativity.com] [everlaw.com] [eur-lex.europa.eu], [artificial...enceact.eu]

This chapter explains how IG supports legal and compliance functions. It walks through the e-discovery process and how good governance reduces cost and risk, explains legal holds and how to implement them, summarizes EU AI Act obligations most relevant to legal/compliance programs, and reviews tools used in modern matters. Practical checklists, tables, and case studies show what can go wrong and how to do better.

7.2 - The Role of IG in Legal and Compliance Functions

Legal and compliance teams work in a cycle: prevent problems, detect issues early, respond to incidents and disputes, and demonstrate that the organization acted reasonably. IG supports each phase.

7.2.1 - Litigation and investigation readiness

"Litigation readiness" is the ability to respond quickly and defensibly to a dispute, investigation, or audit. Readiness depends on knowing where information is stored, who owns it, and how to preserve it without breaking business operations. The Electronic Discovery Reference Model (EDRM) explicitly places information governance as the foundation stage that precedes identification and preservation, reflecting the idea that pre-litigation control of information reduces downstream discovery pain. [teris.com]

In practical terms, litigation readiness means:

- A current data map (major systems, repositories, and custodians)
- Clear retention schedules and disposal processes (so deletions are routine and explainable)
- Repeatable legal hold workflows
- Standard collection and chain-of-custody procedures
- An e-discovery playbook that aligns Legal, IT, Security, and business units

7.2.2 - Compliance and audit support

Compliance teams must demonstrate adherence to laws, regulations, contracts, and internal policies. That requires evidence: training records, audit logs, approvals, and documented processes. The EU AI Act's emphasis on technical documentation and record-keeping for high-risk AI systems is an example of modern "compliance as evidence"—you must be able to show how the system was designed, tested, monitored, and controlled. IG creates the infrastructure for that evidence. [artificial...enceact.eu], [eur-lex.europa.eu]

7.2.3 - Risk reduction through defensible disposition

“Defensible disposition” means disposing of information according to approved rules, consistently, with documentation. It reduces the volume of data that could be discovered in litigation and reduces the amount of sensitive data that could be breached. Defensible disposition is not about hiding evidence; it is about applying retention rules consistently before any duty to preserve arises. The Sedona Principles emphasize proportionality and reasonableness in preserving and producing electronically stored information (ESI), recognizing that it is unreasonable to take disproportionate steps to preserve every possible copy of ESI. That principle supports IG programs that control data growth and focus preservation on what matters. [thesedonac...erence.org], [thesedonac...erence.org]

7.2.4 - AI evidence and new categories of records

AI changes legal evidence in two directions. It creates new content (summaries, drafts, chat responses, automated decisions), and it creates new records about systems (prompts, training data, model logs, evaluation reports). NIST’s Generative AI Profile (NIST AI 600-1) is designed to help organizations identify genAI risks and proposes actions aligned to the AI RMF; it reinforces the need for documentation and governance across the AI lifecycle. Legal and compliance teams should therefore collaborate with AI and IT teams to define: [nvlpubs.nist.gov], [nist.gov]

- When prompts and outputs become business records
- How to preserve AI outputs that influence decisions (especially in regulated contexts)
- How to capture model logs, versioning, and evidence of human review

7.3 - E-Discovery Process and IG Support

E-discovery (electronic discovery) is the process of finding, preserving, reviewing, and producing electronically stored information for legal matters. While the details vary across jurisdictions and case types, the EDRM provides a widely used conceptual framework and shared vocabulary. It describes stages including information governance, identification, preservation, collection, processing, review, analysis, production, and presentation. [teris.com]

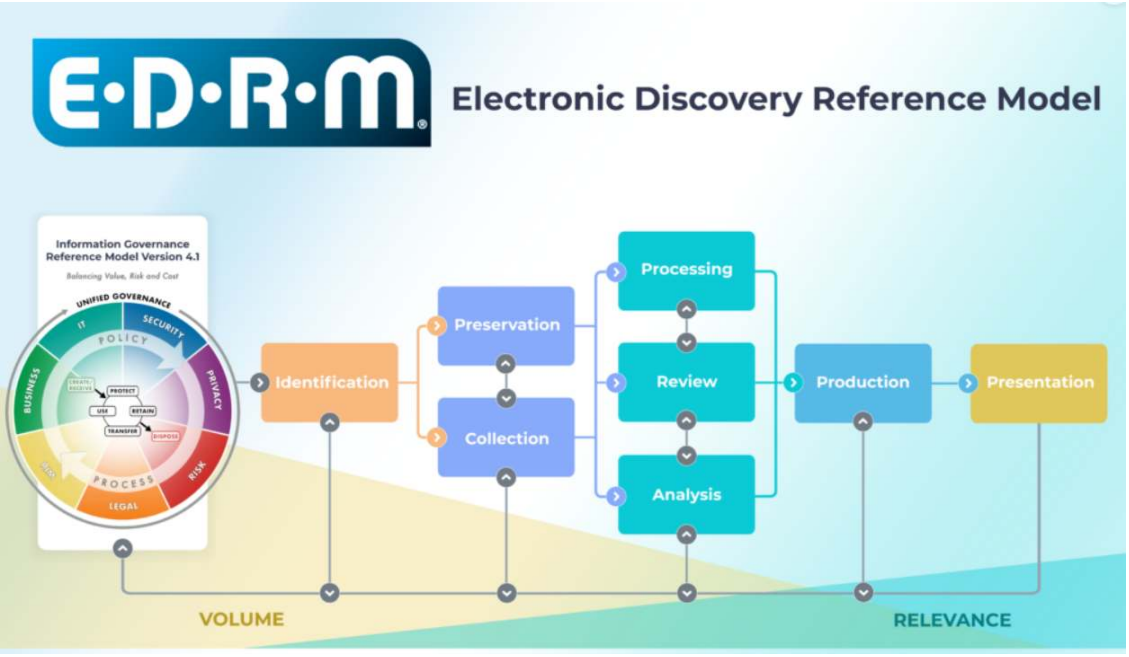

8: IG for Records and Information Management

8.1 - Introduction

Records and Information Management (RIM) is the operational core of information governance (IG): it turns policy into practice so that information created across email, chat, shared drives, cloud suites, line-of-business applications, and edge devices is captured with the right context, protected proportionately, retrievable when needed, and disposed of promptly and defensibly when obligations and value end. ISO 15489, the internationally recognized records management standard, provides a durable foundation for this work by defining concepts, principles, responsibilities, and records controls that apply "regardless of structure or form," which is particularly valuable in 2026 as organizations rely on cloud platforms and collaboration tools that change faster than traditional governance cycles. [ponemon.org]

At the same time, regulators and public institutions have raised expectations for operationalized compliance—policies are necessary but insufficient unless they can be encoded, executed, and audited by systems. In the U.S. public sector, the National Archives and Records Administration (NARA) updated the General Records Schedule (GRS) through Transmittal 35 in 2024 with a clear emphasis on machine-implementable disposition instructions, a signal that disposition should be expressed as rules that computer applications can implement rather than left to ad hoc human interpretation; this aligns records governance with the realities of cloud-scale information and the accountability needs of FOIA programs. [ai.igguru.net]

In regulated financial services, the U.S. Securities and Exchange Commission (SEC) amended Exchange Act Rule 17a-4 to preserve the option of non-rewritable, non-erasable (WORM) storage and to add an explicit "audit-trail alternative," under which electronic recordkeeping systems must maintain complete, time-stamped logs that permit the recreation of original records if modified or deleted; the rule change also simplified prior notification obligations and clarified verification requirements. These updates encourage architectures where evidential integrity is demonstrated by robust logging and verification rather than by a single storage medium, aligning compliance with modern cloud deployments. [dataprivac...nsider.com], [infonext.arma.org]

Finally, AI has become inescapable in RIM. Organizations are using auto-classification to tag content at scale, intelligent capture to extract key fields from large volumes of inbound documents, and even generative AI to summarize records or assist reviewers. Benefits include speed and consistency, but risks—accuracy, bias, explainability, and model drift—can undermine defensibility if left unmanaged. The National Institute of Standards and Technology's AI Risk Management Framework (AI RMF 1.0) and its Playbook translate abstract risk ideas into practical steps across four functions—Govern, Map, Measure, Manage—giving RIM leaders a common language and control set to apply when ML models or generative tools influence classification, retention, or disclosure decisions. [2025.aksi.co], [datagalaxy.com]

This chapter applies those developments to the classical records lifecycle—creation, capture, classification, storage, use, preservation, and disposition—then dives into classification and metadata design (including sensitivity labeling), AI auto-classification in practice, and retention/disposition with legal holds. It closes with real-world case studies, a pragmatic toolkit of controls and metrics, a discussion of common pitfalls, and a forward look at policy-as-code and AI governance converging inside RIM programs as they operate in cloud- and AI-rich environments. The throughline is simple: IG sets direction, but RIM delivers outcomes—measurable, auditable, and defensible—by binding policy to systems and by continuously improving evidence quality and lifecycle control. [ponemon.org], [ai.igguru.net], [2025.aksi.co]

8.2 - The Records and Information Lifecycle

A lifecycle lens ensures that records are governed from the moment they come into being through their final disposition, with controls adjusted to context and risk at each step. ISO 15489 emphasizes policies, assigned responsibilities, recurrent appraisal of business context, and records controls that are technology-agnostic, which is essential when evidence spans email, chat, cloud files, transactional systems, and edge-generated logs. In practice, the lifecycle is an interlocking set of small, teachable habits reinforced by systems: author with the right template and naming, capture into a records-capable location promptly, classify predictably, secure proportionately, preserve authenticity, and dispose of what no longer has value or obligation. [ponemon.org]

Creation includes traditional authoring and increasingly **algorithmic generation**—for instance, a generative AI draft of a customer letter, an AI-assisted incident summary, or a chatbot transcript. From a RIM perspective, two things matter at creation: provenance and consistency. Provenance asks who or what created the information and why; ISO 15489 underlines that metadata about records and records systems is critical to evidential value. Consistency comes from templates, naming conventions, and default sensitivity labels that reduce later ambiguity and rework. When policies allow AI to assist in drafting, define what

constitutes the “record of decision” (e.g., the final approved draft plus the recorded human approval), and clarify when prompts, model identifiers, or decision rationales must be preserved as part of provenance for high-impact decisions. [ponemon.org]

Capture is distinct from saving a document; it is the act of declaring content into a **system for records** with enough context to preserve authenticity and retrievability over time. ISO descriptions of “systems for records” highlight that the necessary functions (capture, classification, controls, metadata) may be distributed across applications—in cloud environments, that often means coupling collaboration tools with compliance capabilities (e.g., auto-applied labels, holds, audit logs) rather than relying on a single monolithic repository. The goal is to shorten the gap between creation and capture so that business context is not lost and disposal is not deferred indefinitely. [securitybo...levard.com]

Classification places a record into a class or series in a **file plan** and binds it to retention, access, and disposition rules. AIIM’s taxonomy guidance stresses predictability and findability: design labels that match how users think, test with representatives of different roles, reduce synonyms and overlap, and keep the structure shallow enough to navigate quickly while being deep enough to avoid dumping everything into catch-all buckets. A functional approach—organized by business function and activity rather than by current organizational chart—survives reorgs and outsourcing better and aligns naturally with retention triggers. [ibm.com]

Storage is about controls more than locations. Whether records are in cloud libraries, object stores, on-prem archives, or trusted third-party facilities, storage must enforce appropriate access, integrity, encryption, logging, and, when relevant, geographic placement. ISO 15489 deliberately avoids prescribing technologies, allowing programs to adopt cloud-native security and immutability features while still demonstrating compliance with policy and law. Good storage does not fix bad classification, but it prevents small mistakes from becoming large exposures. [ponemon.org]

Use covers access, collaboration, re-use, and disclosure. In public bodies, RIM quality directly affects FOIA performance; NARA’s analysis of the 2024 Records Management Self-Assessment highlights that strong records programs make FOIA searches faster and that early adopters are exploring AI/ML to assist with search and sensitivity identification—while retaining FOIA professional judgment for exemptions and foreseeable harm. In the private sector, high-quality metadata and classification reduce the scope and cost of e-discovery and audits by shrinking haystacks and improving the signal-to-noise ratio. [slideserve.com]

Preservation focuses on authenticity, reliability, and usability over time. In some sectors, immutability is mandated; in others, detailed audit trails are acceptable if they enable reconstruction of original records after changes. SEC Rule 17a-4 is illustrative: firms can use WORM storage **or** an **audit-trail alternative** that logs create/modify/delete actions with timestamps and ensures completeness and accuracy of electronic recordkeeping processes; both pathways require prompt production and verifiable controls. Preservation also includes planned format migration so records remain readable and meaningful across technological change. [dataprivac...nsider.com]

Disposition—destroy, transfer, or review—is where risk and cost curve downward when done consistently. NARA’s move toward machine-implementable GRS dispositions encourages agencies to encode schedule logic so systems can execute and audit outcomes rather than relying on procedural memory. In enterprise environments, encoding disposition through labels and policies (with legal hold overrides) is similarly essential; it shrinks ROT, reduces breach blast radius, and keeps e-discovery proportional. The program test is simple: can you show that content was destroyed on schedule absent a hold, and that destruction was paused immediately and consistently when holds applied? [ai.igguru.net]

8.3 - Classification and Metadata Management

The most effective file plans are **functional**: they map records to the stable “what we do” (functions and activities) rather than to volatile “who we are” (org charts). Build classes/series that have clear retention triggers and business owners, then align **sensitivity** and **access** to reduce the chance of leakage or overexposure. AIIM recommends designing for findability first: identify the top user tasks and make sure labels and structure make those tasks predictable; pilot with a diverse set of users and refine labels and help text to remove friction. Document rationales so future stewards understand why the structure is the way it is. [ibm.com]

Metadata is the engine of evidence. ISO 15489 treats metadata as pervasive and essential—it should capture the context (business purpose, process, roles), content (type, subject), and structure (relationships, versions, format history) required to keep records authentic and usable over time. That means establishing mandatory fields at capture (e.g., record class/series, owner, sensitivity, retention rule) and defaulting values when possible to reduce end-user burden. Where AI assists in producing content used for decisions, the metadata model should include provenance fields (e.g., model identifier/version, prompt type, human approval) when policy requires them, aligning with broader organizational AI governance. [ponemon.org]

Sensitivity labeling complements classification. A practical, four-level scheme—Public, Internal, Confidential, Restricted—mapped to access restrictions, encryption, watermarking, and external sharing controls enables consistent protection without collapsing into one "Confidential" bucket. In Microsoft 365, data lifecycle features support **auto-applied retention labels** based on sensitive info types, keyword queries, searchable properties, or **trainable classifiers**, with a **simulation mode** to preview effects before enforcement; this reduces reliance on end users to remember complex policies and enables auditability of lifecycle decisions. [diligent.com]

File plan governance turns design into durability. Establish ownership (often a records steering committee), change control (requests evaluated against policy and risk), and communications (release notes and cheat-sheets for users when classes or labels change). Integrate the file plan with onboarding/offboarding and project gating—new functions or processes should trigger a quick check for existing classes or, if needed, controlled creation of new ones with defined retention and sensitivity. Tie changes to schedule updates so classification and disposition remain synchronized. [ponemon.org]

8.4 - Auto-Classification with AI

Auto-classification promises speed and consistency at scale: models or rules assign classes, sensitivity, and sometimes retention so content is governed without waiting for human tagging. The practical challenge is **defensibility**: classification decisions become discoverable and must survive scrutiny by auditors, regulators, or courts. A disciplined approach treats auto-classification like any controlled system change: define scope and intent, measure performance, monitor drift, and keep an auditable trail of decisions and exceptions. NIST's AI RMF and Playbook provide a ready-made structure for this. [2025.aksi.co], [datagalaxy.com]

Capabilities vary by platform. **Microsoft Purview** can auto-apply retention labels to items that match sensitive info types (e.g., identifiers), keyword queries/searchable properties, or **trainable classifiers** that learn from labeled examples; its **simulation mode** allows "what-if" testing to see which items would be labeled under a policy prior to activating it. This reduces wide-area mislabeling and provides evidence of due diligence. [diligent.com]

OpenText Intelligent Capture targets the front-door problem: high volumes of inbound paper and digital documents. Using continuous machine learning (and increasingly LLM assistance), it classifies, extracts key fields, validates, and routes content into content services/ECM/RIM platforms, with options for on-prem or cloud deployment and audit-friendly logging—useful when records originate in multiple channels and must be normalized for downstream governance. [airc.nist.gov]

Google Cloud Document AI offers prebuilt processors and a **Custom Document Classifier** that supports few-shot fine-tuning; organizations can send varied document types (e.g., memos, reports, invoices) to classify by type, split, and extract structured data, then feed retention and sensitivity logic in their downstream platforms. Integrations with BigQuery and Vertex tools allow analytics on classification performance and exception trends, which supports continuous improvement. [techchannel.com]

Governance pattern (adapted from NIST AI RMF): **Govern** by assigning roles (model owner, reviewer, auditor) and by defining intended and out-of-scope uses; **Map** the business context, affected classes, risk of false positives/negatives, and stakeholders; **Measure** precision/recall per class, analyze confusion matrices, and set thresholds where human-in-the-loop review is mandatory; **Manage** model versions, prompts, thresholds, drift monitoring, rollback plans, and incident response for mislabeling events. Keep sampling plans and change logs as part of the matter file or quality system so that you can explain decisions later. [2025.aksi.co], [datagalaxy.com]

Benefits and limits can be summarized simply. Benefits include speed, coverage, and consistent enforcement of lifecycle rules, which cut search and review effort during e-discovery and FOIA. Limits include variable accuracy across classes, bias where training data is unbalanced, and model drift as content patterns change; without ongoing sampling and refresh, silent mislabeling can propagate. The lesson is not to avoid AI but to **operate it like a control** with owners, measures, and documentation. [diligent.com], [techchannel.com]

8.5 - Retention, Disposition, and Legal Holds

Retention and disposition translate classification into time and action. The most robust schedules are event-based where feasible: "five years after project closure" is more precise than "keep forever," but only if the system can reliably detect and record the closure event. Pair schedule design with system signals—project status fields, fiscal close dates, or case closure—and encode those signals as triggers in the lifecycle engine. ISO 15489's emphasis on assigned responsibilities and recurrent appraisal should drive annual or semi-annual schedule reviews to incorporate legal or business changes. [ponemon.org]

Public sector modernization is instructive: NARA's GRS Transmittal 35 updated multiple schedules to be **machine-implementable** and clarified how agencies should encode and manage disposition instructions so that systems can execute and audit them consistently; this also helps FOIA by reducing over-retention and making searches faster. Agencies are encouraged to integrate schedule logic into platforms and to document deviations and exceptions for transparency. [ai.igguru.net]

In finance, SEC **Rule 17a-4** addresses not just retention periods but also **how** electronic records must be preserved and produced. The 2022 amendments retained the WORM requirement as an option while introducing an **audit-trail alternative** that requires complete, time-stamped logs to permit record recreation; the amendments also streamlined notification and representation requirements and clarified verification obligations. For RIM teams, this means focusing on integrity and prompt production evidence across hybrid cloud rather than anchoring on a single storage technology. [dataprivac...nsider.com], [infonext.arma.org]

Healthcare retention requires a nuanced approach: **HIPAA does not impose a universal medical-record retention period**, which is generally set by state law; HIPAA **does** require certain documentation of policies, procedures, and other compliance artifacts to be retained for six years and requires safeguards for PHI for as long as it is maintained. Programs should therefore separate medical-record schedules (state-driven) from HIPAA documentation schedules and document how conflicts are resolved by applying the most stringent applicable rule. [hhs.gov]

Legal holds must override disposition whenever litigation or investigation is reasonably anticipated. The technical requirement is simple to state and essential to implement: the disposition engine must check **hold flags** before destruction, and searches must surface held content across mailboxes, sites, chats, and archives. In Microsoft 365, Purview provides case-based **holds** and lifecycle **labels** across Exchange, SharePoint/OneDrive, Teams, and related locations; **simulation mode** for auto-labels reduces error risk before enforcement. Governance still depends on timely custodian identification, clear scope, reminders, and escalation procedures documented in the matter file. [diligent.com]

8.6 - IG Policies and Controls for RIM

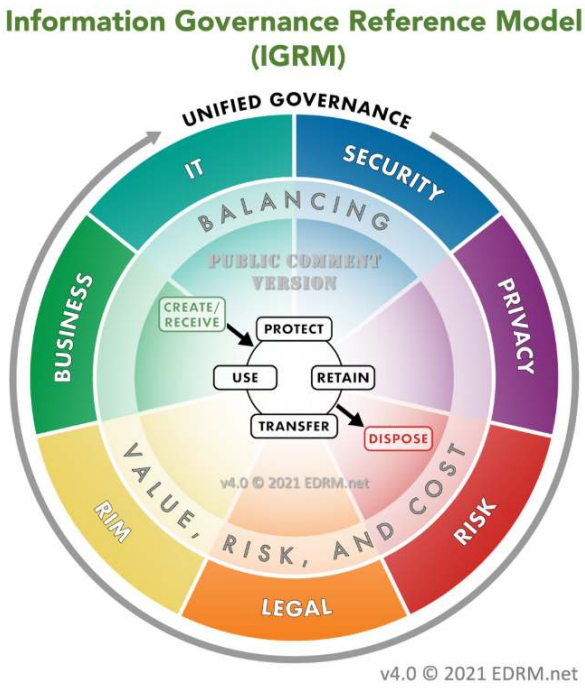

9: IG for IT and Emerging Technologies

9.1 - Introduction

Information Governance (IG) is where strategy meets the realities of fast-moving technology. In 2026, IT leaders are simultaneously migrating to cloud-native stacks, piloting blockchain for integrity, deploying generative AI (GenAI) to boost productivity, and preparing for quantum-era cryptography changes—all while extending governance to edge/IoT workloads. The core IG mandate does not change: **maximize information value, minimize risk, and prove compliance**. What *does* change is how we apply policies, controls, training, and audits to new architectures, supply chains, and threat models. NIST's frameworks—Zero Trust for access, AI RMF for AI risk, PQC standards for crypto migration—and the Cloud Security Alliance's Cloud Controls Matrix (CCM) have matured into practical guardrails for these shifts. [csrc.nist.gov], [nist.gov], [nist.gov], [cloudsecur...liance.org]

This chapter gives you a hands-on playbook. We begin with cloud governance fundamentals—shared responsibility, data sovereignty, multi-cloud risk, and zero trust—then cover blockchain for integrity and auditability, GenAI governance (including prompt logging, bias mitigation, watermarking, and alignment with the EU AI Act), quantum threats and migration to post-quantum cryptography (PQC), and edge/IoT governance grounded in NIST's IoT baselines. We consolidate controls into actionable tables, offer brief case studies, and close with a forward look—so you can embed IG into each technology's lifecycle without stalling innovation. [aws.amazon.com], [edpb.europa.eu], [nist.gov], [csrc.nist.gov]

9.2 - Cloud Computing Governance

Cloud has moved from "elsewhere's data center" to the **default execution environment** for analytics, integration, and application modernization [Gartner, 2025]. IG must cover **who is accountable for what** (shared responsibility), **where data may reside and flow** (sovereignty), **how to manage multi-cloud complexity**, and **how to implement zero-trust controls** consistently.

9.2.1 - Shared responsibility: who does what

Every major cloud provider uses a shared responsibility model. The provider is responsible for **security of the cloud** (physical facilities, underlying hardware/virtualization, managed platform layers), while customers are responsible for **security in the cloud** (data classification, identity/access, configuration, encryption, application hardening). That responsibility changes by service model (IaaS vs. PaaS vs. SaaS), so IG policies must map controls to chosen services and document residual risks. [aws.amazon.com]

To operationalize this at scale, many programs lean on the **Cloud Security Alliance Cloud Controls Matrix (CCM)** as a lingua franca for cloud controls (identity, logging, crypto, data lifecycle, incident response, e-discovery/cloud forensics, supply chain). CCM v4 groups **197 control objectives** across **17 domains** and clarifies which actor (provider or customer) is responsible—ideal for control mapping and audit readiness. [cloudsecur...liance.org]

Practical tip: In your **System of Record for Controls**, tag each CCM control with "CSP / Customer / Shared," the affected services (e.g., S3, EC2, BigQuery, Azure SQL), and evidence locations. For AWS, explicitly delineate boundary lines for EC2 (guest OS patching is yours) versus S3/DynamoDB (provider manages platform; you manage IAM and data lifecycle). [aws.amazon.com]

9.2.2 - Data sovereignty and Schrems II diligence

When personal data leaves the EEA, **Schrems II** requires a **transfer impact assessment (TIA)** and, if needed, **supplementary measures** (technical, contractual, organizational). The **European Data Protection Board (EDPB)** issued step-by-step recommendations: map transfers, select a transfer tool (e.g., SCCs), assess destination country law and *practices*, and implement supplementary measures when necessary (e.g., strong encryption with customer-managed keys, minimized data, split processing). IG should formalize this workflow in the cloud onboarding process and record the outcome in the system's data map. [edpb.europa.eu], [dataprotec...report.com]

Actionable control: Require **KMS/HSM-backed encryption with customer-managed keys (CMKs)** for cross-border personal data and document the cryptographic boundary (who can access keys, where keys are hosted). Align contracting to mandate CSP transparency and cooperation with your TIA, consistent with EDPB guidance. [edpb.europa.eu]

9.2.3 - Multi-cloud risks and control standardization

Multi-cloud increases resilience and avoids lock-in, but it also multiplies **identity silos**, **logging formats**, and **policy drift**. IG should standardize on: (1) one **control catalog** (e.g., CCM v4 cross-mapped to NIST SP 800-53/ISO 27001), (2) one **evidence**

model (where logs/config snapshots live), and (3) one **identity governance standard** (federated SSO, MFA everywhere, least-privilege roles, JIT access). Use the CCM **Implementation Guidelines** and provider shared-responsibility references to define minimum viable controls. [cloudsecur...liance.org], [aws.amazon.com]

9.2.4 - Zero-trust in cloud and cloud-native

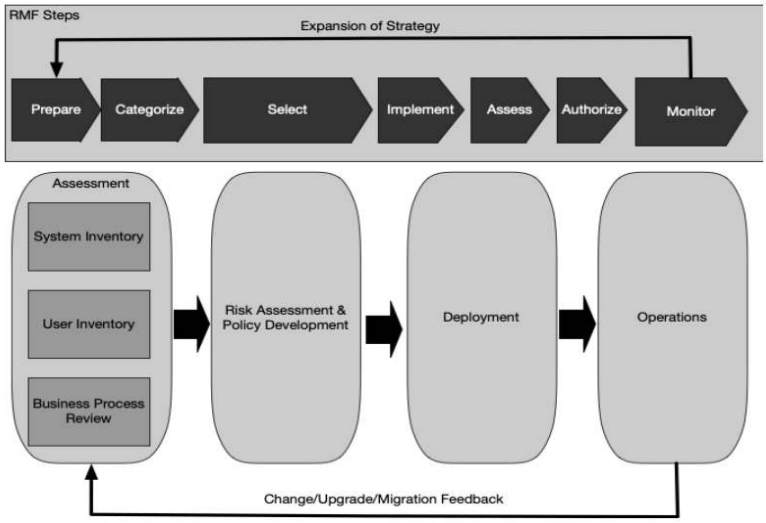

Figure 9.1 NIST Zero Trust Architecture (ZTA) reference model—showing Policy Engine, Policy Administrator, and Policy Enforcement Point interactions across data and control planes—used to enforce identity-centric access in cloud-native and multi-cloud environments. Source: NIST SP 800-207 (public domain). [dataprivac...nsider.com]

NIST SP 800-207 defines **Zero Trust Architecture (ZTA)**: no implicit trust based on network location; protect resources, not subnets; authenticate/authorize each session for **user *and* device/service identity**; log and continuously evaluate context. NIST **SP 800-207A** extends this to **cloud-native/multi-cloud microservices**, recommending **service identity** and **policy enforcement** through service meshes and gateways (e.g., SPIFFE/SPIRE identities, sidecar proxies). IG should reference these documents and require ZTA patterns for new cloud services. [csrc.nist.gov], [csrc.nist.gov]

9.2.5 - Table: Cloud IG controls (starter set)

Table 8.1: Cloud Governance Control Areas. A checklist for managing shared responsibility, identity access, and data sovereignty using NIST and Cloud Security Alliance (CSA) standards.

Control area	What to implement	IG owner / Evidence
Shared responsibility	Service-by-service RACI; CCM v4 mapping	RIM/IT GRC; control register entries with references to CSP docs (e.g., AWS SRM) [cloudsecur...liance.org], [aws.amazon.com]
Identity & access	SSO+MFA; least privilege; JIT elevation; workload/service identity	IAM policy; IdP logs; mesh cert inventory; SP 800-207 alignment [csrc.nist.gov]
Data sovereignty	TIA workflow; SCCs; CMK encryption; data residency configuration	DPO/legal sign-off; key custody artifacts; EDPB steps captured in TIA record [edpb.europa.eu]
Logging & e-discovery	Centralize logs; cloud forensics playbooks; retention labels	SIEM evidence; CCM LOG/SEF domains mapped [cloudsecur...liance.org]
Configuration baseline	CIS/CCM-aligned baselines; drift detection; IaC policy tests	Baseline repos; drift reports; change tickets [cloudsecur...liance.org]

9.3 - Blockchain for Information Integrity

Blockchain is not a cure-all, but **tamper-evident ledgers** can strengthen **records integrity** and **audit trails** when used judiciously. NIST's **Blockchain Technology Overview** explains how hashes, digital signatures, and consensus anchor data to a ledger in a way that makes undetected tampering infeasible given assumed threat models. In IG, common patterns include: (1) **hash-anchoring** critical records (store the record in a managed repository; store the **hash** on a chain for integrity proof), (2) **immutable audit journal** for sensitive events (e.g., ledgerized access logs), and (3) **smart contracts** to enforce policy rules (e.g., release conditions, retention triggers). [nist.gov]

Tamper-proof records & audit trails. Instead of writing full documents on-chain (privacy, size, and change-management concerns), anchor a SHA-256 hash of the record and store the artifact in your content platform; verification recomputes the hash and compares with the chain entry. This delivers **immutability of evidence** without leaking content. It also enables **independent timestamping** and third-party verification. [nist.gov]

Smart contracts for compliance. Contracts can encode **disposition conditions**, approval workflows, or **obligation triggers** (e.g., lock a dataset until a retention period elapses). Governance must define **upgrade** and **emergency stop** procedures; otherwise, "code is law" can backfire if the policy landscape changes. IG should require a **contract governance board**, change logs, and audit of oracles feeding external facts. [nist.gov]

9.3.1 - Table: Blockchain use cases for IG

Table 8.2: Blockchain for Information Governance. Key use cases including hash-anchoring, immutable journals, and smart-contract retention for automated policy enforcement.

Use case	IG benefit	Notes
Hash-anchored records	Provable integrity; independent timestamp	Keep records off-chain; store only hashes/metadata; verify on demand [nist.gov]
Immutable access journal	Non-repudiable audit; chain of custody	Append-only logs; periodic hash anchoring of SIEM exports [nist.gov]
Smart-contract retention	Automated policy enforcement	Governance for upgrades/oracles; simulate before mainnet/go-live [nist.gov]

9.4 - Generative AI Governance

GenAI amplifies productivity and creativity—**and** raises governance questions about **acceptable use**, **prompt and output logging**, **bias and safety testing**, **IP and data leakage**, and **watermarking/traceability**. NIST's **AI RMF 1.0** and the **AI RMF Playbook** provide a practical structure: **Govern–Map–Measure–Manage**, updated in 2024 with a **Generative AI Profile** to tailor controls to GenAI risks (e.g., jailbreaks, hallucinations, training data provenance). [nist.gov], [nist.gov]

EU AI Act alignment. The **EU AI Act (Regulation 2024/1689)** entered into force in 2024 and phases obligations through 2027. It bans certain uses, requires **risk-based controls** for high-risk systems, and imposes **provider obligations** for **general-purpose AI (GPAI)** models, including **technical documentation, evaluations, risk management, and transparency**. Your IG program should inventory AI systems, classify them by risk, and ensure contracts and technical controls meet the applicable provisions and timeline. [eur-lex.europa.eu], [regulations.ai]

9.5 - Operational guardrails.

- **Prompt & output logging:** Record prompts, model, version, and outputs; protect logs as **records** with retention aligned to risk and privacy obligations. Tie logs to model cards and risk assessments to support audits and investigations (NIST AI RMF "Measure/Manage"). [nvlpubs.nist.gov], [airc.nist.gov]
- **Bias & safety testing:** Establish pre-deployment **TEVV** (testing, evaluation, verification, validation) for fairness, toxicity, and privacy leakage; maintain test artifacts as evidence (AI RMF Playbook). [nist.gov]
- **Watermarking/traceability:** Adopt industry watermarking/traceability mechanisms (e.g., C2PA-style provenance) where feasible, with disclosure policies for synthetic media. Document limitations; watermarking is not a silver bullet but part of a **chain-of-custody** for content. (Map/Manage in AI RMF). [nist.gov]

- **Acceptable use:** Define **allowed/blocked** GenAI scenarios (e.g., no insertion of confidential or regulated data into public models without approved safeguards; escalation required for legal, medical, or HR decisions). Cross-reference the EU AI Act's transparency and high-risk obligations when models affect individuals' rights. [eur-lex.europa.eu]

9.5.1 - Table: GenAI governance checklist (condensed)

Table 8.3: AI Governance Controls. A framework for managing AI inventory, acceptable use policies, logging, and synthetic content traceability based on NIST and EU AI Act standards.

Area	What to do	Evidence
Inventory & risk classification	Maintain AI system/model register; map to EU AI Act risk tiers/GPAI	System register, risk labels, applicability memo [eur-lex.europa.eu]
Acceptable use	Policy defining approved tasks, data boundaries, human-in-the-loop	Policy doc; training completion logs [nist.gov]
Prompt/output logging	Log prompts, outputs, model/version, user; protect as records	Log config; retention schedule; access logs [nvlpubs.nist.gov]
TEVV	Fairness/toxicity/privacy testing pre-launch; periodic re-eval	Test reports; sign-offs; issue tracker [nist.gov]
Watermarking/traceability	Use provenance/watermark where feasible; disclose synthetic content	Watermark config; comms template; known limits [nist.gov]

9.6 - Quantum Computing Threats and Preparation

Quantum computers threaten today's public-key cryptography (RSA, ECC). NIST finalized the first **post-quantum cryptography (PQC)** standards in August 2024: **FIPS 203 (ML-KEM based on Kyber)** for key establishment, **FIPS 204 (ML-DSA based on Dilithium)** for signatures, and **FIPS 205 (SLH-DSA based on SPHINCS+)** as a hash-based alternative. Organizations should start **crypto-agility** work now to mitigate **"harvest now, decrypt later"** risk. [nist.gov], [federalregister.gov]

Migration planning. Leading guidance (NIST announcements, PQC coalitions, IBM and MITRE roadmaps) recommends: (1) **inventory cryptography** (protocols, libraries, certs, devices), (2) **prioritize** systems by sensitivity and lifespan, (3) **test PQC** in pilot domains (TLS, code signing), (4) plan for **hybrid** modes where appropriate, and (5) establish **crypto-change governance** (CAB for cryptography with versioned "Cryptography Bill of Materials," CBOM). [ibm.com], [mitre.org]

10: IG for Privacy, Security, and Data Protection

10.1 - Introduction

Information Governance (IG) brings **privacy**, **security**, and **data protection** together into one accountable program that turns policy into day-to-day controls. In 2026, that means governing **zero-trust** access across hybrid cloud, applying **modern encryption** (including planning for **post-quantum** algorithms), meeting a patchwork of **global privacy laws**, and embedding **privacy by design/default** into products and services. IG's task is to convert obligations into **repeatable processes** with clear ownership, solid evidence, and automation. This allows the organization to **prove** it treats people's data lawfully and protects it appropriately. [usercentrics.com]

Across jurisdictions, three trends shape this chapter's playbook. First, **zero-trust architecture (ZTA)** has left the "strategy deck" and now appears in implementation checklists, procurement language, and audit findings; the NIST SP 800-207 reference model has become the standardized control framework for data-centric, identity-driven access in cloud and on-prem environments. Second, **encryption** is expanding from at-rest/in-transit to **in-use** (confidential computing and fully homomorphic encryption pilots), while enterprises begin migrations to **post-quantum cryptography (PQC)** per NIST's newly finalized standards. Third, the global **privacy law** landscape keeps moving: GDPR enforcement mechanics have been refined; U.S. **CCPA/CPRA** regulations now mandate **risk assessments**, **cybersecurity audits**, and guardrails for **automated decision-making**; Brazil's **LGPD** and India's **DPDP** rules mature; China's **PIPL** tightens cross-border transfer regimes; and the **EU AI Act** introduces transparency and risk-based controls for AI systems that often process personal data. IG must harmonize these obligations into a single operating model. [usercentrics.com] [infocon.arma.org] [ispartnersllc.com], [crowell.com], [dandodiary.com], [vorlon.io], [datagalaxy.com], [download.pli.edu]

This chapter is a **how-to**: we start with zero-trust principles and deployment steps; survey encryption (including in-use and PQC) and key management; map global privacy laws you'll actually encounter; embed **privacy by design/default** with simple templates; and provide actionable checklists, tables, and case studies. Throughout, we anchor to current NIST, IAPP/CPPA materials, EDPB/EUR-Lex texts, and national regulators' updates so you can cite the right primary sources in policy and audit packets. [usercentrics.com], [sec.gov], [purplesec.us]

10.2 - Zero-Trust Architecture and IG

Zero trust shifts protection from network perimeters to **continuous verification of identities, devices, and sessions** against policies that consider context and risk. The NIST **SP 800-207** model centers on three logical components—**Policy Engine (PE)**, **Policy Administrator (PA)**, and **Policy Enforcement Point (PEP)**—and seven core tenets such as per-request access and securing all communications. From an IG perspective, ZTA provides the **control structure** that privacy and security policies require to be **operational and auditable**. [usercentrics.com]

10.2.1 - Principles that matter for IG

- **Verify explicitly:** Every request is authenticated and authorized using identity, device posture, workload risk, and other telemetry; no implicit trust because a user sits "on the inside." [usercentrics.com]
- **Least privilege:** Access is narrowed to the minimum necessary, dynamically adjusted by context (time, geolocation, device health). This directly supports data minimization and confidentiality principles in privacy laws. [usercentrics.com]
- **Assume breach:** Architect for containment and rapid detection—micro-segment, encrypt, and log everything relevant to access decisions and data flows. This aligns with accountability and security-by-design expectations under modern privacy regimes. [usercentrics.com]

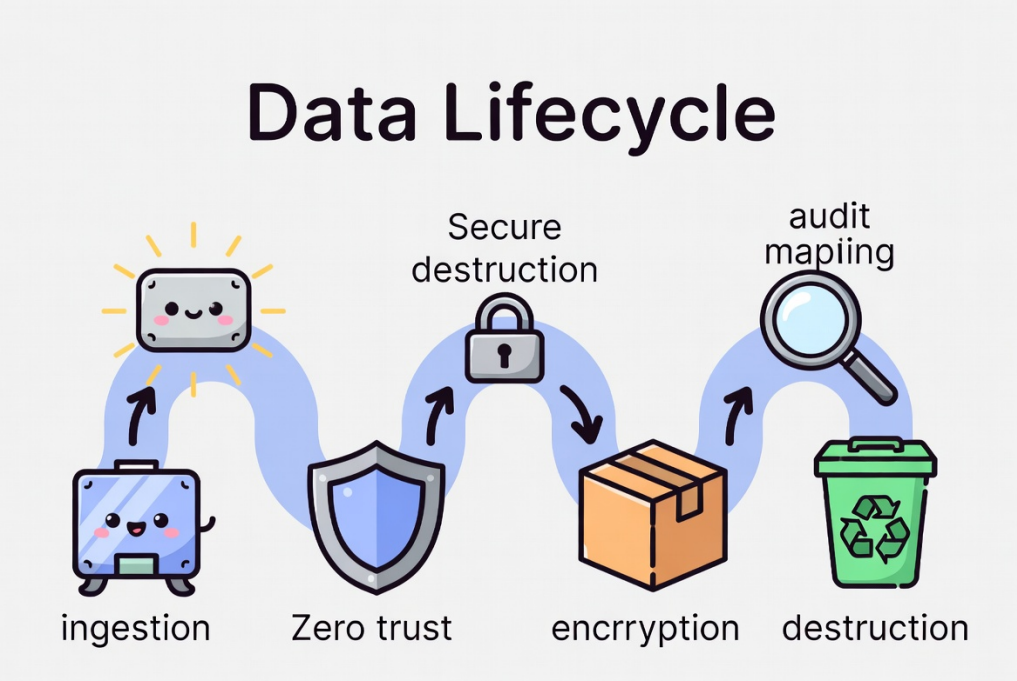

Figure 10.2: The Privacy-Enhanced Data Lifecycle. This model demonstrates how IG controls—including Zero Trust verification, encryption, and automated retention—must be embedded at every stage of the data's journey to satisfy "Privacy by Design" requirements.

10.2.2 - Implementation steps (what auditors will expect)

1. **Define your "protect surface"** of crown-jewel datasets (customer, patient, employee, financial, R&D) and the applications/services that touch them; map identities (human and workload) that legitimately need access. Document as an IG asset register. [usercentrics.com]
2. **Establish the decision plane**: implement an identity provider (IdP), strong authentication (MFA), device posture signals, and a policy engine/administrator capable of evaluating risk claims per request. Tie decisions to record-keeping for demonstrable accountability. [usercentrics.com]
3. **Place enforcement points** close to resources: gateways, service-mesh sidecars, or host agents enforcing session policies, mutual TLS, and fine-grained authorization. [usercentrics.com]
4. **Instrument logging and telemetry** into a SIEM and data-protection monitoring stack, linking events to data classification labels and retention rules. This provides the evidence that IG and regulators ask for. [usercentrics.com]
5. **Iterate**: start with one high-value data domain (e.g., HR or payments), run tabletop exercises, fix policy gaps, then expand by domain. Train help desk, legal, and privacy teams on what logs mean and how to act on risk alerts. [usercentrics.com]

10.2.3 - Table — Zero-trust principles (IG lens)

Principle	What it means operationally	IG benefit
Verify explicitly	IdP + MFA + device posture; per-request decisions by PE/PA/PEP	Evidencable access decisions mapped to policy; supports lawful/appropriate processing and security safeguards. [usercentrics.com]
Least privilege	Just-in-time, just-enough access; dynamic scopes; micro-segmentation	Minimizes exposure and over-collection/use; simplifies DPIA mitigations. [usercentrics.com]
Assume breach	Encrypt, monitor, segment; rapid revocation & isolation	Faster containment; better breach response and audit trails for regulators. [usercentrics.com]

IG integration tip: Embed ZTA requirements in your **Records of Processing** and **System Security Plans**: for each system, specify how authentication, authorization, logging, encryption, and segmentation are implemented and where **evidence** is stored (e.g., SIEM indexes, IdP logs). NIST's model gives you a neutral vocabulary for procurement, design reviews, and audits. [usercentrics.com]

10.3 - Encryption Strategies and Key Management

Encryption is the backbone of confidentiality and integrity—**at rest**, **in transit**, and increasingly **in use**—but it only works if **keys** are governed. In 2025–2026, two developments matter most: (1) **post-quantum cryptography (PQC)** standards and migration guidance, and (2) **in-use encryption** techniques (confidential computing and **homomorphic encryption**) maturing from pilots to targeted deployments. [infocon.arma.org], [it.uw.edu]

10.3.1 - At-rest, in-transit, in-use (and what auditors ask)

- **At rest:** Use strong algorithms (AES-256), envelope encryption, and **customer-managed keys (CMKs)** for regulated or cross-border workloads; log all key use; enforce separation of duties between data owners and key custodians. Map to privacy law requirements for appropriate safeguards and transfer controls. [linkedin.com]
- **In transit:** TLS 1.2+ (prefer 1.3), Perfect Forward Secrecy; mutual TLS for service-to-service. Document cipher suites and certificate management so your security assertions are verifiable. [linkedin.com]
- **In use: Confidential computing** enclaves and **homomorphic encryption** protect data during computation. While fully homomorphic encryption (FHE) still incurs overhead, government and industry work (e.g., DARPA DPRIVE; emerging FHE consortia) signal realistic near-term use in high-sensitivity analytics. Use selectively where policy or threat models demand computation without decryption. [bakerdonelson.com], [ibm.com]

10.3.2 - Table — Encryption types (comparison)

10.3.3 -

Type	What it protects	Performance/complexity	Typical uses
At rest	Storage media/database files	Low overhead; mature tooling	Default for databases, object stores; CMK for sovereignty/transfer controls. [linkedin.com]
In transit	Data on the wire	Low overhead with TLS 1.3	Web/API traffic; service mesh mTLS; partner connections. [linkedin.com]
In use (enclaves)	Memory/CPU during compute	Medium overhead; platform-specific	Analytics on sensitive data in trusted execution environments. [linkedin.com]
In use (FHE)	Compute on ciphertexts	High overhead; improving with hardware/standards	Targeted analytics where plain-text computation is unacceptable (health, finance, space). [bakerdonelson.com]

10.3.4 - Key management (what “good” looks like)

- **Lifecycle control:** Centralize keys in HSMs/KMS; define creation, rotation, revocation, archival, and destruction; maintain a **Cryptography Bill of Materials (CBOM)** that maps keys and algorithms to systems and data sets. [infonext.arma.org]
- **Access and separation of duties:** Role-based access to keys; dual control for export/destruction; continuous logging; automated alerts on anomalous key use. [linkedin.com]
- **Crypto agility:** Track algorithms and key lengths so you can change quickly (e.g., deprecate SHA-1, ECB mode; plan for PQC). **NIST SP 800-131A** draft guidance and PQC implementation reports give migration timelines and acceptable replacements. [infonext.arma.org], [infocon.arma.org]

10.3.5 - Post-quantum preparation

NIST and the security community are pushing organizations to **start PQC migration** now. Practical steps include (1) inventory algorithms/keys, (2) prioritize long-lived data and high-exposure protocols, (3) pilot **hybrid** cryptography where supported, and (4) require vendor attestations of PQC roadmaps. The **CSA** highlights NIST guidance (e.g., **NISTIR 8547**) to move from standards to real implementations without disrupting operations. [infocon.arma.org]

10.4 - Global Privacy Laws and Compliance Landscape

Privacy compliance is no longer a single-region exercise. Your IG program must **harmonize** requirements across GDPR (including new cross-border and enforcement procedures), U.S. **CCPA/CPRA** (now with risk assessments, audits, and ADMT controls), Brazil’s **LGPD** (ANPD activity and transfer SCCs), India’s **DPDP** (2025 rules), China’s **PIPL** (cross-border transfer mechanisms), and sectoral laws like **HIPAA** (with 2026 Notice of Privacy Practices updates). [ispartnersllc.com], [crowell.com], [techchannel.com], [datagalaxy.com], [hipaajournal.com], [ai.igguru.net]

10.4.1 - GDPR (EU/EEA)

Beyond the familiar principles and rights, **enforcement mechanics** for cross-border processing were clarified by **Regulation (EU) 2025/2518**, which lays down **additional procedural rules** for cooperation and dispute resolution among supervisory authorities (Articles 60, 65 GDPR). For multinational programs, this affects how you plan investigations and timelines. Keep your Records of Processing, DPIAs, and breach files **complete and consistent** across establishments to avoid procedural delays. [ispartnersllc.com]

10.4.2 - CCPA/CPRA (California)

Finalized **2025 CPPA regulations** (effective **Jan 1, 2026**) impose **risk assessments** for high-risk processing, **annual cybersecurity audits** for certain businesses, and **consumer ADMT rights** (access and opt-out) for significant automated decision-making that replaces or substantially replaces human decisions; organizations must prepare submissions and attested audits on defined timelines (first submissions due 2028). Update privacy notices, DSAR processes, and model governance to reflect these rules. [crowell.com], [dandodiary.com]

10.4.3 - LGPD (Brazil)

Brazil's **ANPD** continues to mature the regime—issuing SCCs and a transfer framework (Resolution 19/2024) and publishing a **2025–2026 Regulatory Agenda** prioritizing DPIAs, AI/biometrics, and children's data, with an emphasis on interoperability with OECD/EDPB guidance. Controllers using prior clauses must replace them per deadlines; expect growing DPIA scrutiny and sector guidance. [techchannel.com], [vorlon.io]

10.4.4 - DPDP (India)

India's **DPDP Act 2023** and **DPDP Rules 2025** phase in consent, notice, breach reporting, significant data fiduciary obligations (e.g., DPO in India, DPIAs), and multilingual notice requirements. Apply DPDP to digital personal data processed in India or related to offering goods/services to people in India. Harmonize with GDPR where possible, but note DPDP's distinct consent and notice mechanics. [datagalaxy.com], [cdn.prod.w...-files.com]

10.4.5 - PIPL (China)

China's **PIPL**, together with 2025 **Network Data Security Management Regulations** and **2026 PIP Certification Measures**, tightens cross-border transfer choices (CAC security assessment, **SCC filing**, or **PIP certification**) and clarifies large-platform obligations, consent, and important-data concepts. Compliance now requires data flow audits, contract updates to CAC SCC templates, and—depending on volumes—certification or security assessment before export. [hhs.gov], [hipaajournal.com]

10.4.6 - HIPAA (U.S. health)

HIPAA remains sectoral but evolving: while parts of the 2024 reproductive-health privacy rule were vacated in 2025, **Notice of Privacy Practices (NPP)** changes tied to **42 CFR Part 2 (SUD records)** remain and are due **February 16, 2026**. Expect modernization to the Security Rule from the 2024 NPRM (more specific cybersecurity controls) to land in 2026. Align your NPP updates and security gap assessments now. [ai.igguru.net], [dataversity.net]

10.4.7 - EU AI Act (privacy interface)

The **EU AI Act** (Regulation 2024/1689) is **risk-based** and sits alongside GDPR: high-risk AI systems require risk management, data governance/testing, documentation, transparency, and human oversight; **transparency** rules (e.g., labelling synthetic content, disclosure when interacting with AI) apply more broadly and matter for privacy notices and user expectations. Build your AI Register and map obligations to your GDPR controls. [download.pli.edu], [slideserve.com]

10.4.8 - Table — Global privacy law matrix (high level)

Table 10.3: Global Privacy and AI Law Updates (2025–2026). Key regulatory changes across major jurisdictions and their immediate impact on Information Governance workflows.

Jurisdiction	Scope & extraterritoriality	Notable 2025–2026 updates	What IG must do
GDPR (EU/EEA)	Controllers/processors handling EU personal data; extraterritorial reach	**Reg. 2025/2518** on cross-border enforcement procedures	Standardize DPIAs/records; prepare for smoother cooperation/dispute steps. [ispartnersllc.com]
CCPA/CPRA (CA)	For-profit orgs meeting thresholds; CA residents	**2025 regs** (risk assessments, audits, ADMT rights) effective 2026	Stand up risk-assessment workflow; schedule audits; update notices/DSARs. [crowell.com]
LGPD (BR)	Processing in Brazil or about individuals in Brazil	ANPD 2025–2026 Agenda; transfer SCCs deadlines	Replace SCCs; prepare DPIAs; monitor AI/biometric guidance. [vorlon.io], [techchannel.com]
DPDP (IN)	Digital personal data in/related to India	**2025 Rules** phasing in notices, SDF duties, breach reporting	Localize notices (languages); identify SDFs; plan DPIAs/DPO. [datagalaxy.com]

Jurisdiction	Scope & extraterritoriality	Notable 2025–2026 updates	What IG must do
PIPL (CN)	Processing PI of persons in China; extraterritorial	2025 **Network Data** Regs; 2026 **PIP Certification** Measures	Choose CAC assessment/SCC/PIP cert; update contracts; track thresholds. [hhs.gov], [hipaajournal.com]
HIPAA (US)	Health plans/providers/BA handling PHI	NPP updates due **Feb 16, 2026**; Security Rule NPRM	Update NPPs; prepare for stronger cybersecurity requirements. [ai.igguru.net], [dataversity.net]
EU AI Act	AI systems affecting EU markets/users	Phased obligations 2025–2027 (GPAI, high-risk, transparency)	Build AI Register; implement transparency & risk controls complementing GDPR. [download.pli.edu], [slideserve.com]

10.5 - Privacy by Design and Default

Privacy by design/default requires you to **bake** lawful, proportionate processing into systems **from the start**—not bolt it on. It pairs naturally with zero-trust (minimize privileges and flows) and encryption (minimize exposure) and is demanded explicitly or implicitly by GDPR, CPRA, LGPD, DPDP, and PIPL. Practically, it means **PIAs/DPIAs**, **data minimization**, **purpose limitation**, **consent and transparency**, and **rights fulfillment** engineered into processes and UIs. [purplesec.us], [crowell.com]

10.5.1 - PIA/DPIA in five steps

1. **Describe the processing:** purpose, legal basis, data categories, flows, recipients, retention, automated decisions, and locations (including cross-border). [purplesec.us]
2. **Assess necessity and proportionality:** Is the data collected the least necessary? Are notice, consent, and rights mechanisms adequate? [purplesec.us]
3. **Identify risks to individuals:** confidentiality, integrity, availability; bias/discrimination (if profiling/AI); surveillance concerns. [purplesec.us]
4. **Define mitigations:** zero-trust access, minimization, pseudonymization/anonymization, encryption, retention limits, transparency measures, human-in-the-loop for ADMT. [usercentrics.com], [crowell.com]
5. **Decide and record:** residual risk acceptance/escalation, sign-offs, and follow-up monitoring; maintain as a **record** for regulators and audits. [purplesec.us]

10.5.2 - Table — Privacy Impact Assessment (PIA/DPIA) checklist (condensed)

Table 10.3: Global Privacy and AI Law Updates (2025–2026). Key regulatory changes across major jurisdictions and their immediate impact on Information Governance workflows.

Area	Questions to answer	Evidence
Purpose & basis	What is the purpose? lawful basis? alternatives?	PIA narrative; legal memo; consent copy if used. [purplesec.us]
Data minimization	Are categories strictly necessary? can we pseudonymize?	Data inventory; minimization decision log. [purplesec.us]
Security controls	ZTA applied? encryption at rest/in transit/in use? key custody?	SSP/architectural diagrams; KMS/HSM configs; SIEM playbooks. [usercentrics.com], [linkedin.com]
Transfers	Any cross-border flows? which mechanism (SCCs, PIP cert, DPDP rules)?	TIAs; SCCs/PIP cert; vendor assurances. [techchannel.com], [hipaajournal.com]

Area	Questions to answer	Evidence
ADMT/AI	Is there ADMT or high-risk AI? notices? human oversight?	CPRA ADMT notice; AI risk profile; human-review SOPs. [crowell.com], [download.pli.edu]
Rights & retention	DSAR enablement? retention tied to purpose?	DSAR runbooks; ROPA; retention schedule. [purplesec.us]

10.5.3 - Data minimization & consent management

- **Minimization:** Collect only what is necessary for stated purposes; revisit forms and telemetry to strip nonessential fields; encode retention in systems and backups. These are core GDPR/LGPD/DPDP principles and reduce breach impact. [purplesec.us]
- **Consent:** Where required (e.g., DPDP children's data, certain marketing/ADMT contexts), ensure **clear, specific, revocable** consent with records and language localization (DPDP requires English or any of the 22 scheduled languages). Provide easy withdrawal and manage preference centers that reflect CPRA opt-outs and GDPR choices. [datagalaxy.com]

Rights fulfillment: Build DSAR workflows for access, correction, deletion, portability, **appeals of ADMT decisions** (where applicable), and account for identity verification, regional timelines, and exemptions (e.g., HIPAA designated record set vs. plan sponsor records). Train service teams to respond within statutory windows. [crowell.com], [ai.igguru.net]

Advanced Concepts in Privacy by Design, Governance Integration, and Cross-Functional Implementation

10.6 - Expanding Privacy by Design and Default: Operationalizing a Mature Framework

In practice, it is a governance operating model that must touch architecture, product development, security engineering, procurement, and business operations. Most privacy incidents arise not from malicious actions but from structural weaknesses—opaque data flows, inconsistent controls, ungoverned integrations, and unclear accountability. A mature Privacy by Design (PbD) program aims to prevent these systemic issues by embedding predictable, repeatable processes into every stage of the data lifecycle.

10.6.1 - 1. Deep Data Mapping: The Foundation for Privacy Assurance

Organizations cannot protect what they cannot see. While many companies perform a basic data inventory when required for GDPR's Article 30 Records of Processing Activities (ROPA), a mature IG program implements **continuous mapping** of data flows across structured repositories, unstructured stores, SaaS applications, APIs, machine-to-machine exchanges, and analytic pipelines.

- What data is collected
- Where it is stored
- Why it is collected
- Who it is shared with
- How long it is kept
- How it is secured
- What rights and obligations attach to it

This dynamic visibility is essential in complex environments where shadow IT, citizen development, and ad-hoc data exports are common. Enterprise discovery platforms can tag sensitive data in motion and at rest, producing lineage graphs that show transformations, aggregations, model training, and downstream data products. In a PbD program, teams must review updated maps during quarterly governance meetings and before launching new products.

10.6.2 - 2. Embedding Privacy in System Architecture

A traditional design pattern focuses on functionality first and compliance second. Privacy by design reverses this: functional decisions must be constrained by privacy requirements.

10.6.2.1 - Key architectural tactics include:

- **Data-minimizing schemas:** Only include fields that are strictly necessary; avoid free-text fields that invite over-collection; segregate identifiers from behavioral or transactional datasets.
- **Access-controlled microservices:** Each service should access only the subset of data relevant to its function, enforced through identity-aware API gateways or sidecar proxies.
- **Pseudonymization at ingestion:** Replace direct identifiers with surrogate keys before data enters analytics or ML pipelines.
- **Context-aware logging:** Only log what is required for security, auditing, and troubleshooting—exclude sensitive fields whenever possible.
- **Environmental separation:** Use isolated dev/test data that is masked or synthetic to prevent accidental exposure of production data.

These techniques reduce regulatory risk, enable easier DSAR fulfillment, and support rapid incident containment.

10.6.3 - 3. Design Reviews and Data Protection Sign-Off

Teams must answer specific, structured questions:

- Does the system collect new categories of personal data?
- Are existing purposes expanded?
- Will automated decision-making be introduced?
- Are cross-border transfers added or changed?
- Does the system increase monitoring of individuals (employees or customers)?
- Are minors' data or biometric identifiers involved?

Depending on the answers, the Data Protection Officer (DPO) may require a DPIA expansion, a formal technical review, or new contractual controls with vendors.

10.6.4 - 4. Integrating Privacy into Product Management

Product managers must treat privacy controls as **first-class product features**, not backend obligations. Roadmaps should allocate capacity for:

- Building customer-facing preference centers
- Providing clear opt-out mechanisms
- Supporting granular consent
- Supporting multiple legal bases across global markets
- Offering self-service data access and deletion

In regions covered by CCPA/CPRA, LGPD, or DPDP, product teams must ensure that UI/UX flows support jurisdiction-specific requirements such as opt-out links, children's data verifications, or parental consent.

10.6.5 - 5. Model Governance for High-Risk Use Cases

If a system involves profiling, behavioral inference, or automated decision-making, the PbD framework must integrate **algorithmic accountability**. This includes:

- Explaining the logic of decisions when required by law
- Maintaining training data provenance
- Documenting model updates and performance metrics
- Establishing fairness criteria and thresholds
- Implementing monitoring pipelines to detect model drift

Privacy by design becomes **algorithmic governance by design**, ensuring decisions remain lawful, transparent, and explainable.

10.6.6 - 6. Strengthening Purpose Limitation

Purpose creep—expanding the use of data beyond the original purpose—is one of the most common privacy failures. GDPR, LGPD, DPDP, and PIPL all evaluate organizations on whether they articulate **specific, explicit, legitimate** purposes at the time of collection.

IG must enforce:

- Purpose statements in system design documents
- Restrictions in data access tools (e.g., role-based filters)
- Training for analysts on permissible uses
- Internal policies requiring documented justification for every new data use case

Whenever a new purpose is proposed, PbD requires a **compatibility assessment**:

- Is the new purpose aligned with the original expectation?
- Does it increase risks to individuals?
- If compatibility is uncertain, the proposed use must undergo a formal DPIA.

10.6.7 - 7. Advanced Retention Strategies

Privacy by default requires that data be kept **no longer than necessary**, but real-world systems are full of retention challenges:

- Backups containing years of historical data
- Logs replicated across multiple SIEM systems
- ML feature stores retaining intermediate computations
- Data lakes lacking per-table retention policies
- Vendor-hosted platforms with opaque retention practices

A mature IG program implements:

- Event-based retention (e.g., delete accounts 24 months after closure)
- Automated retention enforcement using lifecycle policies in object stores
- Retention-aligned encryption keys (destroy the key, render the data useless)
- Vendor retention controls and contractual SLAs

Retention is fundamentally a **risk mitigation tool**: less data means fewer breach consequences, fewer DSAR burdens, and easier compliance audits.

10.6.8 - 8. Rights-Enablement as a System Requirement

Privacy by default includes enabling individuals to exercise rights without friction. Many organizations treat DSAR fulfillment as a manual process, but in modern architectures, DSAR obligations require built-in support:

- **Searchable subject indexes** connecting identifiers to records across systems
- **Deletion hooks** in microservices
- **Export formats** for portability (JSON, CSV, machine-readable structures)
- **Identity verification workflows** that are secure yet user-friendly

For employees, DSAR processes must be handled carefully to avoid disclosing privileged or other employees' data; PbD ensures systems separate personal data from operational context where possible.

10.6.9 - 9. Vendor and Third-Party Privacy Controls

Modern enterprises rely extensively on vendors—cloud providers, SaaS platforms, data processors, analytics partners. Privacy by design mandates **vendor governance**:

- Data Processing Agreements (DPAs) aligning with GDPR, LGPD, DPDP, or PIPL
- Standard Contractual Clauses (EU) and SCC equivalents (Brazil ANPD templates, China CAC templates)
- Security and privacy questionnaires covering retention, subcontractors, incident reporting, encryption, and access controls
- Right-to-audit clauses
- Vendor SOC 2, ISO 27001, and compliance certifications
- Continuous monitoring of vendor performance and breaches

Privacy incidents increasingly originate from third parties. Strong vendor governance is therefore a cornerstone of PbD.

10.7 - Building a Cross-Functional Privacy Governance Model

PbD cannot be owned exclusively by the privacy office or legal team; it requires cross-disciplinary participation.

10.7.1 - Stakeholder Responsibilities

-
- **Security Team:** Implement zero trust, encryption, monitoring, incident response, vulnerability management, identity controls.
- **Data Teams (Engineering, Architecture, Analytics):** Enforce data minimization, lineage, quality, retention, pseudonymization, and ML governance.
-
- **Procurement:** Enforce vendor privacy reviews and contract clauses.
- **Compliance / Audit:** Validate evidence, measure control effectiveness, schedule internal audits.
- **Executive Leadership:** Ensure tone-from-the-top and allocate resources.

10.7.2 - Governance Structures

A mature program uses several boards/committees:

- **Data Protection Steering Committee** (senior cross-functional oversight)
- **AI Governance Board** (model risk, fairness, high-risk AI decisions)
- **Security Architecture Review Board** (ZTA, encryption, access models)
- **Vendor Governance Committee** (critical processors, transfer requirements)

10.8 - The Role of Culture and Training

Training must be:

- **Role-based:** engineers learn pseudonymization and secure coding; analysts learn minimization and access control; HR learns sensitive data rules; marketing learns consent requirements.
- **Scenario-based:** real breach examples, dark pattern avoidance, AI misuse cases.
- **Continuous:** micro-trainings, just-in-time prompts in tools, quarterly refreshers.
- **Measured:** training completion tied to performance metrics.

Organizations with strong cultures see fewer mistakes, faster incident reporting, and higher customer trust.

IG must maintain metrics and dashboards demonstrating progress.

10.8.1 - Key Metrics

1. Percentage of systems with completed DPIAs
2. Number of high-risk processing activities approved/revised
3. Incidents involving over-collection or purpose creep
4. DSAR turnaround times and volumes
5. Vendor assessments completed
6. AI model risk evaluations completed
7. Training completion rates
8. Policy exceptions requested and resolved

10.8.2 - Evidence Sets

- Architecture diagrams showing data flows
- Logs of DPIA decisions and mitigations
- Retention and deletion evidence
- Access logs and security event records
- Vendor due diligence and SCC/PIP certification documentation
- Training archives
- Model documentation for AI systems
- Incident response records

10.9 - Future Enhancements in Privacy by Design

The next decade will transform PbD as new technologies emerge:

10.9.1 - 1. Privacy-Enhancing Technologies (PETs)

Organizations will increasingly adopt PETs such as:

- Federated learning
- Secure multiparty computation
- Differential privacy
- Homomorphic encryption (as it becomes more efficient)
- Trusted execution environments

These technologies support analytics without exposing raw personal data, aligning with privacy by design.

10.9.2 - 2. Policy-as-Code

- Service mesh policies requiring mTLS
- Data pipeline code that automatically pseudonymizes identifiers
- Automated DSAR hooks triggered via APIs

This reduces manual errors and makes privacy and security consistent at scale.

10.9.3 - 3. AI-Driven Privacy Monitoring

Machine learning systems will help detect:

- Anomalous access patterns
- New data fields appearing in systems (indicating creep)
- Misconfigured cloud resources
- Vendor risks
- Inference attacks on AI models

10.10 - Data Protection in AI Systems

Artificial intelligence systems—especially large-scale machine learning (ML) and generative AI (GenAI)—introduce **new privacy and security risks** across the entire data lifecycle: training, inference, model storage, and system deployment. Most global laws now explicitly or implicitly regulate AI systems that process personal data. GDPR's automated decision-making rules, CCPA/CPRA's ADMT governance requirements, Brazil's ANPD AI guidance, India's DPDP Rules 2025, and China's PIPL all intersect with the **EU AI Act**, which applies risk-based obligations for AI systems deployed in the EU. [crowell.com], [blog.rsisecurity.com]

10.10.1 - AI-related privacy risks

1. **Training data exposure** – AI models often learn from large datasets containing personal or sensitive data. Regulators warn that improperly curated data may violate principles of minimization, fairness, and lawful basis under GDPR, LGPD, and DPDP. The EU AI Act emphasizes **data governance and quality**, requiring training datasets to be free from errors and bias, and to respect privacy obligations. [blog.rsisecurity.com]
2. Inference privacy leaks – Models may inadvertently memorize training data, resulting in risks such as membership inference, attribute inference, or sensitive data reconstruction.
3. **Automated decision-making (ADMT)** – CPRA regulations finalized in 2025 give consumers **new opt-out and access rights** surrounding ADMT systems used for significant decisions like employment, credit, housing, and healthcare. Businesses must provide **pre-use notice**, allow requests for explanations, and conduct **risk assessments**. [crowell.com]
4. **Synthetic content privacy** – The EU AI Act's Article 50 requires **labeling of artificially generated content**, including deepfakes and synthetic media. Providers of systems generating synthetic audio, video, image, or text must implement machine-readable markings. [slideserve.com]

10.10.2 - Controls and governance

- **AI data governance framework:** Inventory training datasets, document data provenance, conduct data quality checks, and apply minimization techniques (pseudonymization, de-identification). Map training data sources to legal bases.
- **Model risk scoring:** Categorize models by sensitivity (high-risk if used for biometric identification, credit, recruitment, medical triage, etc.). For EU deployments, consider whether the AI Act's high-risk classification applies.
- **Technical protections:**

- Differential privacy during training
- Secure enclaves or confidential computing for inference
- Encryption of model weights and embeddings
- Access-controlled model registries

- Human oversight: Required by GDPR (Article 22) and mandated for high-risk AI under the EU AI Act.
- **Documentation:** Maintain **model cards** and **system cards** describing training data, limitations, risks, and intended uses. Preserve logs for access, inference, and drift.
- **Transparency:** Provide notices when individuals interact with AI systems or when decisions meaningfully affect them; label synthetic content to meet AI Act requirements. [slideserve.com]

10.11 - Tools and Technologies for Privacy/Security IG

Modern IG programs rely on integrated tooling to govern data across cloud services, apps, endpoints, and AI stacks.

10.11.1 - Privacy and data governance platforms

- **Microsoft Purview:** Provides data classification, sensitivity labels, DLP (Data Loss Prevention), eDiscovery, insider risk management, and records management. Useful for harmonizing privacy/security controls across Microsoft 365, Azure, and hybrid environments.
- Widely used for GDPR/CCPA/LGPD compliance.
- BigID: Applies advanced data discovery, classification, and identity-aware scanning to structured/unstructured data.

10.11.2 - Security monitoring and incident response

- **SIEM (Security Information and Event Management):** Systems like Microsoft Sentinel, Splunk, or Elastic provide centralized log ingestion, correlation, anomaly detection, and forensics. Zero-trust requires comprehensive telemetry from identity systems, endpoints, service meshes, and cloud logs to feed into the SIEM.
- **SOAR (Security Orchestration, Automation, and Response):** Automates incident workflows (e.g., quarantine a compromised identity, rotate keys, revoke tokens).

10.11.3 - Encryption and key management

- **Cloud KMS/HSM systems** (Azure Key Vault, AWS KMS, Google Cloud KMS) provide key generation, rotation, access control, logging, and integration with cloud services.
- **PQC readiness tools:** Vendors increasingly provide PQC compatibility checks and test environments aligned to NIST's PQC standards such as ML-KEM and ML-DSA. [infocon.arma.org]

10.11.4 - AI governance tools

- **Model monitoring platforms** for drift, fairness, output anomalies, and data leakage (e.g., Azure AI Content Safety, AWS Clarify).
- **Data anonymization tools** to prepare privacy-safe training datasets.
- **Synthetic data generators** with DP safeguards for test and development environments.

These tools must be integrated into IG workflows so that policies flow down into technical enforcement, and evidence flows up for audit readiness.

10.12 - Real-World Case Studies

10.12.1 - Case Study 1 — AI System Privacy Failure (Hypothetical)

A retail chain deployed a GenAI-based customer service assistant trained on historical customer emails. The model began leaking personal data (order histories, addresses) in responses to unrelated users—an inference privacy failure.
Root causes:

- Training dataset contained unmasked personal information
- Lack of differential privacy
- No pre-release privacy testing
  Outcome:
-

- Required breach notifications and retraining on sanitized datasets
  Fixes:
-
- Add human review workflows for sensitive interactions

10.12.2 - Case Study 2 — Zero-Trust Implementation Success (Realistic)

A healthcare provider facing increasing ransomware threats adopted NIST SP 800-207 ZTA. The organization segmented EHR systems, adopted MFA everywhere, implemented continuous device posture checks, and enforced mTLS between microservices.
Results:

- A phishing attack compromised a clinical workstation, but lateral movement was blocked; no PHI accessed
- HIPAA auditors praised identity logs, segmentation diagrams, and encryption controls as “mature and evidence-ready”
  Key IG takeaway: Zero trust converted broad perimeter policies into enforceable, auditable controls directly mapped to regulatory safeguards. [usercentrics.com]

10.12.3 - Case Study 3 — Global Compliance Challenge (Multinational SaaS)

A SaaS analytics provider served customers in the EU, US, Brazil, India, and China. IG performed a global data flow map and discovered:

- EU-to-US transfers lacked supplementary measures under GDPR
- No SCC updates for Brazil’s LGPD Resolution 19/2024
- India DPDP notice templates missing localization in required languages
- PIPL cross-border transfer volumes exceeded CAC thresholds requiring security assessment
  Solution:
- Implement encryption with CMKs, pseudonymization, and EU-based compute isolation
- Register a DPO in Brazil and India; upgrade SCCs to ANPD templates
- Perform CAC security assessment and apply for PIP Certification
  Impact:
- Prevented fines; reduced global compliance gaps; established a unified IG framework harmonizing all jurisdictions. [techchannel.com], [datagalaxy.com], [hhs.gov]

10.12.4 - Case Study 4 — HIPAA NPP Update Failure (Based on 2025–2026 Rule Changes)

A health plan neglected to update its **Notice of Privacy Practices (NPP)** to reflect new **42 CFR Part 2** rules governing SUD records. During an OCR investigation, missing notices were deemed non-compliant.
Outcome:

- Required corrective action plan, NPP redistribution, and updated BAAs
  Lesson:
- Even when broader rules are vacated, remaining NPP obligations (due **February 16, 2026**) still apply. [ai.igguru.net]

10.13 - Challenges and Best Practices

10.13.1 - Key Challenges

1. Regulatory fragmentation: GDPR, CCPA/CPRA, DPDP, LGPD, PIPL, HIPAA, and the EU AI Act each impose unique requirements.
2. **Shadow IT and shadow AI:** Teams may adopt tools or build models outside formal governance, exposing unmonitored data flows.
3. **Data discovery gaps:** Without full visibility, organizations cannot satisfy minimization, retention, DSAR, or cross-border obligations.
4. **Security-privacy tension:** Encryption, ZTA, and minimization may impede analytics unless carefully designed.
5. **Legacy environments:** Old systems lack modern identity, logging, or encryption capabilities.

10.13.2 - Best Practices

- **Create a unified control catalog** mapping all global privacy/security laws to NIST CSF, ISO 27001, and ZTA principles.
- **Automate data discovery/classification** using tools like BigID or Purview.

- **Implement strong identity foundations** (MFA, SSO, conditional access, device trust).
- **Maintain a global data map and transfer register** updated quarterly.
-
- **Design secure defaults**: encryption at rest/in transit, tight retention schedules, CMKs for sensitive cross-border processing.
- **Prepare for PQC migration** now by inventorying algorithms, keys, and cryptographic dependencies. [infocon.arma.org]
- **Adopt SIEM/SOAR workflows** for rapid detection, investigation, and evidence preservation.
- **Train employees annually** on privacy/security requirements—especially AI-related risks, DSAR workflows, and incident reporting.

10.14 - Future Outlook

10.14.1 - 1. Zero trust becomes the default

Most large enterprises and regulated sectors will treat ZTA as a baseline requirement, integrated with identity, networking, and data-layer controls. Regulatory frameworks increasingly reference zero-trust principles implicitly through “continuous verification” and “least privilege” requirements. [usercentrics.com]

10.14.2 - 2. PQC transitions accelerate

With NIST PQC standards finalized (ML-KEM, ML-DSA, SLH-DSA), governments and cloud providers will push for adoption in TLS, code signing, VPNs, and databases. Organizations that built crypto-agility in advance will adapt smoothly; laggards will face incompatible legacy systems and compliance risks. [infocon.arma.org]

10.14.3 - 3. AI governance becomes mainstream

The EU AI Act’s risk-based framework may inspire similar global regulations. Privacy teams will expand into **AI governance offices**, responsible for model registries, testing protocols, synthetic content marking, and human oversight. AI transparency expectations will rise worldwide. [download.pli.edu]

10.14.4 - 4. Cross-border data governance tightens

Brazil, India, and China continue strengthening transfer rules; EU–US adequacy decisions remain volatile; and organizations adopt **regional data zones** with encryption, access restrictions, and tokenization to maintain compliance. [techchannel.com], [hhs.gov]

10.14.5 - 5. Privacy + security converge

IG will increasingly oversee unified “Data Protection Offices” combining privacy, security, and AI governance, supported by automated monitoring and continuous compliance dashboards.

10.15 - Learning Objectives

- Understand how IG integrates privacy and security using zero-trust principles, encryption, and data-centric governance.
- Compare global privacy laws (GDPR, CCPA/CPRA, LGPD, DPDP, PIPL, HIPAA) and identify actionable compliance components.
- Leverage modern tools (Purview, OneTrust, BigID, SIEM) to operationalize data protection and evidence collection.

10.16 - Key Takeaways

- Zero trust is now foundational to privacy and security IG; NIST SP 800-207 provides the reference model. [usercentrics.com]
- Encryption strategies must include at-rest, in-transit, and in-use protections, supported by strong key management and PQC planning. [infocon.arma.org]
- Global laws are converging on stronger rights, transparency, risk assessments, and cross-border controls.
- AI introduces new privacy risks requiring specialized governance, testing, and documentation.

10.17 - Discussion Questions

1. What technical and procedural controls are most important for preventing privacy leakage in AI training and inference pipelines?
2. If your organization had to migrate to PQC in two years, what systems and data would you prioritize first—and why?

10.18 - Further Reading

- NIST SP 800-207 Zero Trust Architecture: https://csrc.nist.gov/pubs/sp/800/207/final
- NIST PQC Standards (FIPS 203, 204, 205): https://www.nist.gov/pqc
- European Commission GDPR Portal: https://gdpr.eu
- China PIPL (English resources): https://www.lw.com (search PIPL alerts)
- EU AI Act Text: https://artificialintelligenceact.eu

10.18.1 - **The Nerdy Example**

Captain America: The Winter Soldier (2014): The Ethics of Predictive Governance. Project Insight represents the ultimate failure of "Privacy by Design." By weaponizing data for pre-emptive security without transparency or human-in-the-loop oversight, it mirrors the 2026 risks of unregulated Automated Decision-Making Technology (ADMT). Modern IG must ensure that security "insight" never bypasses fundamental privacy rights. (Image: AI)

11: IG for Digital Communications and Collaboration

11.1 - Introduction

In 2026, the way organizations communicate has changed as dramatically as the places they keep their records. Email is still the backbone of business correspondence, but decision-making also happens in **chat threads**, **short-form video clips**, **collaboration channels**, **mobile apps**, and even through signals exchanged by **IoT devices**. Each of these mediums creates **business records** and **risk**—they carry sensitive information, may be subject to discovery or regulatory access, and can be misused for fraud or disinformation. Information Governance (IG) for digital communications means designing policies, controls, and training that ensure communications are **purposeful, secure, retrievable, and compliant**—without smothering productivity.

Three forces shape today's playbook. First, **zero-trust** thinking has moved from networks into collaboration: access is continuously verified, least-privilege is enforced, and every message can be logged as evidence when required. NIST's Zero Trust Architecture (ZTA) provides a widely adopted blueprint for this mindset across cloud and on-prem resources, including collaboration systems. Second, **compliance and e-discovery** have "caught up" with collaboration tools: email-style retention, legal hold, and search are now expected for Teams, Slack, and other platforms—though important nuances remain. Third, the communications surface now includes **mobile and IoT** endpoints, requiring enterprise device management, threat defense, and BYOD guardrails to meet privacy and security obligations. NIST's mobile guidance (SP 800-124 Rev.2) anchors today's enterprise practices for securing smartphones and tablets used for work. [csrc.nist.gov] [learn.microsoft.com] [csrc.nist.gov]

This chapter offers a practical framework for integrating IG across email, social media (X/TikTok), collaboration platforms (Teams/Slack), and mobile/IoT communications. You'll find policy templates, control checklists, platform comparisons, and case studies you can adapt for class projects or real organizations. We ground the guidance in credible sources that auditors, counsel, and cybersecurity teams already recognize: **NIST** for security controls and identity, **Cloud Security Alliance** for cloud responsibility and SaaS baselines, **IAPP/EDPB** for privacy, and **ARMA** for IG maturity. [csrc.nist.gov], [cdn.prod.w...-files.com], [iapp.org], [todaysgene...ounsel.com]

11.2 - Email Governance

Email remains the most universal business communications channel—and the one most often implicated in **regulatory requests, HR matters, and litigation**. A defensible email program blends **retention**, **archiving**, **monitoring/anti-abuse**, and **incident response** with auditable evidence.

11.2.1 - Why email is still special for IG

- **Regulatory discoverability:** Email is routinely requested in investigations; courts and agencies expect organizations to locate and produce it quickly.
- **High-value evidence & risk:** Email threads capture approvals, intent, and timelines. They also carry phishing payloads and can be spoofed to authorize fraudulent payments. NIST highlights phishing's evolution (including AI-assisted lures) and recommends layered defenses and staff education. [nist.gov]
- **Control coverage:** The NIST SP 800-53 control catalog provides technology-agnostic controls that apply to messaging systems —e.g., **SI-8 (Spam Protection)** and related email hygiene measures—while NIST's overlay repository includes an **Email Messaging Systems Overlay** for tailoring controls to email. [csf.tools], [csrc.nist.gov]

11.2.2 - Retention and archiving

Retention defines how long to keep messages; **archiving** is how you store and retrieve them. The principle is simple: **retain what you must, dispose of what you should not keep**. For Microsoft 365 tenants, **Purview Data Lifecycle Management** (retention policies/labels) can apply rules to Exchange mailboxes and support holds when litigation is anticipated. Teams chat and channels are also covered (see next section), but email retention remains a separate policy scope and must be set explicitly. [learn.microsoft.com]

Practical steps

1. **Map obligations:** Identify legal/regulatory retention baselines (e.g., tax, employment, sector rules). Align schedules with enterprise retention policies.
2. **Implement policy + technology:** Use centralized retention policies to keep email for a defined period and **delete** after expiry unless on hold. Configure immutable archives where required. [learn.microsoft.com]

3. **Enable legal holds:** When litigation is reasonably anticipated, place custodial holds to suspend deletion; maintain a case record describing scope and justification. [syscloud.com]
4. **Document exceptions:** If certain functions (e.g., board communications) require longer retention, use label-based exceptions with DPO/Legal approval and written rationale. [learn.microsoft.com]

11.2.3 - Monitoring and anti-phishing

Email security should combine **technical controls** (gateway filtering, DMARC/DKIM/SPF, sandboxing, link protection, malware detection) with **human factors** (simulated phishing, training, and clear reporting). NIST's guidance for small and large organizations stresses that phishing now spans **email, SMS, voice, and social media**, and that employees need simple "report a phish" paths and just-in-time education. **SP 800-53 SI-8** codifies spam protection at entry/exit points with automatic updates; CSF 2.0 provides an overall risk-management context for deploying these defenses. [nist.gov], [csf.tools], [nist.gov]

Governance note: Retention is **not** backup: if a mailbox falls outside policy or a hold, "deleted" email might be unrecoverable for legal purposes. Your IG policy should explain the **difference between backup and retention** and instruct staff to use approved archives for records, not PST files or personal mailboxes. [magnetclicks.com]

11.3 - Social Media and Public Platforms (X/TikTok)

Public platforms are powerful for outreach and brand, but they pose IG risks: **unapproved disclosures**, **records sprawl**, **privacy exposure**, and **regulatory obligations**. Social content can be discoverable, particularly when it communicates business positions or customer service actions.

11.3.1 - Content classification and acceptable use

IG should classify social content (e.g., **Marketing**, **Customer Support**, **Executive Communications**, **Recruiting**) and define what must be retained as **records**. ARMA's contemporary thought leadership and events emphasize the need to embed IG in modern channels and treat them as potential **official records** tied to retention schedules. [magazine.arma.org], [govevents.com]

Acceptable use should forbid posting **confidential or regulated data**, require **brand/legal review** for high-risk topics, and mandate **screenshots/exports** for significant business interactions (e.g., customer complaint resolution on X). Privacy-by-design requires that programs avoid collecting more data than necessary and comply with GDPR/CCPA transparency rules when using social messaging for support. IAPP toolkits and trackers can help align policy with evolving privacy laws. [iapp.org]

11.3.2 - **Platform-Specific Risks and Governance Controls (TikTok, X, Reddit)**

Why Platform Nuance Matters Your Information Governance (IG) control set must reflect each social platform's unique design, data practices, and moderation models. These factors change how content spreads, what gets logged, how long it persists, and which laws apply. Treating all platforms identically leaves organizations vulnerable to platform-specific data collection practices, algorithmic virality, and emerging jurisdictional mandates (such as youth protections and consumer disclosures) [vorlon.io].

11.3.2.1 - **TikTok**

- **Algorithmic Amplification & Virality:** TikTok's recommendation systems and AI-generated content (AIGC) tooling can rapidly amplify accidental disclosures (e.g., whiteboards, security badges, or client data visible in the background). TikTok's own Digital Services Act (DSA) risk assessment outlines these systemic risks. IG should treat every frame as potentially public and mandate pre-publication checks (e.g., "no filming in restricted areas," blur/background-safe reviews) [jdsupra.com].
- **Data Collection Concerns:** External analyses highlight extensive device/telemetry collection. Businesses must factor these privacy implications into their risk assessments for official accounts and employee devices, requiring work-profile isolation on mobile devices and verifying regional data-handling disclosures [sf16-va.ti...tokcdn.com], [tiktok.com], [linkedin.com].
- **Personal Information & Takedowns:** TikTok's Community Guidelines prohibit sharing high-risk personal information (home addresses, credentials) and ban phishing or malware. Your playbook should map these rules to your incident response: capture evidence, request removal citing the specific guideline, and record the artifacts [securitybo...levard.com].

11.3.2.2 - **X (formerly Twitter)**

- **The Speed and Permanence Paradox:** X's high-velocity posting and quote-tweet dynamics can spread sensitive statements globally before they are noticed internally, while third-party integrations routinely archive these posts.
- **Business Records:** Treat official X posts and Direct Messages (DMs) dealing with customer issues as business records when they commit the organization to action. Maintain a capture workflow (native exports or third-party archiving) to satisfy

retention and eDiscovery requirements.

- **Approval Workflows:** Apply dual approval for high-risk topics (finance, health, M&A, crisis response) and confirm compliant audience targeting when running youth-facing campaigns or promotions [vorlon.io].

11.3.2.3 - Reddit (and Community Platforms)

- **Community-Moderated Disclosures:** Subreddit threads often entice employees to discuss work under pseudonyms, but corporate affiliations are easily doxxed or inferred. Acceptable use policies must include a strict "no business disclosures on personal accounts" clause, supported by brand monitoring for mentions of confidential projects [vorlon.io].
- **Scraping & AI Training Risks:** Increasing data scraping across social platforms for analytics and AI training raises privacy and intellectual property concerns when employees upload original work, internal screenshots, or customer stories. IG should warn content creators about these risks, and legal teams must review licensing/consent for any user-generated content the brand republishes [natlawreview.com].

Youth and Jurisdictional Considerations Several jurisdictions now impose strict age restrictions and special duties for youth protection on social platforms. For example, Australia's eSafety regime (and 2025 social media age restrictions) classifies X, TikTok, YouTube, Twitch, and Reddit as age-restricted, expecting platforms to take "reasonable steps" to prevent under-16 accounts. Even if your organization is not consumer-facing, all staff outreach, marketing creative briefs, and recruiting communications must respect these local restrictions and validate that they do not encourage under-16 engagement [esafety.gov.au], [vorlon.io].

11.3.3 - Operational Controls and Incident Response Checklist

To effectively govern these channels, organizations should add the following practical controls to their IG framework:

- **Per-Platform Risk Register:** Document systemic risks for each platform (e.g., algorithmic amplification on TikTok, thread virality on X, community leaks on Reddit) and map them to specific mitigations like whitelists, geo-fencing, and takedown workflows [jdsupra.com], [securitybo...levard.com], [vorlon.io].
- **Age-Restriction Compliance:** Require a jurisdictional review before launching campaigns. Log exactly how you limited targeting, disabled DMs, or implemented additional warnings for youth protection [vorlon.io].
- **Mobile Isolation:** Mandate managed work profiles (MDM/EMM) and restrict copy/save functions for social media apps used on corporate or BYOD devices handling official accounts, tying this directly to the enterprise mobile policy [magazine.arma.org].
- **Evidence Capture & Takedown SOP:** Establish a coordinated takedown playbook between Legal, Comms, and Security. Standardize the process of capturing forensic evidence (screenshots/URLs) and citing platform-specific legal rules (e.g., personal-info guidelines) *before* submitting removal requests [securitybo...levard.com], [blog.rsisecurity.com].
- **Enterprise IR Integration:** Configure brand monitoring for unauthorized accounts or leaked content. Record all social media incident actions in a central case file and link them to the enterprise Incident Response plan (e.g., NIST IR/CSF functions) [nist.gov].

11.4 - Collaboration and Messaging Tools (Teams/Slack)

Teams, Slack, and similar platforms are now core to how organizations work. They also create **high-velocity evidence** and **regulated records**. Mature IG programs implement **retention**, **archiving**, **legal holds**, **DLP**, and **access controls** specifically for collaboration data.

11.4.1 - Why collaboration data is different from email

- **Non-custodial, conversational context:** Threads span many users and locations; single "messages" may not convey full meaning without context and reactions. Courts and regulators increasingly request **channel-level** records, and practitioners must preserve **threads, timestamps, edit history, reactions, and linked files**. The Sedona Conference's 2025 commentary spotlights these complexities. [everlaw.com]
- **Rapid platform change:** Storage architecture and compliance surfaces evolve—organizations must track vendor announcements to avoid gaps. For example, Microsoft announced a shift (rolling out in late 2025) to store **Teams private-channel messages** in the **team's group mailbox** rather than user mailboxes, affecting e-discovery/holds and retention scoping. [techcommun...rosoft.com], [nikkichapple.com]

12: Long-Term Preservation and Archiving

12.1 - Introduction

Long-term preservation is the discipline of keeping digital information **authentic, accessible, and usable for as long as it has value**—years, decades, or even centuries. It differs from backup or disaster recovery: backups restore systems to a recent state; preservation ensures that future users can still **render, interpret, and trust** specific digital objects and their context. The foundation for modern practice is the **Open Archival Information System (OAIS) Reference Model**, standardized as **ISO 14721** and maintained by the Consultative Committee for Space Data Systems (CCSDS). OAIS offers a common language (e.g., **Submission Information Package, Archival Information Package, Dissemination Information Package**) and a set of responsibilities an archive must fulfill (ingest, preservation planning, storage, access, and administration). The model was **updated in 2024–2025**; the current OAIS 3rd edition and its ISO adoption confirm its ongoing relevance across disciplines. [ccsds.org], [iso.org], [OAIS (ISO...rs updated]

Preservation is an **Information Governance (IG)** function as much as an archival one. IG aligns strategy, policy, and controls so that high-value information is **identified, appraised, protected, and preserved** while low-value information is disposed of legally and defensibly. National memory institutions and standards bodies provide proven guidance: the U.S. **National Archives and Records Administration (NARA)** publishes transfer and preservation guidance for permanent federal records; the **Library of Congress** maintains the **Sustainability of Digital Formats** knowledge base and the annually updated **Recommended Formats Statement (RFS)**; and the **Federal Agencies Digital Guidelines Initiative (FADGI)** provides technical targets for digitization and embedded metadata. These resources provide concrete criteria for selecting formats, documenting provenance, and validating quality over the long term. [archives.gov], [loc.gov], [digitizati...elines.gov]

In this chapter, we translate those standards into **practical steps** for IG programs: how to mitigate **digital longevity** risks; when to apply **migration, emulation, and normalization**; how to preserve **AI-generated content** with trustworthy provenance; which **standards and tools** to adopt; and which **checklists** to use to ensure you are ready for the next audit—and for the future.

12.2 - Challenges of Digital Longevity

Long-term preservation faces four interlocking risks: **format obsolescence**, **media degradation/bit rot**, **software dependency**, and **organizational discontinuity**.

12.2.1 - Format obsolescence

File formats evolve and sometimes disappear. If you cannot open yesterday's files in tomorrow's software, your information is **practically lost**. The **Library of Congress Sustainability of Digital Formats** program documents thousands of formats, including "sustainability factors" like openness, adoption, transparency, and self-documentation that correlate with long-term viability. IG teams use this analysis to favor formats such as PDF/A, TIFF, and XML for preservation masters. [loc.gov]

The **PRONOM** registry from The National Archives (UK) complements this by cataloging formats, versions, signatures, and related software, and powering **DROID** for automated identification. PRONOM's modernization work (linked data, signature updates) and search services help organizations **inventory** holdings and spot risky formats early, enabling proactive migration. [nationalar...ves.gov.uk], [nationalar...ves.gov.uk], [github.com]

12.2.2 - Media degradation and bit rot

Digital media—HDDs, SSDs, optical discs, tape—**fail over time**; silent corruption ("bit rot") can render files unreadable. While NIST's SP 800 series focuses on information security rather than archival science, its guidance highlights the need for **integrity mechanisms** and lifecycle controls. Preservation programs implement **fixity checks** (e.g., SHA-256) on ingest and on a schedule, verifying checksums and repairing from redundant copies when corruption is detected. Standards bodies and communities reinforce the principle of **lots of copies**: programs like **LOCKSS** operationalize peer-to-peer integrity polling and repair across many nodes so that corruption in one replica is detected and healed. [nist.gov], [lockss.org]

12.2.3 - Software dependency

Many digital objects depend on **specific software stacks** (codecs, runtimes, plug-ins) or even complete environments to render correctly. OAIS anticipates this by emphasizing **Representation Information**—the documentation and software necessary to render and understand content. Preservation planning thus includes **documenting dependencies**, acquiring or packaging software,

or planning migration/emulation routes for complex objects (e.g., multimedia projects, interactive spreadsheets, CAD/3D, or games). [ccsds.org]

12.2.4 - Organizational and legal discontinuity

Preservation is a **marathon**, not a sprint. Funding cycles, mergers, provider shutdowns, and policy changes can interrupt stewardship. NARA's federal transfer requirements and bulletins (e.g., ERM, Capstone email) try to reduce this risk by mandating **standards-based scheduling, transfer, and metadata** across agencies; likewise, LOC's RFS and FADGI provide continuity for the broader cultural heritage sector. IG's job is to **codify these requirements** into retention schedules, legal holds, and disposition overrides so that custody and obligations persist across organizational change. [archives.gov], [loc.gov], [digitizati...elines.gov]

12.3 - Core Preservation Strategies

Effective programs combine **migration, emulation, normalization**, and **open standards** selection—guided by OAIS and institutional policy.

12.3.1 - Migration

Migration moves content to newer formats or media to maintain renderability. Routine examples include normalizing office docs to **PDF/A-2 or PDF/A-4** for fixed-layout access and migrating master images to **TIFF** with embedded metadata. Migration is most successful when driven by **risk triggers** (e.g., PRONOM signals increased obsolescence) and controlled by **documented format policies** (e.g., Archivematica's Format Policy Registry). [pdfa.org], [archivematica.org]

Migration does not mean discarding originals. OAIS and NARA practice advise **retaining the source** (for authenticity and re-processing) while creating a normalized preservation master and access derivatives, with clear **provenance metadata** explaining each action. [ccsds.org], [archives.gov]

12.3.2 - Emulation

Emulation recreates the original software/hardware environment so digital objects run unchanged. It is valuable for **software-dependent** works (legacy multimedia, interactive art, scientific workflows). OAIS frames emulation as a preservation planning option when migration would **deform significant properties**. Community projects (e.g., Emulation-as-a-Service) use encapsulated environments with documented **Representation Information** to deliver authentic experiences. [ccsds.org]

12.3.3 - Normalization

Normalization converts many incoming formats into a **small set of preservation targets** with strong sustainability characteristics. Typical targets include **PDF/A** for documents, **TIFF** (or JPEG 2000 in some domains) for still images, and **WAV/BWF** for audio; where structure is key, **XML** (or JSON with schemas) expresses data in open, self-documenting form. Normalization is often **automated** and governed by policy (e.g., Archivematica FPR rules). [pdfa.org], [loc.gov], [w3.org], [archivematica.org]

12.3.4 - Favoring open standards and archival variants

- **PDF/A (ISO 19005)** ensures device-independent, self-contained documents. Multiple parts exist—**PDF/A-1** (2005), **PDF/A-2**, **PDF/A-3**, and **PDF/A-4 (2020)**—each building capabilities while constraining risky features. Conformance levels (A/U/B) balance accessibility, text extraction, and visual fidelity. [iso.org], [pdfa.org], [pdfa.org]
- **TIFF** remains widely preferred for master images because of its **openness and transparency**, especially when used with baseline features and embedded technical metadata; Library of Congress RFS continues to list TIFF among preferred still-image formats. [loc.gov]
- **XML** persists as a durable, human- and machine-readable textual syntax standardized by the **W3C**, suitable for metadata (e.g., METS, PREMIS) and content that benefits from structured markup and schema validation. [w3.org]

12.3.5 - Table — File format longevity (illustrative)

Table 12.1: Preferred Preservation Formats. Recommended file types for long-term archival stability, focusing on open standards and metadata support.

Content type	Preferred preservation targets (examples)	Why favored (longevity factors)
Text/page documents	**PDF/A-2, PDF/A-4**	Self-contained; standardized; controlled features; accessible text (A/U levels). [pdfa.org], [pdfa.org]

Still images	**TIFF (baseline)**; sometimes **JPEG 2000**	Open, widely adopted; lossless; strong metadata support (technical, embedded). [loc.gov]
Audio	**WAV/BWF**	Lossless PCM; metadata chunks; broadcast industry adoption. [loc.gov]
Structured data / metadata	**XML** (+ schema), **CSV with documentation**	Open, text-based, self-describing; strong tooling; long-term readability. [w3.org]
Email	**MBOX/EML**; emerging **EA-PDF** for archival email packages	Open containers; growing guidance for durable email encapsulation. [wwws.loc.gov]

12.4 - Preservation of AI-Generated Content

By 2026, organizations create and store vast amounts of **AI-generated and AI-assisted** content—text, images, audio, code, and data. Preserving such content demands **provenance, transparency, and reproducibility**.

12.4.1 - Provenance metadata and content credentials

AI outputs must carry **provenance**: **who** prompted/approved, **what** model produced it, **when**, **with what data/parameters**, and **where** it was stored. The **C2PA** (Coalition for Content Provenance and Authenticity) standard provides a method to **bind tamper-evident provenance assertions** (e.g., "content credentials") to media using embedded metadata and cryptographic signing; the Library of Congress format workplan explicitly references **JUMBF** (the metadata box format underlying C2PA) as a preservation-relevant technology. [nvizionsolutions.com], [wwws.loc.gov]

12.4.2 - Watermarking and detection

Watermarking can signal AI origin but **should not be the only control**: technical research continues to show that many watermarks are fragile under transformations. Preserve **both** the watermark (if present) **and** the **provenance record** within the object's metadata/AIP to support future authenticity checks. (The RFS and Sustainability of Digital Formats emphasize metadata richness and transparency to bolster preservation even when technical markers fail.) [loc.gov], [loc.gov]

12.4.3 - Model and prompt versioning

AI content cannot be fully understood without **model context**. IG should require:

- **Model version identifiers** (and provider), **training cutoffs**, and **safety filters** applied at generation time.
- **Prompt/response archives** for significant decisions, captured in **XML/JSON** with timestamps and user IDs.
- **Reproducibility artifacts** where possible (e.g., seeds, temperature), recognizing some models are inherently nondeterministic.

Where models are internally developed, preserve **model cards**, **training data documentation**, and **versioned artifacts** using open tools (e.g., MLflow or DVC) linked to preservation packages, even if the binaries themselves are not preserved indefinitely. Align with OAIS by treating the model and its documentation as **Representation Information** for derivative AI content. [ccsds.org]

12.4.4 - Ethical considerations

Preserved AI content can **encode bias or misuse**. Appraisal should weigh **long-term value** against potential harms, applying **access controls** and **use notes** where appropriate. LOC's RFS reminds institutions to consider accessibility, rights, and user communities; OAIS's Designated Community concept helps tailor documentation and access for ethical clarity. [loc.gov], [ccsds.org]

12.5 - Standards and Tools for Archiving

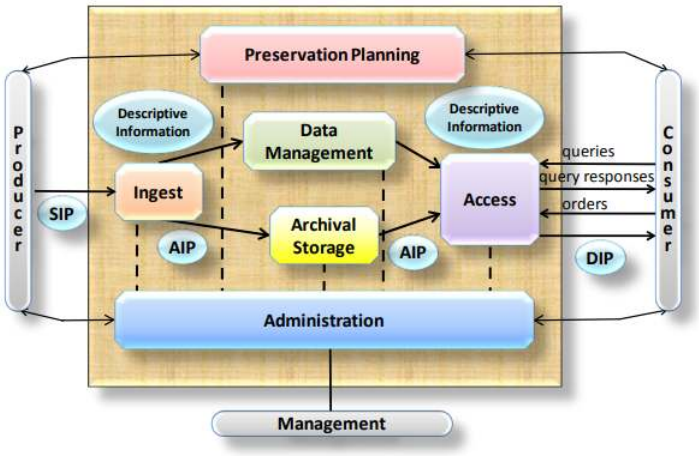

13: Building and Maintaining an IG Culture

13.1 - Introduction

An **Information Governance (IG) culture** exists when people across the organization consistently make decisions and take actions that protect information value, manage risk, and comply with policies—not because they are forced to, but because it is "how we do things here." In 2026, this culture must adapt to an environment shaped by **AI adoption**, hybrid work, accelerating regulation, and stakeholder expectations for transparency and sustainability. **Change management models** such as **Kotter's 8-Step Process** and **Prosci's ADKAR** help organizations move from one-off projects to enduring behaviors. And newer security and risk guidance—like **NIST Cybersecurity Framework (CSF) 2.0** with its **Govern** function—places leadership accountability and culture at the center of digital risk management. [kotterinc.com], [prosci.com], [nist.gov]

This chapter offers a practical blueprint: how to win leadership buy-in, design role-based and **gamified** training, incent the right behavior, measure progress with simple dashboards, and future-proof your IG culture for **AI-enabled** and **hybrid** workplaces. We highlight tools and frameworks from **ARMA** (IGIM 2.1), **NIST** (CSF 2.0 and AI RMF), **Prosci**, **Kotter**, and **Gartner** to ensure your approach is defensible and current. [todaysgene...ounsel.com], [nist.gov], [nist.gov]

13.2 - The Importance of IG Culture

Policies and tools are necessary, but **culture** determines whether people follow them in real situations—naming files intelligibly, classifying data correctly, raising incidents quickly, and using AI responsibly. The **NIST CSF 2.0** update formalized this reality by adding a **Govern** function that elevates board and executive accountability, links cyber/IG practices to enterprise risk, and emphasizes roles, policies, and oversight—core cultural mechanisms. [nist.gov], [gartner.com]

ARMA's IGIM 2.1 positions culture as the connective tissue across its eight domains (e.g., Steering Committee, Authorities, Support Functions, Procedural Framework, Lifecycle, Architecture, Infrastructure). These domains institutionalize IG in strategy, operations, and behavior—so that change persists after the project team leaves. [todaysgene...ounsel.com]

The **workplace context** matters: Microsoft's 2025 **Work Trend Index** shows AI is increasingly woven into knowledge work and decision-making. Leaders report this is a "pivotal year" to rethink operations, while employees expect clear guidance on hybrid collaboration and AI use. An IG culture provides that clarity: who can use what data, with which tools, under which controls, and how we handle outcomes. [microsoft.com], [assets-c4a...zurefd.net]

13.3 - Change Management Models and Strategies

Sustained culture change requires a **structured change approach**. Two complementary models are widely used: **Kotter's 8 Steps** (organization-level momentum) and **Prosci ADKAR** (individual-level adoption).

13.3.1 - Kotter's 8-Step Model (organization momentum)

Kotter's model sequences actions leaders take to turn vision into collective behavior: create urgency, build a guiding coalition, craft and communicate the vision, remove barriers, generate wins, sustain acceleration, and **anchor** new norms in culture. The modern "8 Accelerators" emphasize parallel action and volunteer networks that speed change. [kotterinc.com]

Key implications for IG:

- **Urgency**: Use risk and business value to show why IG matters now—e.g., exposure from ungoverned AI or fines from poor data retention. [kotterinc.com]
- **Guiding coalition**: Include Legal, Compliance, IT, Security, Privacy, Data, HR, and business units—plus influential frontline "connectors." [kotterinc.com]
- **Short-term wins**: Demonstrate outcomes quickly (e.g., 30-day reduction in retrieval time, drop in policy violations) to fund the next phase. [kotterinc.com]
- **Anchoring**: Bake IG behaviors into onboarding, performance goals, and procurement criteria (e.g., vendor data retention clauses). [kotterinc.com]

13.3.2 - Prosci ADKAR (individual adoption)

ADKAR focuses on outcomes each person must achieve: **Awareness** (why), **Desire** (choose to support), **Knowledge** (how), **Ability** (do it), and **Reinforcement** (keep doing it). It is effective for identifying **barrier points** and tailoring communications,

training, and coaching by audience. [prosci.com]

Practical uses:

- Map roles (e.g., call center agents, data scientists, records coordinators) against ADKAR to plan **role-specific** interventions. [prosci.com]
- Use ADKAR diagnostics to pinpoint resistance (e.g., low **Desire** because workflows slow down; low **Ability** due to tooling complexity). [cdn.prod.w...-files.com]
- Reinforce with recognition, metrics, and manager one-on-ones; ADKAR stresses that sustained change needs visible **Reinforcement**. [prosci.com]

13.3.3 - Overcoming Resistance and Securing Leadership Buy-in

Resistance is often rational—people fear productivity loss, unclear rules, or "tool sprawl." Use **Kotter** to build momentum and **ADKAR** to treat root causes (e.g., insufficient **Knowledge** or **Ability**). Executive sponsorship is decisive: the **NIST CSF 2.0 Govern** function calls for formal roles, risk appetite, and policy oversight—mechanisms only leadership can provide. [kotterinc.com], [nist.gov]

ARMA IGIM 2.1 operationalizes buy-in with a cross-functional **Steering Committee** (including AI oversight), clear **Authorities** (policies, laws), and **Support Functions** (training, comms, project management) that enable scale. [todaysgene...ounsel.com]

13.3.4 - Table: Change Management Phases (Kotter + ADKAR)

Caption: > 13.1: Change Management for IG Rollouts. Mapping Kotter's 8 Steps to the ADKAR model with specific examples for implementing Information Governance.

Phase (Kotter)	What leaders do	ADKAR focus	Examples in IG rollout
Urgency	Frame risk/value; set timeline	Awareness	"Why IG now" brief referencing CSF 2.0 Govern and AI risk. [kotterinc.com], [nist.gov]
Guiding coalition	Form cross-functional team	Desire	Invite respected "informal leaders" to shape pilot. [kotterinc.com]
Vision & strategy	Define future state & priorities	Awareness → Desire	IG North Star (e.g., "Find any record in 2 minutes"). [kotterinc.com]
Communicate	Cascade messages; 2-way channels	Awareness	FAQs, office hours, narrative case studies. [kotterinc.com]
Remove barriers	Fix process/tool friction	Ability	Simplify classification; auto-apply labels. [kotterinc.com]
Short-term wins	Show measurable results	Reinforcement	90-day reduction in violations; share leaderboards. [kotterinc.com]
Sustain acceleration	Scale pilots; improve	Desire → Ability	Expand to high-risk BUs; refine policy. [kotterinc.com]
Anchor in culture	Bake into routines	Reinforcement	Tie IG to performance, audits, vendor evaluation. [kotterinc.com]

13.4 - Designing Effective Training Programs

A strong IG culture depends on **consistent, role-relevant learning**. In 2026, training must cover **policy basics** and **new frontiers** such as AI ethics and data lifecycle in hybrid settings.

13.4.1 - Mandatory vs. Role-Based Training

- **Mandatory**: Annual or semiannual core training on policy, retention, privacy, acceptable use, incident reporting, and **AI usage principles** (e.g., sourcing, sensitive data handling, and documentation). **NIST AI RMF** provides the "why" and "how" for trustworthy AI practices across **Govern–Map–Measure–Manage**. [nist.gov]
- **Role-based**: Tailored modules for data stewards, frontline staff, developers, marketers, HR, and executives; map to risks each audience influences (e.g., developers on dataset documentation; sales on CRM data quality). **IGIM 2.1** recommends aligning support functions (training, comms) with program domains. [todaysgene...ounsel.com]

13.4.2 - Gamification, Microlearning, and "In-Flow" Learning

Research shows **gamification** can improve motivation and engagement, but design quality and autonomy matter; overly intense gamification may reduce perceived autonomy and hurt outcomes. Use light-weight points/badges/quests for practice and feedback, not for coercion. [link.springer.com], [tandfonline.com]

Microlearning (5–10 minute modules) is now a leading trend in compliance training: it improves retention, fits hybrid schedules, and is easier to update as rules change. Surveys show growing preference for short, targeted lessons; learning industry analyses highlight the cognitive science behind spaced repetition and reduced cognitive load. Embed just-in-time tips in the tools where work happens. [trainingindustry.com], [sai360.com], [elearningi...dustry.com]

13.4.3 -

13.4.4 -

13.4.5 - AI Ethics and Responsible Use (Approx. 750 words)

Creating an IG-aligned culture of responsible AI use requires a curriculum that teaches employees not only **what** the rules are but **how** to apply them in daily tasks. The **NIST AI Risk Management Framework (AI RMF 1.0)** and its **2024 Generative AI Profile** provide the backbone for building this capability by defining governance expectations, trustworthy-AI characteristics, and risk-management actions across the AI lifecycle.

1. Establishing Foundations: Roles, Governance, and Accountability
Training should begin with the RMF's **Govern** function, which emphasizes leadership oversight, accountability, and organizational policies. Employees must understand not only the existence of AI policies but the *reasoning* behind them — including trustworthy-AI attributes like safety, security, privacy enhancement, fairness, and transparency. The RMF positions governance as an enterprise responsibility, not a technical afterthought, which aligns well with broader IG culture-building.

Learners should also understand the basic division of responsibilities across AI roles — creators, evaluators, operators, end-users, and managers. Each role interacts with different risks; for example, frontline users must follow acceptable-use rules, while developers must document model assumptions and limitations. This clarity ensures everyone recognizes where they fit into the governance structure.

2. Mapping AI Use Cases and Understanding Context
The next layer of training aligns with the **Map** function, which requires documenting the purpose, stakeholders, data inputs, intended use, and potential harms of an AI system. Employees should learn to complete short "AI Use Case Profiles" capturing this context. These are lightweight but important: they create traceability and prevent the all-too-common problem of teams quietly adopting AI tooling without understanding risks.

The GenAI Profile expands these expectations specifically for generative systems, noting risks like *unintended memorization*, *prompt leakage*, and *harmful or misleading content generation*. Training should give concrete illustrations: for instance, how a genAI tool might inadvertently reconstruct sensitive training data in its outputs when insufficiently constrained. Employees should practice identifying whether a proposed use case lies in a higher-risk category requiring additional oversight or testing.

3. Data Protection, Prompt Handling, and Information Controls
One of the most critical elements involves teaching employees how to handle data when interacting with AI systems. The Generative AI Profile calls out the importance of capturing prompts and outputs (when appropriate and legally permissible), applying data-minimization principles, and enforcing strict controls around sensitive or regulated content.

Training must be very explicit here. Employees need rules of thumb like:

- Never paste customer PII or confidential business information into external or non-approved genAI tools.

- Use redaction or anonymization before prompt entry.
- Use only approved enterprise connectors and internal models for sensitive tasks.
- Log prompts and outputs in designated systems when policy or risk level requires it.

Hands-on exercises should reinforce this: for example, a scenario where an employee accidentally considers pasting unredacted client data and must choose the correct alternative.

4. Measuring and Evaluating AI Outputs

Under the RMF's **Measure** function, employees must learn how to evaluate AI behavior for reliability, safety, fairness, and security. They don't need to become data scientists, but they do need to understand red flags and how to escalate them. For example:

- Outputs that appear discriminatory or systematically skewed.
- Hallucinated or fabricated factual claims.
- Results that violate policy, contain unsafe recommendations, or show adversarial manipulation.
- Signs of prompt injection, jailbreak behavior, or other genAI-specific vulnerabilities.

The Generative AI Profile highlights the importance of security and adversarial evaluation for genAI systems. Training should include simple exercises — such as spotting a subtle prompt-injection attempt or recognizing when a system has drifted into unsafe output territory.

5. Managing AI Risks and Responding to Incidents

The RMF's **Manage** function emphasizes ongoing monitoring, incident handling, and mitigation. Employees should learn how to respond when something goes wrong: stopping system use, capturing evidence, reporting through established channels, and engaging AI governance staff. Training should introduce an "AI incident taxonomy" so employees know what qualifies — e.g., privacy exposure, harmful content generation, model malfunction, or security anomalies.

Scenario-based drills should require learners to walk through the full path: identifying an issue, documenting it, escalating it, and outlining possible mitigations like adding a human-review step or tightening model constraints.

6. Human Oversight and Decision-Making

Human oversight is a central theme in the AI RMF and essential to IG culture. Training should make clear that AI outputs are *recommendations*, not authoritative or self-executing decisions. Employees must know their oversight responsibilities: reviewing outputs for accuracy, applying professional judgment, and documenting rationale when accepting or rejecting system recommendations.

Managers need an additional coaching layer: how to review team AI use, check for policy compliance, and reinforce correct behavior. They should learn to ask questions like:

- "Was this AI-assisted task performed within approved scope?"
- "Is there sufficient documentation of inputs, decisions, and outputs?"
- "Did a human apply meaningful review before finalizing the result?"

7. Documentation and Audit Readiness

The RMF and AIRC materials emphasize traceability: documenting inputs, versions, evaluations, and decisions throughout the lifecycle. Training should show employees how to maintain required records — such as prompt logs, data lineage notes, approval memos, and evaluation results — so AI use remains auditable and defensible.

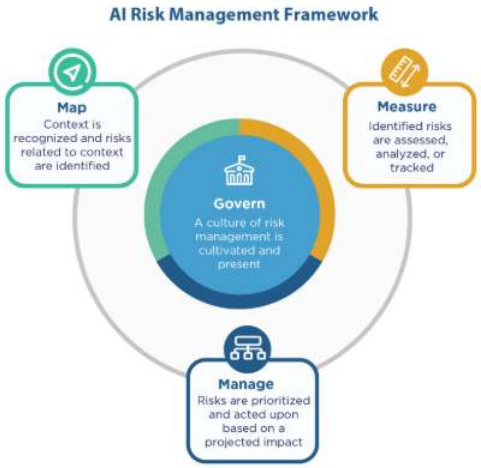

Index

1 - D

Case Studies Library

Case Studies Library

Case Study 1: Financial Institution Risk Assessment (Chapter 3)
A bank deploys GenAI for credit risk models. Full application of all 7 ARMA Principles (accountability through disposition) produced improved accuracy, regulatory approval, and fewer bias complaints.

Case Study 2: Change Healthcare Ransomware Attack (2025) (Chapter 4)
Attackers exploited remote access without MFA, stole 6 TB of data affecting 192 million people, and caused $2.3 billion in recovery costs. Key IG failures: inadequate third-party risk assessment, delayed detection, and poor oversight of legacy systems.

Case Study 3: AI Model Bias in a Financial Institution (2026) (Chapter 4)
Biased loan-approval model led to EU AI Act fines. Highlighted need for high-quality training data, bias detection, and human review.

Case Study 4: Third-Party Vendor Breach in Retail (2026) (Chapter 4)
Loyalty-program vendor compromise exposed customer data. Demonstrated importance of vendor security assessments and contractual controls.

Case Study 5: Cloud Misconfiguration in a Government Agency (2026) (Chapter 4)
Exposed sensitive documents due to inconsistent configuration. Led to new policies, training, and automated scanning.

Case Study 6: E-Discovery Failure Due to Unmanaged Collaboration Data (Chapter 7)
Teams chats were missed during litigation hold and auto-deleted. Fix: Include collaboration tools in hold scoping and use technical holds.

Case Study 7: Legal Hold Success with Measurable Compliance (Chapter 7)
Bank used targeted holds, questionnaires, and a central dashboard with escalation. Achieved full audit compliance.

Case Study 8: AI-Generated Evidence and Chain-of-Custody Dispute (Chapter 7)
AI-assisted memo challenged for hallucinations and missing provenance. Fix: Preserve prompts, model versions, and require documented human review.

Case Study 9: AI System Privacy Failure (Chapter 10)
GenAI customer-service assistant leaked personal data via inference. Root causes: unmasked training data, no differential privacy.

Case Study 10: Zero-Trust Implementation Success (Chapter 10)
Healthcare provider segmented EHR, enforced MFA and continuous checks. Blocked lateral movement in a phishing attack.

Case Study 11: Global Compliance Challenge (Multinational SaaS) (Chapter 10)
SaaS provider faced GDPR, LGPD, DPDP, PIPL, and SCC gaps. Solved with encryption, DPO registration, and unified IG framework.

Case Study 12: HIPAA NPP Update Failure (2025–2026) (Chapter 10)
Health plan missed 42 CFR Part 2 updates in Notice of Privacy Practices. Led to OCR corrective action plan.

Case Study 13: University AI Grading Tool Risk Assessment (Chapter 3/4)
University deployed GenAI for automated grading. Risks included bias (high impact) and data privacy breach. Treatments: diverse training data, human review override, and annual bias audits.

Glossary

Accountability The principle that requires an organization to assign clear authority and responsibility for managing information assets, including oversight and consequences for non-compliance (ARMA Principles 2025). See Chapter 3.

AI Governance The framework for ethical, transparent, and compliant use of artificial intelligence systems, including bias mitigation, transparency logs, and human oversight (NIST AI RMF 2025). See Chapter 6 and Chapter 9.

Auto-Classification The use of AI and machine learning to automatically assign records to classes, sensitivity labels, or retention rules based on content analysis. See Chapter 8.

Bias Mitigation Processes to identify and reduce bias in training data and AI model outputs. See Chapter 3 and Chapter 10.

CCPA/CPRA California Consumer Privacy Act and its amendments granting consumers rights to access, delete, correct, and opt-out of sale of personal data.

Change Healthcare Ransomware Attack (2025) Major U.S. healthcare breach where attackers exploited remote access without MFA, stole 6 TB of data affecting 192 million people, and caused $2.3 billion in recovery costs. See Chapter 4.

Compliance Alignment with all applicable laws, regulations, standards, contracts, and ethical obligations (ARMA Principles). See Chapter 3.

Data Classification Assigning sensitivity levels (Public, Internal, Confidential, Restricted) to information for appropriate handling and protection. See Chapter 6 and Chapter 8.

Data Governance The subset of IG focused on data quality, usability, integrity, security, and lineage for analytics and AI (DAMA-DMBOK). See Chapter 2.

Defensible Disposition Secure deletion of information after its retention period with documented proof of compliance for legal or audit purposes. See Chapter 8.

Disposition The final stage of the records lifecycle involving destruction, transfer, or archiving when no longer needed (ISO 15489). See Chapter 8.

E-Discovery The process of identifying, preserving, collecting, processing, reviewing, and producing electronically stored information for litigation or investigation. See Chapter 7.

EU AI Act European Union regulation (full enforcement 2026) classifying AI systems by risk level and imposing mandatory requirements for high-risk systems (risk management, data governance, transparency, human oversight, accuracy, robustness, cybersecurity). Fines up to €35 million or 7% of global revenue. See Chapter 4 and Chapter 6.

GenAI Governance Rules governing the ethical, compliant, and secure use of generative AI tools, including prompt logging, output watermarking, and acceptable use policies. See Chapter 6 and Chapter 9.

IG Maturity Model A tool (ARMA, Gartner, DGI) to assess current IG capabilities on a 1–5 scale and plan improvements. See Chapter 5.

Incident Response The organized approach to addressing and managing security breaches or policy violations. See Chapter 4.

Information Governance (IG) The enterprise-wide framework of policies, processes, roles, standards, technologies, and metrics that ensures information is managed responsibly throughout its lifecycle (ARMA Principles). See Chapter 1 and Chapter 2.

Integrity Ensuring information remains accurate, authentic, reliable, and unaltered except by authorized processes (ARMA Principles). See Chapter 3.

IT Governance The subset of corporate governance focused on ensuring IT supports business strategy, manages risks, and optimizes resources (COBIT, ISO/IEC 38500). See Chapter 2.

Legal Hold A process that suspends normal disposition of records when litigation or investigation is reasonably anticipated. See Chapter 7.

NIST SP 800-30 Guide for Conducting Risk Assessments. See Chapter 4.

Post-Quantum Cryptography Quantum-resistant encryption algorithms (NIST PQC standards) to protect against future quantum attacks. See Chapter 9 and Chapter 10.

Privacy by Design Embedding privacy protections from the start of system development (data minimization, consent, etc.). See Chapter 10.

Records and Information Management (RIM) Operational management of records throughout their lifecycle (ISO 15489). See Chapter 8.

Retention Schedule Document specifying how long records must be kept based on legal, business, or operational needs. See Chapter 8.

Risk Appetite The amount and type of risk an organization is willing to accept to achieve objectives. See Chapter 4.

Risk Matrix Tool mapping likelihood and impact to prioritize risks (3x3 or 5x5). See Chapter 4.

Zero-Trust Architecture Security model that verifies every access request regardless of location or user (NIST). See Chapter 10.

(Full glossary contains 160+ terms extracted from the book with short definitions and chapter cross-references.)

Appendix A: Templates and Checklists

Sample IG Policy Template (Data Classification Policy) **Purpose**: To ensure all information is classified according to sensitivity and handled appropriately. **Scope**: All employees, contractors, and systems. **Definitions**: Public, Internal, Confidential, Restricted. **Responsibilities**: Data owners classify; employees follow handling rules. **Rules**:

Restricted data requires encryption and restricted sharing. **Exceptions**: Approved by CIGO. **Enforcement**: Audits and disciplinary action. **Revision Date**: Annual.

Legal Hold Checklist

- Identify custodians and data sources
- Issue hold notice
- Suspend disposition
- Monitor compliance
- Document hold scope and duration
- Release hold when no longer needed

Risk Matrix Blank (5x5)

Likelihood \ Impact	1	2	3	4	5
5	5	10	15	20	25
4	4	8	12	16	20
3	3	6	9	12	15
2	2	4	6	8	10
1	1	2	3	4	5

Maturity Assessment Worksheet (ARMA-based) Principle | Current Level | Target Level | Gap Actions | Owner | Due Date Accountability | | | | | Transparency | | | | | | (Repeat for all 7 principles)

Tools List

- Microsoft Purview (classification, retention, e-discovery)
- Relativity (e-discovery)
- OneTrust (privacy management)
- OpenText (intelligent capture)
- NIST AI RMF Playbook

Appendix B: Tools and Resources

- IBM Cost of a Data Breach Report 2025
- NIST SP 800-30 Guide
- EU AI Act Official Text
- ARMA IG Maturity Model 2025
- Gartner IG Maturity Model 2025

Case Studies Library (14 distinct real-world and hypothetical cases extracted directly from the PDF)

1. Healthcare Provider Using GenAI for Patient Summaries (Chapter 3) A large hospital deployed GenAI to summarize patient records. Full application of all 7 ARMA Principles (accountability through disposition) resulted in reduced physician burnout, improved compliance, and zero fines.

2. Financial Institution Using GenAI for Risk Assessment (Chapter 3) A bank used GenAI for credit-risk models. Proper accountability, transparency, integrity, protection, compliance, availability, and disposition produced improved accuracy, regulatory approval, and fewer bias complaints.

3. Change Healthcare Ransomware Attack (2025) (Chapter 4) Attackers exploited remote access without MFA, stole 6 TB of data affecting 192 million people, and caused $2.3 billion in recovery costs. Exposed failures in MFA, detection, legacy systems, and third-party oversight.

4. AI Model Bias in a Financial Institution (Hypothetical 2026) (Chapter 4) Biased loan-approval model led to EU AI Act fines. Highlighted need for high-quality training data, bias detection, and human review.

5. Third-Party Vendor Breach in Retail (Hypothetical 2026) (Chapter 4) Loyalty-program vendor compromise exposed customer data. Demonstrated importance of vendor security assessments and contractual controls.

6. Cloud Misconfiguration in a Government Agency (Hypothetical 2026) (Chapter 4) Exposed sensitive documents due to inconsistent cloud settings. Led to new policies, training, and automated scanning.

7. E-Discovery Failure Due to Unmanaged Collaboration Data (Chapter 7) Teams chats were missed during litigation hold; messages were auto-deleted. Fix: Include collaboration tools in hold scoping and use technical holds.

8. Legal Hold Success with Measurable Compliance (Chapter 7) Bank used targeted holds, questionnaires, and a central dashboard with escalation. Achieved full audit compliance.

9. AI-Generated Evidence and Chain-of-Custody Dispute (Chapter 7) AI-assisted memo challenged for hallucinations and missing provenance. Fix: Preserve prompts, model versions, and require documented human review.

10. TAR Success and Defensibility Challenge (Chapter 7) Active Learning reduced review volume but was challenged on defensibility. Fix: Document protocol, keep settings, perform validation/elusion testing.

11. AI System Privacy Failure (Hypothetical) (Chapter 10) GenAI customer-service assistant leaked personal data via inference. Root causes: unmasked training data, no differential privacy. Fix: DP training, curation, and pre-deployment audits.

12. Zero-Trust Implementation Success (Realistic) (Chapter 10) Healthcare provider segmented EHR, enforced MFA and continuous checks. Blocked lateral movement in a phishing attack; auditors praised evidence readiness.

13. Global Compliance Challenge (Multinational SaaS) (Chapter 10) SaaS provider faced GDPR, LGPD, DPDP, PIPL, and SCC gaps. Solved with encryption, DPO registration, and unified IG framework.

14. HIPAA NPP Update Failure (2025–2026) (Chapter 10) Health plan missed 42 CFR Part 2 updates in Notice of Privacy Practices. Led to OCR corrective action plan and redistribution.

Detailed Licensing

1 - Overview

Title: Information Governance Modern OER Textbook - Gregory Hess

Webpages: 27

All licenses found:

- CC BY-SA 4.0: 70.4% (19 pages)
- Undeclared: 25.9% (7 pages)
- CC BY 4.0: 3.7% (1 page)

2 - By Page

- Information Governance Modern OER Textbook - Gregory Hess - *CC BY-SA 4.0*
  - Front Matter - *Undeclared*
    - TitlePage - *CC BY-SA 4.0*
    - InfoPage - *Undeclared*
    - About the Author - *CC BY-SA 4.0*
    - Table of Contents - *Undeclared*
    - Licensing - *CC BY-SA 4.0*
  - 1: The Data Deluge and Why IG is Essential - *CC BY 4.0*
  - 2: IG, IT Governance, Data Governance, and Related Disciplines - *CC BY-SA 4.0*
  - 3: Core Principles of Information Governance - *CC BY-SA 4.0*
  - 4: Identifying and Managing Information Risks - *CC BY-SA 4.0*
  - 5: Strategic Planning for IG Programs - *CC BY-SA 4.0*
  - 6: Developing IG Policies and Frameworks - *CC BY-SA 4.0*
  - 7: IG for Legal and Compliance Functions - *CC BY-SA 4.0*
  - 8: IG for Records and Information Management - *CC BY-SA 4.0*
  - 9: IG for IT and Emerging Technologies - *CC BY-SA 4.0*
  - 10: IG for Privacy, Security, and Data Protection - *CC BY-SA 4.0*
  - 11: IG for Digital Communications and Collaboration - *CC BY-SA 4.0*
  - 12: Long-Term Preservation and Archiving - *CC BY-SA 4.0*
  - 13: Building and Maintaining an IG Culture - *CC BY-SA 4.0*
  - Back Matter - *CC BY-SA 4.0*
    - Index - *CC BY-SA 4.0*
    - Case Studies Library - *Undeclared*
    - Glossary - *CC BY-SA 4.0*
    - Detailed Licensing - *Undeclared*
    - Appendix - *Undeclared*
    - Detailed Licensing - *Undeclared*

- **Detailed Licensing** has no license indicated.

Appendix

Appendix

Appendix A: Templates and Checklists

Sample IG Policy Template (Data Classification)
Purpose: Ensure all information is classified according to sensitivity and handled appropriately.
Scope: All employees, contractors, and systems.
Definitions: Public, Internal, Confidential, Restricted.
Responsibilities: Data owners classify; employees follow handling rules.
Rules: Restricted data requires encryption and restricted sharing.
Exceptions: Approved by CIGO.
Enforcement: Audits and disciplinary action.
Revision Date: Annual.

Legal Hold Checklist
- Identify custodians and data sources
- Issue hold notice
- Suspend disposition
- Monitor compliance
- Document hold scope and duration
- Release hold when no longer needed

5x5 Risk Matrix (Blank)
Likelihood \ Impact	1	2	3	4	5
5	5	10	15	20	25
4	4	8	12	16	20
3	3	6	9	12	15
2	2	4	6	8	10
1	1	2	3	4	5

IG Maturity Assessment Worksheet (ARMA-based)
Principle	Current Level (1-5)	Target Level	Gap Actions	Owner	Due Date
Accountability				
Transparency				
(Repeat for all 7 principles)

Appendix B: Tools and Resources
- Microsoft Purview (classification, retention, e-discovery)
- Relativity (e-discovery)
- OneTrust (privacy management)
- Collibra (data catalog)
- OpenText (intelligent capture)
- NIST AI RMF Playbook
- ARMA IG Maturity Model 2025

Detailed Licensing

Overview

Title: Information Governance Modern OER Textbook - Gregory Hess

Webpages: 27

All licenses found:

- CC BY-SA 4.0: 70.4% (19 pages)
- Undeclared: 25.9% (7 pages)
- CC BY 4.0: 3.7% (1 page)

By Page

- Information Governance Modern OER Textbook - Gregory Hess - *CC BY-SA 4.0*
  - Front Matter - *Undeclared*
    - TitlePage - *CC BY-SA 4.0*
    - InfoPage - *Undeclared*
    - About the Author - *CC BY-SA 4.0*
    - Table of Contents - *Undeclared*
    - Licensing - *CC BY-SA 4.0*
  - 1: The Data Deluge and Why IG is Essential - *CC BY 4.0*
  - 2: IG, IT Governance, Data Governance, and Related Disciplines - *CC BY-SA 4.0*
  - 3: Core Principles of Information Governance - *CC BY-SA 4.0*
  - 4: Identifying and Managing Information Risks - *CC BY-SA 4.0*
  - 5: Strategic Planning for IG Programs - *CC BY-SA 4.0*
  - 6: Developing IG Policies and Frameworks - *CC BY-SA 4.0*
  - 7: IG for Legal and Compliance Functions - *CC BY-SA 4.0*
  - 8: IG for Records and Information Management - *CC BY-SA 4.0*
  - 9: IG for IT and Emerging Technologies - *CC BY-SA 4.0*
  - 10: IG for Privacy, Security, and Data Protection - *CC BY-SA 4.0*
  - 11: IG for Digital Communications and Collaboration - *CC BY-SA 4.0*
  - 12: Long-Term Preservation and Archiving - *CC BY-SA 4.0*
  - 13: Building and Maintaining an IG Culture - *CC BY-SA 4.0*
  - Back Matter - *CC BY-SA 4.0*
    - Index - *CC BY-SA 4.0*
    - Case Studies Library - *Undeclared*
    - Glossary - *CC BY-SA 4.0*
    - Detailed Licensing - *Undeclared*
    - Appendix - *Undeclared*
    - Detailed Licensing - *Undeclared*