

Hype Cycle for Privacy, 2026

24 June 2026 - ID G00846812 - 98 min read

By: Stefan Dumitrescu, Shadrock Roberts

Initiatives: Strategy, Risks and Opportunities; Build and Evolve a Resilient and Agile Cybersecurity Program

Privacy remains a top concern for boards, executives, and consumers as AI draws attention to the need for additional safeguards against the misuse of personal data. Cybersecurity leaders should use this Hype Cycle to prioritize investments that drive scalable protection.

Strategic Planning Assumptions

By 2030, all organizations must include age assurance mechanisms in their privacy UX to comply with restrictions in processing minors' personal data.

By 2030, unused data will represent one of the top largest liabilities on the balance sheet of consumer-facing companies.

By 2030, 50% of the world's population will be covered under modern AI regulations in one form or another, up from 7% today.

By 2029, most privacy incidents will result from AI-generated inferences about individuals, not direct exposure of personally identifiable information.

Analysis

What You Need to Know

AI adoption is leading CISOs to take a greater role in enterprise privacy, with 41% already accountable for global privacy outcomes. ¹

This shift is accelerated by evolving privacy regulations, as well as emerging AI-specific laws such as the EU AI Act and a growing patchwork of regional AI governance requirements (see [2026 AI Regulation: Comparative Guide to U.S. State and EU Rules](#)). In addition, there is growing scrutiny around automated decision making, model training, inference, and autonomous agent behaviors that can materially impact individuals and their personal data.

To deliver an effective privacy program in an era of growing AI influence, cybersecurity leaders must leverage emerging privacy-enhancing technologies, inculcate privacy-by-design principles, and drive coordinated participation from functions across the enterprise.

The Hype Cycle

Cybersecurity leaders should leverage a variety of technologies to support privacy goals across AI, analytics, and business intelligence. Privacy-enhancing technologies, such as synthetic data and differential privacy, enable privacy-preserving solutions by allowing AI and analytics to leverage data without directly exposing sensitive personal information. Their effectiveness varies depending on the use case and requires careful implementation to balance privacy and utility (see [Use-Case Guidance for Choosing Privacy-Enhancing Technologies](#)).


Technology is not enough to establish a mature privacy program. Privacy is a multidisciplinary practice, combining technical capabilities with strong governance, privacy-by-design principles, and cross-functional collaboration. While protecting data through cybersecurity controls is fundamental to privacy, this alone does not guarantee privacy outcomes. Effective privacy programs also require organizations to apply operational disciplines such as the application of data minimization, purpose limitation, retention governance, and responsible data use practices.

This Hype Cycle examines the technologies, governance practices and operational capabilities organizations are adopting to:

- Enable trusted access to data by maintaining appropriate controls and regulatory safeguards (e.g., data discovery and subject rights requests)
- Improve governance transparency and audit readiness, and demonstrate regulatory defensibility across business operations (e.g., privacy management tools)
- Mitigate the risk of unauthorized data exposure and the risk of misuse or unsafe AI processing activities (e.g., confidential computing, synthetic data, homomorphic encryption)
- Limit the operational and reputational impact of unintended data processing outcomes and AI-driven decisions (e.g., format-preserving encryption)
- Improve consistency, accountability, and repeatability in privacy and AI risk decision making across the organization (e.g., AI trust, risk and security management [AI TRiSM] and privacy by design)
- Strengthen customer trust, digital transparency, and user experience while supporting responsible innovation (e.g., influence AI, consent and preference management)

Figure 1: Hype Cycle for Privacy, 2026

Hype Cycle for Privacy, 2026



The Priority Matrix

The Priority Matrix maps the benefit rating of each innovation or technology against the expected time it takes that innovation or technology to achieve mainstream adoption. The benefit rating indicates the innovation's or technology's potential, but the rating may not apply equally to all industries and organizations.

Transformative technologies such as sovereign cloud, decentralized identity, homomorphic encryption, and influence AI have the potential to fundamentally change how organizations govern, share, and protect data. While adoption may require significant investment and organizational change, early adopters can gain competitive advantages through improved data sovereignty, stronger privacy protections, greater trust by individuals, and enablement of new forms of data collaboration and AI innovation.

High-impact technologies are expected to mature significantly over the next three years. Key areas of advancement include privacy operations, such as breach response and subject rights automation; governance capabilities, such as AI TRiSM and digital ethics; and privacy-enhancing technologies, including confidential computing and federated machine learning.

Organizations are increasingly opting for bundled privacy user experience (UX) solutions that integrate subject rights request (SRR) handling, consent and preference management (CPM), and tracker cookie management (TCM) into a single, cohesive user experience. This approach improves compliance posture, reduces overheads, and offers greater transparency than stitching together stand-alone tools.

Table 1: Priority Matrix for Privacy, 2026

(Enlarged table in Appendix)

Benefit ↓	Years to Mainstream Adoption			
	Less Than 2 Years ↓	2 to 5 Years ↓	5 to 10 Years ↓	More Than 10 Years ↓
Transformational	Sovereign Cloud	Decentralized Identity Homomorphic Encryption	Influence AI	
High	Data Breach Response Digital Ethics Privacy Management Tools Subject Rights Requests	Age Assurance and Parental Consent Tracking AI TRISM Confidential Computing Privacy Impact Assessments	Federated Machine Learning	
Moderate	Format-Preserving Encryption Personification	5G Network Security AI Security Posture Management Consent and Preference Management Data Clean Rooms Data Discovery Data Security Governance Privacy by Design Synthetic Data		
Low		Differential Privacy		

Source: Gartner (June 2026)

Off the Hype Cycle

- Data classification has been moved off this Hype Cycle because it has reached the Plateau of Productivity and achieved mainstream adoption levels.

On the Rise

AI Security Posture Management

Analysis By: Dennis Xu, Neil MacDonald, Andrew Bales

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Definition:

AI security posture management (AISPM) helps enterprises scan their infrastructure, AI development and engineering pipelines, data pipelines, and run-time environment to identify risks in both third-party and custom-built AI applications and AI agents. AISPM provides visibility into misconfigurations and vulnerabilities related to AI and helps maintain a secure AI environment. AISPM is offered by a wide range of solutions such as AISP, CNAPP, DSPM, SSPM, and application security products.

Why This Is Important

AI risks can arise when third-party or enterprise-deployed AI applications are unknown or misconfigured, or contain vulnerabilities. Classic SPM tools do not cover AI applications. AISPM tools address this by helping to evaluate AI infrastructure, applications, and agents for exposures and risk, prioritizing the findings and aiding remediation. More comprehensive tools are offered as a part of a broader AI security platform solution to secure AI applications from build time to run time.

Business Impact

Business teams are rapidly deploying AI and AI agents to drive innovation and efficiency. Cybersecurity leaders need to support this rapid adoption while ensuring that critical systems and enterprise data remain protected. Misconfigured AI infrastructure, excessive access privileges, or rogue agents can expose sensitive assets and introduce compliance risks, damage competitive advantage, and harm organizational reputation and human safety.

Drivers

- AI assistants and AI agents are increasingly deployed across the enterprise infrastructure, often without proper assessment of their configurations, permissions, exposures, or data they access. This creates a need for security posture management focused specifically on AI.
- AI agents add complexity beyond traditional chatbots because they can access a wide range of native and third-party tools, increasing their exposure to various agentic and tool-based security risks.
- There is a growing need to evaluate the confidentiality and commercial sensitivity of data being used in the models to prevent noncompliance, privacy violations, data poisoning risk, and security breaches.
- AISPM helps organizations answer critical questions, such as “What AI applications and AI agents are in use across my organization, and do they represent unknown or unacceptable security risks?”

Obstacles

- Realizing the full benefits of AISPM requires cooperation among business, AI engineering, data and analytics, legal, and cybersecurity leaders to establish a multidisciplinary governance approach.
- AISPM relying on a rule-based, deterministic approach can struggle to effectively identify security risks in autonomous agents where the agent execution steps are planned at run time.
- An AI agent can update itself during run time, such as by rewriting or creating new agent skills, which challenges traditional build-time security posture management approaches.
- AISPM products struggle to support the vast array and rapid change of AI agents. They also have limited or no capability to assess models developed in-house or open-source models for correct configuration or vulnerabilities.
- Gaining visibility into the data used for model creation in SaaS is challenging, as integrating with the other tens of thousands of SaaS applications via APIs or agent communication protocols can be complex.

User Recommendations

- Start with the incumbent vendor’s AISPM feature to fulfill your AI security needs.

- Prioritize AISPM tools that continuously evaluate the security posture of AI agents in run time, beyond the build-time posture check.
- Prefer an AISPM that is offered as part of a broad and integrated AI security platform to enable end-to-end AI security defense.
- Ensure the vendor can support all types of AI and AI agents used and developed by the enterprise.
- Assess capabilities to identify security risks from all relevant AI models, assistants, and agents.
- Identify capabilities that also track overprivileged or vulnerable agents and data workloads and any potential misconfigurations that impact sensitive data used in training and reference data, libraries, APIs, and across AI data pipelines.

Sample Vendors

Cato Networks; Cisco Systems; Cyera; Noma Security; Obsidian Security; Palo Alto Networks; Reco; Varonis Systems; Google (Wiz); Zenity

Gartner Recommended Reading

[Top Strategic Technology Trends for 2026: AI Security Platforms](#)

[AI Security Platforms Are Core to Cybersecurity Revenue Growth Strategy](#)

[AI Vendor Race: Palo Alto Networks Is the Company to Beat in AI Security Platforms](#)

[AI Vendor Race: GMs: Win the AI Security Battle With an AI Security Platform](#)

[Create a Data Risk Assessment That Drives Business Decisions](#)

[Quick Answer: What Are the Major Data Risks for AI?](#)

Influence AI

Analysis By: Andrew Frank

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Definition:

Influence AI is the creation and deployment of models and algorithms designed to automate elements of digital experience that guide people's choices at scale by learning and aligning with their intents and motivations. For instance, an AI agent may achieve a desired outcome for its organization by framing a set of recommendations based on a user's inferred motives and predicted behaviors, or inducing an external agent to do so.

Why This Is Important

AI already influences customer choices through answer engines and AI sales agents; marketers who fall behind risk losing control of brands and reputations as competitors and AI platforms seek the upper hand in AI-led influence operations. As more buyers entrust agents to act as marketplace advisors and commerce proxies, established trust factors destabilize and new patterns of persuasion emerge.

Business Impact

Pressure is mounting on marketers to turn AI adoption into topline growth, pushing them to look beyond operational efficiency toward AI uses that increase sales and loyalty. As LLM providers amass pervasive power over consumer choices, marketers need to adopt new disciplines of data, content and creativity to forge new trust relationships both directly and through these mediated channels by using AI and scaled personalization to align offers and dialogs with people's values, mindsets and circumstances.

Drivers

- **Rising pressure on AI to deliver growth.** AI is under pressure to deliver more tangible benefits. Executive leaders expect marketing to use AI to meet growth goals for revenue and profit. Influencing customer choices offers the clearest path toward meeting these expectations.
- **Evolving user expectations.** Consumers and business buyers increasingly expect digital experiences tailored to their preferences and needs. Influence AI enables hyperpersonalization by dynamically adapting recommendations, content, and interactions to context and history.
- **Advancements in AI and data analytics.** Improved modeling techniques, natural language processing and behavioral analytics provide the technical foundation to infer user intent and motivations at scale.
- **Marketing's head start in AI.** Marketing's dependence on communication and content to influence behavior made it an early adopter of machine learning and GenAI for segmentation and creative optimization. This builds confidence in the possibilities of AI agents for influence.
- **Omnichannel engagement.** Influence AI can operate across multiple channels (web, mobile, voice, etc.), ensuring consistent and context-aware guidance throughout the customer journey.
- **New intermediaries.** AI answer engines and sales agents will amplify Influence AI's impact on consumer choices with more personalized, persistent and persuasive guidance. This raises the stakes for marketers to understand and engage these intermediaries by targeting and measuring influence effects.
- **Competitive imperative.** AI's rapidly growing digital market presence is reaching a tipping point where the gap between leaders and laggards produces clear wins and losses. Companies that fail to adapt to AI's influence face strategic peril.
- **Governance adoption.** Gartner surveys and inquiries show AI governance is a top enterprise action and priority, building tolerance for potentially risky experiments with safety policies and oversight.

Obstacles

- **Ethical and privacy concerns and backlash.** Influence AI operates by inferring and acting on user motives, raising significant questions about manipulation, consent, and transparency. Regulatory scrutiny is likely to increase as techniques improve. Models may inadvertently reinforce biases, make unfair recommendations or over-personalize, exposing brands to significant backlash.
- **Explainability and trust.** Users and regulators demand transparency in how AI-driven recommendations are made, especially in regulated industries. Black-box models can undermine trust and hinder adoption. Reports of backlash against excessive personalization may reinforce hesitation.
- **Data quality and concentration.** Influence AI relies on large volumes of high-quality behavioral data. Fragmented or poor-quality data limits effectiveness. Hyperscalers with massive volumes of behavioral data need to be cautious of regulatory crackdowns.
- **Complexity and risk.** Deploying, monitoring, and continuously improving influence models requires significant resources, specialized talent, and a robust operational infrastructure. While agentic AI promises to democratize creation of autonomous agents, they can backfire in complex or adversarial situations.
- **Machine customers.** Buyers' use of agent proxies may block AI's psychological influence, favoring differentiation through product innovation and operational efficiency over persuasive interactions.

User Recommendations

- Establish or locate a Trust Council within your organization (see [Every Enterprise Needs a Trust Council to Combat the Deluge of AI Disinformation](#)) and add Influence AI governance to its agenda. Articulate how to distinguish beneficial nudges from those that exploit vulnerabilities. Document potential use cases and subject them to neutral, non-commercial ombuds review.
- Recruit user test and focus groups from within and outside your organization for AI influence-related market research, or seek providers for such projects. Be transparent about goals and technologies. Apply strict privacy controls and use synthetic data or digital twins to supplement findings. Devise extensive, adversarial test cases for experimental designs.
- Probe partners, vendors and martech applications for signs of unapproved AI Influence techniques emergent in customer interactions. They may produce strong, but ethically or legally compromised, marketing results. Require compliance with influence policies in service agreements and respond quickly and decisively to perceived violations. Assess competitors' and platform providers' activities.

Gartner Recommended Reading

[Top Marketing Trends for 2024: AI for Privacy, Personalization and Influence](#)

[CMO's Guide to Protecting Trust in Content With TrustOps](#)

At the Peak

Federated Machine Learning

Analysis By: Bart Willemsen

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Obsolete

Definition:

Federated machine learning is a decentralized approach to machine learning that enables multiple clients to collaboratively train a shared model without sharing their raw data, enhancing privacy and overcoming data transfer constraints. This process involves clients performing local training on their data and sending updates to a central server, which aggregates these updates to refine the global model.

Why This Is Important

Federated learning facilitates the decentralized training of ML algorithms without disclosing sensitive or personal data. It enhances model personalization and contextualization by allowing local data processing on IoT, isolated enterprises, and end or edge devices, including vehicles. It also facilitates collaborative learning model building across data silos and helps improve model accuracy and robustness by exchanging model coefficients rather than identifiable data.

Business Impact

Federated learning reduces costs, ensures compliance with data regulations and preserves privacy and security. It does so by allowing enhancements to model accuracy and operational efficiency without transferring sensitive data but instead based on exchange of model coefficients. Federated learning is used on a project basis across healthcare, finance, telecommunications, and IoT in diagnostics, fraud detection, network optimization and personalized services.

Drivers

- The proliferation of privacy and legislative regulations requires protection of local data. Federated learning is one type of solution to protect data privacy.

- Growth of edge computing and IoT requires data processing and model training directly and in real time on distributed devices, reducing the need for data centralization.
- Exploding data volumes and data gravity make large-scale data transfers challenging, but federated learning can resolve these challenges by processing data locally and avoiding the need for centralization.
- Reinvigorated interest in privacy-enhancing-technology (PET) use for compliance and risk reduction reasons brings it back under consideration for select use cases.
- Centralized architectures of machine learning have certain limitations, including scalability, power consumption and latency issues. Federated learning can help address these challenges.
- Advancements in enabling technologies, such as differential privacy and blockchain, extend federated learning adoption by enhancing privacy and facilitating decentralized coordination.

Obstacles

- Adoption is hindered by lack of awareness and trust issues that are rarely solved without adding one or more other PET to the architecture mix in projects.
- The streamlined pipeline needed for federated learning spans DataOps, ModelOps, deployment and continuous tracking/retraining. This maturity is often missing.
- Creating new, accurate and unbiased models from local model improvements can be cumbersome and readily available models lure users into less-controlled adoption.
- Diverse device capabilities cause client challenges, such as requiring asynchronous training and adaptive algorithms to handle variability in computation and connectivity, which isn't easily achieved.
- Ensuring equitable model performance across diverse clients is challenging due to statistical heterogeneity, necessitating fairness-aware algorithms to address disparities during training and aggregation.

User Recommendations

- Focus on scenarios where data decentralization is required by privacy regulations, competitive sensitivities or logistical challenges, leveraging its advantage in creating decentralized smart services with diverse user data.
- Use centrally pretrained models to start federated learning processes, capturing general patterns and focusing federated rounds on fine-tuning and personalization with local data.
- Begin with small pilot projects or simulations to gain experience and understand algorithm and framework nuances before large-scale deployments.
- Establish governance, trust and incentive mechanisms early to facilitate partnerships with suppliers, customers or peers, enhancing operational efficiency and product offerings through shared insights.
- Implement strategies to prevent excessive divergence from the global model, using central reference models or algorithms like FedProx to maintain cognitive cohesion.

Sample Vendors

Devron; Duality; Ederlabs; FedML; Google; Intel; TripleBlind

Gartner Recommended Reading

[Reference Architecture for Federated Analytics](#)

[Use-Case Guidance for Choosing Privacy-Enhancing Technologies](#)

[What Data Governance Leaders Must Know About Privacy-Enhancing Technologies \(PETs\)](#)

[Privacy Imperatives for Executive Leaders](#)

5G Network Security

Analysis By: Sylvain Fabre, Nat Smith, Peter Liu

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Definition:

The 5G network offers security enhancements over 4G like mutual authentication and enhanced subscriber privacy. However, 5G's cloud-native architecture, edge processing, increased network capacity and low latency can amplify risks, given the proliferation of vulnerable IoT devices. Although 3GPP strengthens the air interface security, the distributed nature of 5G, with network function security left to various international forums, can lead to inconsistencies and a wider attack surface.

Why This Is Important

Securing 5G networks is a priority for CSPs' public networks. The increasing adoption of private mobile networks (PMNs) deployments, vertical applications, cloud architecture, and massive Internet of Things (IoT) connections is creating new vulnerabilities and challenges, such as distributed denial of service (DDoS) attacks. 5G network security offers a single, unified framework to manage how users and devices are authenticated when connecting to the network, regardless of their cellular 3rd Generation Partnership Project (3GPP) or other (non-3GPP) access methods.

Business Impact

5G security:

- Enhances subscriber privacy with unified authentication and Subscription Concealed Identifier (SUCI) encryption, mitigating international mobile subscriber identity (IMSI) catchers.
- Adds user plane integrity protection and a security edge protection proxy (SEPP) for roaming.
- Secures network access of external entities via network exposure function (NEF) and verifies users and devices across different access types through Extensible Authentication Protocol-Authentication and Key Agreement (EAP-AKA).
- Enables application layer security via service-based architecture (SBA).
- Verifies device location using home network public key (HNPK) authentication when in a visited network, preventing spoofing attacks.
- Mitigates bidding-down attacks and neutralizes fake base stations (aka IMSI catchers).

Drivers

- The wider industrial 5G ecosystem demands security tools with varying capabilities to deliver SLAs and end-to-end security for diverse use cases. This includes an application-focused approach at the edge, ensuring security is tailored to specific industry needs.
- Liability exposure in enterprises deploying private 5G necessitates robust network security.
- Increasing regulations around user data privacy are a key driver, further amplified by legislation like the EU's Cyber Resilience Act (CRA) and Network and Information Systems (NIS2) directive, which mandate stronger cybersecurity measures and reporting obligations.
- Heightened scrutiny on 5G infrastructure vendors and competition based on perceived security levels continue to drive security enhancements.
- Interconnection with legacy networks, satellites and less secure IoT devices requires robust mechanisms to handle such connections safely. The proliferation of poorly secured IoT devices significantly expands the attack surface.
- Delivering cloud-based 5G infrastructure and services (core, slicing, edge) necessitates comprehensive cross-domain security. Network virtualization and cloudification, while offering flexibility, introduce inherent security challenges.
- Demand for edge enterprise solutions will accelerate 5G security adoption, particularly with the need to secure application workloads closer to the data source.
- 5G SBA infrastructure virtualization, automation, orchestration and multivendor solutions increase risk exposure for CSPs and enterprises.
- CSPs can offer differentiated security options, such as using slicing, as a value-added service.
- Growing reliance on 5G for critical infrastructure (smart grids, transportation, healthcare) and the current geopolitical climate make robust security paramount to protect essential services and public safety.

Obstacles

- 5G end-to-end security presents new challenges for operators (e.g., hacks on misconfigured containers) but an IT-centric approach to network functions and back-end/core can improve security posture.
- Some security settings are optional, so implementation security levels may vary.
- Chief information security officers (CISOs) may not realize the need to secure their on-premises 5G PMNs or understand additional threats in 3GPP-based infrastructures. They rely on their vendors, who tend to focus on performance and project delivery rather than security.
- Supply chain security concerns remain for 5G as open-source approaches add complexity. Not all network components are (yet) known and defense models may have to broaden as implementation progresses.
- Backward compatibility with 4G/LTE means some legacy security issues affect 5G; for example, general packet radio service (GPRS) Tunneling Protocol [GTP] attacks.
- 5G runs on commodity hardware and containerization, which creates a larger attack surface compared to 4G.
- Deployment scenarios for private 5G have different risk profiles and security needs.
- The increased network complexity of 5G, coupled with lack of integration with legacy infrastructure and standardization, poses further obstacles.

User Recommendations

- Avoid overreliance on 3GPP 5G security. Solve all 4G issues and continue ongoing security initiatives, since many 5G security features are designed to protect the network rather than secure user data.
- Assess and anticipate increased risk of attacks such as DDoS by improving 5G network design infrastructure, endpoints, and recovery procedures.
- Implement layered DDoS defense by using best-of-breed scrubbing, a cloud-web-application firewall, bot mitigation, DNS protection, ISP and on-premises DDoS appliances.
- Initiate effective mitigation by correlating network traffic, application availability, server performance, and other relevant data.
- Complement 5G infrastructure security with distributed AI/ML anomaly detection algorithms for zero days, covering newer domains such as core slicing.
- Implement edge protection concepts used in application security for elasticity, focusing on app-layer protection.
- Implement 5G end-to-end security by managing algorithm strength, secret keys negotiation, confidentiality protection, cross-domain slice orchestration, and heterogeneous network layers.
- Use a clear model of shared responsibility between operators, vendor partners, and end-user organizations.

Sample Vendors

A10 Networks; Fortinet; Microsoft; Nokia; OneLayer; Palo Alto Networks; SentinelOne; Trend Micro; VIAVI Solutions; Zscaler

Gartner Recommended Reading

[8 Critical Functionalities for Enterprise 5G Private Mobile Network Management and Orchestration](#)

[Navigating the Complexities of 5G Private Network Security](#)

[Market Trend: CSPs Need Tailored Security for 5G Networks](#)

Confidential Computing

Analysis By: Bart Willemsen, Mark Horvath

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Early mainstream

Definition:

Confidential computing is a hardware security mechanism that executes code in a trusted execution environment (TEE), also called an enclave. Enclaves isolate and protect in use the confidentiality of data, content, AI, and other workloads from the host system and its owners, and may also provide code integrity and attestation.

Why This Is Important

- Confidential computing combines a chip-level TEE with key management and cryptographic protocols. It enables computation where cooperation among groups is critical without sharing identifiable data or intellectual property.
- Adoption of the cloud and the increased viability of enclave technology enable the trusted use of data in a third-party's infrastructure, including defense against the provider.
- Confidential computing can address data residency and sovereignty concerns.

Business Impact

- Confidential computing helps mitigate confidentiality concerns, which is a major security barrier to adopting cloud services and AI. This is critical for regulated businesses or any organization concerned about unauthorized third-party access.
- Confidential computing enables a level of data confidentiality and privacy controls among competitors, data processors, and analysts that is difficult to achieve with traditional cryptographic methods.

Drivers

- Concerns about potential access to personal data by cloud service providers (CSPs) and external actors, such as foreign governments, increasingly influence cloud adoption.
- Global data residency restrictions and sovereignty concerns drive the need to isolate content from the CSP with a level of independent assurance.
- Competitive concerns — not just about personal data, but also intellectual property (IP) — are spurring the adoption of confidential computing. This includes confidentiality protection against any third-party access, protection of the processing method (including algorithmic functions), and protection of the data itself. As a result, CSPs have at least an implied competitive interest in demonstrating they do not have or need access to such sensitive workloads when providing infrastructure only.
- Hyperscaler cloud providers offer options for virtualized confidential computing, enabling apps to run without recoding or refactoring. Domains like Microsoft Azure can enable confidential computing with as little as a single line in a configuration file.
- The protection provided by confidential computing facilitates the increased use of externally facilitated AI capabilities. It also allows, for example, the use of AI for text summarization or translation without sharing key information with the provider (“confidential AI”).

Obstacles

- It is rare to find SaaS or business-process-as-a-service (BPaaS) vendors offering integration — which reduces protection choices.
- Lack of trained staff or understanding of implementation methods can weaken deployment (e.g., incorrect key handling, overlooked side-channel vulnerabilities).
- Confidential computing isn't always plug and play. It varies by vendor and technology and may require a high level of effort while delivering only marginal improvement over controls like Transport Layer Security (TLS), multifactor authentication (MFA), and customer-controlled key management services.
- Offerings directly from CSPs vary in robustness, performance, and reliability. Not every solution labeled “confidential computing” offers equal protection or facilitates multicloud.
- Confidential computing using cloud-native key management (KM) may be at risk of inadequate privacy because the CSP manages and has access to the keys. Thus, using third-party key management as a service (KMaaS) becomes more important.

User Recommendations

- Determine whether confidential computing affects application performance by designing (or duplicating) a sample application using one of the available abstraction mechanisms and deploying it into an instance with an enclave. Perform processing on datasets that represent the sensitive information you expect in real production workloads. This step will help you surface and mitigate negative results.
- Examine confidential computing for projects where multiple parties, who might not necessarily trust each other, need to process but not access sensitive data to achieve a result that benefits all parties. In this scenario, none of the parties should control the TEE.
- Look for vendors that offer integration with a broader enterprise key management system, as well as complementary encryption and privacy-enhancing computation techniques.

Sample Vendors

Alibaba Cloud; AMD; Fortanix; Google; IBM; Intel; Microsoft; NVIDIA; VectorZero

Gartner Recommended Reading

[Use-Case Guidance for Choosing Privacy-Enhancing Technologies](#)

[Top Strategic Technology Trends for 2026: Confidential Computing](#)

[CISO Effectiveness: Drive Good Privacy Practices by Consistently Asking Key Questions](#)

[Market Guide for Cloud Infrastructure Sovereign Solutions](#)

[What Data Governance Leaders Must Know About Privacy-Enhancing Technologies \(PETs\)](#)

[Top Strategic Technology Trends for 2026](#)

Homomorphic Encryption

Analysis By: Bart Willemsen, Mark Horvath

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Definition:

Homomorphic encryption (HE) uses cryptographic algorithms to enable computations on data while it remains in a protected state. Partial HE (PHE) supports only limited use cases, such as subtraction and addition, but with little performance impact. Fully homomorphic encryption (FHE) supports a wider range of repeatable and arbitrary mathematical operations; however, there is usually a performance trade-off. Forthcoming hardware accelerators are expected to improve this.

Why This Is Important

HE offers an advance in privacy and confidential data processing. Benefits include the ability to:

- Perform analytics on data while in an encrypted state, so that the processor never sees the data in the clear, yet delivers accurate results.
- Share and pool data among competitors (secure multiparty computation [SMPC]).

- Share all or part of users' data, while protecting their privacy.
- Base the system on lattice encryption, which is used by some postquantum cryptography (PQC) algorithms.

Business Impact

Even in restricted form (PHE), HE enables businesses to use data, send it to others for processing and return accurate results, without fear it will be compromised or stolen. Any data intercepted by a malicious actor is encrypted and unreadable, even by the coming generation of quantum hardware. Applications include:

- Encrypted search
- Data analytics
- Machine learning (ML) model training
- Multiparty computing
- Securing long-term record storage without concerns about unauthorized decryption

Drivers

- Increased geopolitical volatility and a focus on data and operational sovereignty worldwide are forcing organizations to protect data in use, rather than merely in transit and at rest.
- Globally, privacy and data protection frameworks demand that more precise attention be paid to sensitive data. As a result, data pooling, sharing and cross-entity analysis use cases increasingly benefit from forward-looking and sustainable technologies such as HE.
- Aside from primarily financial use cases (e.g., cross-entity fraud analytics), other industries benefit as well. One example is healthcare, where analysis of sensitive data across various entities happens often with data protected while in use (e.g., clinical trials and cross-entity knowledge sharing).
- Solving issues of trust and cooperation with SMPC will benefit internal and external protection of data.
- With the adoption of new NIST PQC algorithms, the homomorphic properties of lattice are being explored by more clients. Timely adoption of HE in data protection will sustainably protect data against the threat of “harvest now, decrypt later” attacks.
- Previously underestimated because of latency challenges, hardware accelerator technology currently forthcoming is reinvigorating interest and viability.

Obstacles

- The application of various forms of HE to common use cases leads to a degree of complexity, slows operations and requires highly specialized staff. While remarkable progress in commercial uses has made it a more viable option, client adoption is still lagging. Several innovations for dedicated hardware accelerator chips are being released in response to this obstacle.
- Some scenarios will never be a good match for HE — for example, those that require security in components beyond analytics and processing, such as production databases and proprietary algorithms.
- The market’s unfamiliarity with this technology stands in the way of speedy adoption.

User Recommendations

- Brainstorm opportunities with your technical and executive teams. For example, come up with a list of five to 10 practical use cases for HE to improve the adoption of core solutions.
- Treat potential HE projects as experiments, keeping in mind the early stage of the technology's development and the significantly non-real-time nature of HE products. Consider these experiments proofs of concept to build experience until the technology matures.
- Continue with existing security controls. HE does not necessarily negate the need for other security controls, observance of data residency requirements or access control.
- Assess the core benefits of using HE in combination with other quantum-safe or privacy-enhancing computation techniques.
- Integrate in-use protection via forms of HE into messaging and third-party analytics services.
- Assess the merits of piloting HE by using a vendor's solution, which could offer functionality without the time investment associated with a custom solution.

Sample Vendors

Duality; Enveil; IBM; LiveRamp; Microsoft; Niobium Microsystems; Thales; Zama

Gartner Recommended Reading

[What Data Governance Leaders Must Know About Privacy-Enhancing Technologies \(PETs\)](#)

[Use-Case Guidance for Choosing Privacy-Enhancing Technologies](#)

[What Executives Need to Do to Support the Responsible Use of AI](#)

[Top Strategic Technology Trends for 2025: Postquantum Cryptography](#)

[Top Strategic Technology Trends for 2026: Confidential Computing](#)

Sovereign Cloud

Analysis By: Rene Buest, Gregor Petri

Benefit Rating: Transformational

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Sovereign cloud is the provision of services by a cloud provider to ensure that data, operational and technological sovereignty meet the sovereignty requirements of its intended customer base. Sovereign cloud ensures that data, applications, infrastructure and operations are immune from control by external jurisdictions and are fully protected from foreign government interference and access.

Why This Is Important

Public sector and commercial organizations depend on cloud services, leading to greater demand for control and operational autonomy and thus offerings that provide a high degree of sovereignty. In many jurisdictions, data residency, data protection requirements and privacy laws are increasing. Increasing geopolitical tensions, different economic ideologies and proliferating cybersecurity risks from a variety of directions, including state actors, are steadily elevating interest in sovereignty.

Business Impact

Concerns about data, operations and technology being controlled by foreign-owned cloud providers are increasing the demand for true sovereign clouds. Regulators are introducing requirements to restrict organizations from using services provided by vendors outside their own country. These regulations impact choices organizations make when using cloud services. Consequently, end users may face difficulties in accessing resources and innovation to support their digital business initiatives.

Drivers

- Ongoing geopolitical turmoil and uncertainty accelerate sovereignty concerns, particularly among governments, regulated industries and critical infrastructure, especially in Europe, the Middle East, Latin America and Asia/Pacific. Organizations in these regions are increasingly concerned about their dependence on foreign cloud infrastructure providers and the sovereignty of data, applications, infrastructure and operations hosted in foreign-owned cloud service offerings.
- Concerns over digital sovereignty lead countries and confederations of states to pursue their own sovereign digital strategies. Many commercial organizations are increasingly impacted by digital geopolitical issues, such as trade disputes, legislation coming from one country that impacts global operations, and government-imposed restrictions on the acquisition and use of digital technology.
- Given the critical role that digital technology plays in the economic, cultural, social and government institutions of a country, great power competition sharpens the focus on digital sovereignty. Particularly in a world of public clouds, every organization needs to be in control to be able to address its own risk management, compliance policies, culture or business strategy and to address multiple regulatory requirements.
- Organizations are increasingly concerned that cloud and digital infrastructure will be used as a tool to exercise pressure, coerce, and enforce political influence. This has led to a rethinking of cloud strategies and the acceleration of geopatiation. The call for a higher degree of digital sovereignty aims to preserve autonomy by developing and maintaining one's own capabilities and avoiding dependencies.

Obstacles

- Domestic cloud providers are currently not able to fully compete with hyperscale cloud offerings, as the hyperscaler capabilities still exceed most domestic cloud offerings. Considerable technical obstacles exist if domestic clouds are expected to deliver the maturity and level of the scalability and functionality of hyperscale offerings.
- The growing influx of sovereign cloud options from global cloud providers, all based on various approaches and delivery models, is making it hard for end users to understand the actual value of each offering and how they differ.
- Individual governments defining their own requirements for sovereign cloud offerings may lead to compliance regimes that impede public cloud scale and innovation and lead to conflicting regulations.

User Recommendations

- Define the sovereignty requirements and agree on the sovereignty level required by your organization by country, vertical industry and workload type by working with your legal counsel and risk departments.
- Differentiate between different sovereign cloud approaches by the type of workload, data and infrastructure when making cloud deployment decisions. Classify various delivery models between the cloud provider approaches. Establish a consistent, repeatable and defensible process when assessing sovereign offerings.
- Evaluate cloud services for offerings that can be provided locally and leverage third-party solutions to protect data at rest and in transit.
- Evaluate the different cloud provider approaches to meet your defined sovereignty requirements by quantifying their technological limitations that might limit the delivery of your mission-critical solutions.
- Consider sovereignty as only one of several critical dimensions that must be considered when evaluating a cloud service provider.

Gartner Recommended Reading

[Build a Sovereign Cloud Strategy](#)

[Prioritize on Sovereign Capabilities When Choosing Your Cloud Provider](#)

Sliding into the Trough

AI TRiSM

Analysis By: Avivah Litan, Jeremy D'Hoinne, Bart Willemsen, Lauren Kornutick

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Definition:

AI trust, risk, and security management (AI TRiSM) comprises four layers of technical capabilities that support enterprise policies for all AI use cases and help assure AI governance, trustworthiness, fairness, safety, reliability, security, privacy, and data protection. The top two layers — AI governance, and AI runtime inspection and enforcement — are purposefully directed to support AI. The bottom two layers represent information governance and traditional technology in the context of AI.

Why This Is Important

AI brings new trust, risk, and security management challenges that conventional controls do not address. Top concerns for enterprises include data compromise, third-party risks, undesired outcomes, and the need to ensure enterprise AI actions align with deployment purposes. Organizations must retain independence from any single AI model or hosting provider to ensure scalability, flexibility, and trust (and cost control as a derivative benefit) as AI markets rapidly change.

Business Impact

Those failing to manage AI risks experience project failures, underperformance, and compromised data. Inaccurate, unethical/unintended AI outcomes, and interference from malicious actors can result in financial and reputational loss, liability, and social harm. AI underperformance can also lead to poor business decisions and uncontrolled costs. Applying AI TRiSM improves project timelines, operational precision, product durability, and enhanced customer trust and overall AI investment ROI.

Drivers

- The increasing use of AI, GenAI, and AI agents is limited by a lack of trust in AI as a safe and ethical option for supporting critical business processes. Fewer than 5% of respondents fully trust their vendor's hallucination safeguards, AI security, and governance controls according to the 2025 GenAI and Agentic AI in Enterprise Apps Survey (see [Assessing the Impact of Generative AI and Agentic AI In Enterprise Applications](#)).
- Enterprises face multiple AI risks and are most concerned with data compromise, third-party risks, and inaccurate or unwanted output.
- Large language models are nondeterministic, and their output and behavior are unpredictable.
- Regulations for AI risk management (such as the AI Act in Europe or Local Law 144 in New York) are driving businesses to institute measures for managing AI risk. Such regulations define new compliance requirements that organizations will have to meet on top of existing ones, like those pertaining to privacy protection.
- Malicious hacks against enterprise AI are still uncommon, while incidents of unconstrained harmful chatbots are well-documented, and internal oversharing data compromises are prevalent.
- AI agents pose new risks of aberrant behavior deviating from human instructions and greatly expand the attack surface.
- The rapid proliferation of AI agents will create more need for governance than human-in-the-loop oversight can fulfill alone.
- User demand for AI TRiSM solutions is increasing, and providers of all sizes are competing for this new enterprise business. Many AI TRiSM startups have been acquired by large security vendors.
- Some organizations are mostly concerned with security and risk mitigation, while others also focus on supporting ethical or safe practices and regulatory compliance.
- AI trust, risk, and security issues surface organizational silo issues, pushing teams to realign to solve problems that cross departmental boundaries and to implement technical measures that address them.

Obstacles

- Adopting AI TRiSM technology is often an afterthought. Many organizations don't consider it until AI applications or agents are in production, when it becomes challenging to retrofit.
- Many enterprises are resource-constrained and don't have the skills or capacity to implement AI TRiSM.
- Enterprises often rely on their incumbent vendors to provide AI TRiSM capabilities, although they often lack it and must rely on vendor licensing agreements to ensure their confidential data remains private in the host environment.
- Most embedded software and SaaS services use AI but they often do not support APIs to third-party AI TRiSM products that can enforce enterprise policies.
- AI TRiSM requires a cross-functional team, including legal, compliance, cybersecurity, IT, and data analytics staff, to establish common goals and use common frameworks. Coordination between these teams might be lacking, leading to competing or overlapping technologies. This fragmentation of TRiSM efforts hinders AI progress.

User Recommendations

- Set up an organizational unit to manage AI TRiSM and include members with a vested interest in AI projects.
- Discover and inventory all AI used in the organization, leveraging the capabilities of TRiSM vendors who support this.
- Define acceptable use policies that are flexible with just the right level of granularity for agile enforcement.
- Revisit and implement data classification, protection, and access management across all enterprise information that can potentially be used by AI. Collaborate across the teams involved in governance.
- Work with legal and compliance to contractually define accountability for unacceptable AI use or behavior in third-party-embedded AI applications.
- Obtain vendor attestation to meet legal requirements.

- Evaluate and implement layered AI TRiSM technology to support and enforce policies across all AI in use. This includes enterprise-owned services, along with controls offered by incumbent model and platform providers. However, do not solely rely on the latter.

Gartner Recommended Reading

[Market Guide for AI Trust, Risk and Security Management](#)

[Top Strategic Technology Trends for 2024: AI Trust, Risk and Security Management](#)

[The Cybersecurity Leader's Playbook: Navigating the EU AI Act](#)

[How to Build a Lightweight Organizational Structure for AI Governance](#)

[Market Guide for Guardian Agents](#)

Differential Privacy

Analysis By: Bart Willemsen

Benefit Rating: Low

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Differential privacy is an approach to using data and extracting information while withholding or distorting certain (personally) identifiable elements in the dataset. When applied correctly, this prevents the disclosure of personal data while ensuring the resulting analysis does not significantly change informationwise. It uses mathematical algorithms that randomly insert noise into the data and add parameters for distinguishability, closeness and diversity of outcomes at each query.

Why This Is Important

Privacy and data protection, as well as a variety of confidentiality concerns, impact how to use personal data for trend analyses and serving content or personalizing recommendations. As regulatory measures are employed to prevent unauthorized use of people's information, organizations are looking for ways to protect personal data while still using that data. Differential privacy is one of the available privacy-enhancing technologies (PETs) that aid in doing so.

Business Impact

Business data holds value and much of it is personal data. Regulations constraining the use of identifiable data are now in effect worldwide. Liability for misusing personal data can be substantial. There are several techniques to address problems in preserving privacy when training AI models. Differential privacy is among the easier, accessible ones and prevents singling out individuals while supporting meaningful analysis of aggregated data.

Drivers

- Differential privacy helps reduce risk but also unlock data for AI and analytics use cases. This relates to both creating and sharing insights internally in the organization, as well as externally across entities.
- Organizations can extract valuable insights from data while respecting ethical standards and adhering to regulatory guidelines on the use of (in)directly identifiable data.
- More restrictive requirements are forthcoming on state and country levels regarding the use of personal data in training of algorithms and on how, in turn, algorithms handle personal data.
- Brand reputation and customer trust are significantly damaged with data breach or misuse, but positively impacted when leaders are vocal about their proactive privacy protection technologies.
- With differential privacy, source data is not altered because the answer to each query is treated “on the fly,” protecting the data in use while retaining source data integrity. This is a differentiating factor as opposed to alternative PET options. It can also be used at the information ingestion point, to prevent identifiable data processing in service of data minimization and privacy by design efforts.
- Information value can be controlled relatively easily via the privacy budget in differential privacy, delivering the desired level of anonymity. This option makes it among the more accessible of PETs.

Obstacles

- Tools that reference the use of differential privacy are rarely comparable.
- PETs vary in effectiveness. Organizations often lack a framework to consistently determine the appropriate approach based on use-case requirements, technology maturity and fit.
- Many deployments solely address (k-)anonymity. Deployments with added control measures to (l-)diversity and (t-)closeness of outcome deliver more robust results, but require more work too.
- Lack of familiarity with differential privacy — and the skilled staff to deploy and manage it — hinders adoption. This is exacerbated by how jurisdictions differ in defining “anonymous” or “pseudonymous” data.
- Lack of transparency around setting of the privacy budget (the extent to which controls are implemented, or “where does one put Epsilon”) undermines trust, whereas increased transparency elevates trust.

User Recommendations

- Prioritize the use of differential privacy techniques to decrease the likelihood of sensitive data exposure as they are among the easier PETs to get started with.
- Use a privacy impact or data protection impact assessment to establish whether additional means are necessary and relevant to the use case.
- Compare differential privacy with other PETs when operating in high-performance environments that require a high level of precision in analytics models.
- Prioritize differential privacy techniques if you’re operating in a highly regulated industry, such as financial services or healthcare.
- Explore differential privacy techniques when using data across regions where privacy regulations may vary and always be transparent about where you have set the privacy budget.

Sample Vendors

Dynamo AI; Immuta; LinkedIn (Tumult Labs); LiveRamp; Snowflake (LeapYear)

Gartner Recommended Reading

[Predicts 2025: Privacy in the Age of AI and the Dawn of Quantum](#)

[CISO Effectiveness: Drive Good Privacy Practices by Consistently Asking Key Questions](#)

[Use Consistent Privacy and AI Impact Assessments to Ensure Control Criteria](#)

[Tool: AI and Privacy Impact Assessment](#)

[Use-Case Guidance for Choosing Privacy-Enhancing Technologies](#)

Data Security Governance

Analysis By: Rahul Balakrishnan, Joerg Fritsch, Andrew Bales, Chiara Girardi

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Data security governance (DSG) is the specification of policies, processes, decision rights, and organizational accountability needed to assess, prioritize, and mitigate business risks related to data security, privacy, and compliance, while ensuring data security objectives support business priorities.

Why This Is Important

The proliferation of data across on-premises and multicloud environments increases business risks, including security, privacy, and compliance concerns, making it challenging to create and enforce consistent data security policies for prioritized risk minimization. DSG helps organizations navigate the tension between effective risk mitigation and the pursuit of business priorities, enabling both through consistent, well-governed data security policies.

Business Impact

DSG provides a clear framework for managing data access and usage while enforcing security and privacy measures to manage data security risk. It requires collaboration among chief information security officers (CISOs), chief data and analytics officers (CDAOs), and business leaders. This collaboration helps organizations:

- Overcome cross-functional communication barriers

- Reduce duplication of efforts
- Identify when governance initiatives may be making it harder for employees to get work done

Drivers

- Guiding the implementation of consistent data access controls across a portfolio of disparate datasets requires risk-focused data security policies.
- Ensuring compliance with emerging AI regulations and mitigating the amplified data security risks introduced by GenAI tools require creating and implementing consistent, DSG-driven data security policies.
- Business continuity requires effectively managing and harmonizing internal and external requirements for user access entitlements to individual datasets. This approach involves considering factors such as privacy, confidentiality, integrity, availability, business purpose, and life cycle risks.
- Minimizing data security policy gaps and inconsistencies demands processes for creating and orchestrating policies across multiple independent data security and identity and access management (IAM) products.
- No single product comprehensively mitigates the business risk associated with data, emphasizing the need for the centralized creation and coordination of data security functionalities.
- An increasing number of privacy regulations have introduced risks associated with data residency and sovereignty requirements. Leveraging the DSG process, including an effective privacy impact assessment (PIA), is essential to mitigate these risks.

Obstacles

- Business stakeholders have fragmented responsibilities for managing data and its associated risks, which creates challenges for implementing consistent data security policies.
- Robust DSG becomes more complex as organizations deploy multiple data security, IAM, and application-native security products that act independently on unstructured or structured data, or leverage proprietary data classification.
- Dynamic changes to data storage and processing across on-premises and multicloud services can hinder consistent enforcement of data security policies.
- DSG requires security teams to orchestrate data security policies throughout the portfolio of available security controls. This obstacle is compounded by the requirement for regular data risk assessments (DRAs) to uncover gaps and inconsistencies from increased business risks or to identify unmet policy or product deployment needs.
- Since they accomplish similar objectives, PIAs are incorrectly viewed as a replacement for a DRA.

User Recommendations

- Use DSG to analyze risks and their business impacts. Practices like infonomics and data monetization can be used to correlate the financial impact of data risks on business outcomes.
- Establish a balanced and unified data security posture using findings from risk and impact assessments and business opportunities presented by DSG stakeholders. At a minimum, include the CDAO and the CISO.
- Create, manage, and deploy consistent data security policies throughout your portfolio of datasets, as well as across data security, IAM, and application management products that interact with each dataset.
- Align the work of data stewards throughout the organization in support of all governance work, including data security, quality, privacy, and AI ethics.
- Gather feedback (e.g., biannual surveys, focus groups) from business unit and cross-functional leaders to understand if data security governance initiatives are making it harder for employees to get work done. If a data security steering committee (DSSC) exists, use it to surface and interpret these insights. Apply the findings to better align governance with employee workflows and business objectives.

Gartner Recommended Reading

[Use Agile, Adaptive, AI-ready \(3A\) Data Security Governance to Secure Shadow AI](#)

[Cybersecurity Trend: AI Democratization Drives Collaborative Data Security Governance](#)

[4 Steps to Accelerate Adoption of Data Security Governance](#)

Data Clean Rooms

Analysis By: Eric Schmitt

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Adolescent

Definition:

A data clean room is a secure, isolated platform that links anonymized marketing and advertising data from multiple parties. For advertisers, data clean rooms are distinguished from other data-sharing methods by the inclusion of detailed advertising impression data, with privacy-safe restrictions on outputting user-level results.

Why This Is Important

Data clean rooms offer a privacy upgrade to advertising capabilities by allowing marketers to rebalance privacy with new data access and utilization policies that are fit for a consent-driven advertising world.

Business Impact

At scale, the opportunity to access ad impression and interaction data presents digital marketing leaders with the potential to get more from advertising spend and to enable insights and AI use cases. Examples include data governance, audience segmentation for advertising activation, campaign measurement, data enrichment and forecasting.

Drivers

- **Demand for data access:** Shifts in end-user behavior across apps and platforms, and the siloed nature of data collection create an opportunity for new approaches to data-driven advertising and marketing.
- **Data appending and enrichment:** Clean rooms can enrich first-party customer and prospect profiles with additional data attributes or elements. These attributes may be sourced from the clean-room provider or second- and third-party data sources, or developed/derived via statistical modeling and processing.
- **Advertising execution enablement:** Clean rooms offer a pathway for addressable and look-alike targeting in the data-rich retailer and publisher environments, without the need for those parties to expose any data.
- **Ad measurement:** On an anonymized basis, clean rooms can measure how many and which people were contacted, and how many times over the duration of the campaign. This sets the stage for target audience penetration analysis, unique unduplicated reach, and refined campaign activation.
- **Campaign performance measurement:** Clean rooms enable marketers to link ad impressions data to sales and engagement data. This can feed campaign performance analysis and attribution, and measurement reports and dashboards. Clean rooms can also inform media planning, financial and management reporting, and business forecasting.
- **Projections and forecasts:** The ability to stack data for ongoing, time-series analysis has strong appeal for many marketing leaders. This can create data views for use in advertising, sales projections and forecasts.

Obstacles

- **Lack of standards:** The absence of data standards for audience targeting and ad measurement force marketers to incur greater cost and complexity related to multiple, platform-specific implementations. The pool of big data processing and analytics talent available to advertisers is limited. Talent shortages combined with procurement inexperience can add costs and delays.
- **Privacy risks:** As with any sensitive data processing, breaches or circumvention of privacy controls could result in reidentification, and therefore liability.
- **Disruption distractions:** Although privacy-related data gaps may boost clean rooms in the long run, in the near term, many advertisers are likely to focus on tactical operational adjustments to media plans, and not speculative and resource-intensive data infrastructure investments.
- **Cost, resourcing and relationship requirements:** Complex database investments like clean rooms incur substantial costs related to technology, staff or services partner time to deploy, configure and maintain. In addition, clean rooms typically require business and legal agreements that span two or more parties.

User Recommendations

- Define clean-room goals. The menu includes reach/frequency measurement, campaign performance analysis, projections and forecasts, governance, and campaign segmentation.
- Explore the landscape of clean-room providers. Brands with heavy-paid media investments in Google or Amazon should start by looking at these vendors' clean-room offerings. Independent and agency solutions are a better fit for cross-media applications.
- Embrace agile principles during clean-room development. Staging, ingesting and linking data — especially from new, high-volume sources — often takes longer than planned and requires multiple iterations.
- Apply good governance and privacy by design. Despite technical safeguards such as aggregate reporting outputs with minimum cell sizes, risks like reidentification exist. Mitigate risks with documented use cases and detailed data flow diagrams that cover processing and storage, and by building up internal data skills like SQL and BigQuery.

Sample Vendors

Omnicom (Acxiom); Amazon; AppsFlyer; Decentriq; Dentsu (Merkle); Google; LiveRamp; Publicis (Epsilon); Snowflake; WPP (InfoSum)

Gartner Recommended Reading

[Navigate the Marketing Data Collaboration Landscape](#)

[Don't Let Cookies' Reprieve Stop Your Ad Strategy Evolution](#)

[The CMO's Key to First-Party Data's Triple Lock](#)

[Quick Answer: Estimate the Costs of a First-Party Data Strategy](#)

[5 Steps for CMOs Exploring First-Party Data for Advertising and Targeting](#)

Synthetic Data

Analysis By: Arun Chandrasekaran, Alys Woodward, Anthony Mullen

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Definition:

Synthetic data is a class of data that is artificially generated rather than obtained from direct observations of the real world. Synthetic data is used as a proxy for real data in a wide variety of use cases, including data anonymization, AI and machine learning development, data sharing, scenario exploration and generative AI model pretraining and fine-tuning.

Why This Is Important

Obtaining and labeling real-world data for AI development is a time-consuming and expensive task. Synthetic data enables rapid data generation, cost-effectively and without personally identifiable information (PII) or protected health information (PHI), making it a valuable technology for privacy preservation. The rise of frontier AI models has highlighted synthetic data as a cost-effective means to build scalable models.

Business Impact

- Avoids using PII when training AI models via synthetic variations of original data or the synthetic replacement of parts of data.
- Reduces costs and saves time in AI development.
- Improves AI performance as more “fit-for-purpose” training data leads to better outcomes.
- Enables organizations to pursue new use cases through intelligent simulation, for which minimal real data is available.
- Addresses fairness issues, such as bias and toxicity.

Drivers

- In regulated industries, such as healthcare and finance, synthetic tabular data is used to preserve privacy in AI training data.
- To meet the increasing demand for synthetic data for natural language automation training, especially for chatbots and speech applications, vendors are introducing new offerings to market, often to train domain generative AI (GenAI) models.
- Synthetic data applications have expanded beyond automotive and computer vision use cases to include platform evaluation and the development of test data.
- Transformer and diffusion architectures, the architectural foundations for GenAI, are enabling synthetic data generation at quality and precision levels not seen before. AI simulation techniques are improving synthetic data quality by better recreating real-world representations.
- There is scope for expansion to other data types. While tabular, image, video, text and speech applications are common, R&D labs are expanding the concept of synthetic data to graphs and multimodal AI. Synthetically generated graphs will resemble, but not overlap, the original. As organizations begin to use graph technology more, this method is expected to mature and drive adoption.
- As data providers for training frontier AI models raise their data access costs, synthetic data will gain traction as an economic alternative.

Obstacles

- Most vendors offering commercial synthetic data solutions directly to enterprises haven't been able to build successful business models.
- While synthetic data reduces privacy risks, some industries (e.g., healthcare, finance) still face regulatory uncertainty about whether it can fully replace or augment real data for compliance purposes.
- Validating the accuracy of synthetic data is difficult. It can be challenging to confirm whether a synthetic dataset accurately captures the underlying real-world environment.
- Buyers are still confused about when and how to use the technology due to a lack of skills.
- Enterprises have legacy data pipelines, and integrating synthetic data into existing data lakes, analytics systems and AI workflows often requires additional effort, tooling and infrastructure.
- There may be a level of user skepticism because data may be perceived to be "inferior" or "fake."

User Recommendations

- Identify areas in your organization where data is missing, incomplete or expensive to obtain and is thus currently blocking AI initiatives.
- Establish clear policies on synthetic data life cycle management, storage and access control.
- Work with legal and compliance teams to ensure synthetic data aligns with industry regulations, such as General Data Protection Regulation, Health Insurance Portability and Accountability Act and the California Consumer Privacy Act.
- Educate internal stakeholders through training programs on the benefits and limitations of synthetic data. Institute guardrails to mitigate challenges such as user skepticism and inadequate data validation.
- Measure and communicate the business value as well as the success and setbacks of synthetic data initiatives.

Sample Vendors

Anyverse; Betterdata; Bifrost; MOSTLY AI; NVIDIA; Parallel Domain; Rendered.ai; Toloka AI; Tonic.ai; YData

Gartner Recommended Reading

[Emerging Tech: Top-Funded Startups in Synthetic Data](#)

[Executive Briefing on Emerging Technology: Synthetic Data](#)

[Use Synthetic Data to Improve Software Quality](#)

Data Breach Response

Analysis By: Nader Henein, Bernard Woo

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Data breach response and disclosure refer to the essential steps taken to evaluate a data breach and, if necessary, notify regulatory authorities as well as affected individuals when personal data is compromised. Disclosure is required under comprehensive laws such as the EU's General Data Protection Regulation (GDPR), Australia's Notifiable Data Breaches (NDB) scheme, as well as subject- or region-specific laws like individual U.S. state breach notification statutes.

Why This Is Important

Appropriate management of a reportable breach can:

- Substantially reduce regulatory fines.
- Reduce financial impact.
- Strengthen trust with affected consumers.

Prompt, clear communication demonstrates that the organization is proactively taking ownership of the situation. However, delayed communication, limited transparency, and overly legalistic communications often elicit regulatory investigations, resulting in reputational damage and customer loss.

Business Impact

Data breach response can have a critical impact on an organization's operational cyber resilience. Breaches often create significant chaos as key executive team members pivot from preexisting priorities to address the reputational, regulatory and likely financial impacts of the breach. Further, more recent privacy legislation in Africa and the Middle East often imposes statutory sentences on company directors for inadequate or negligent handling of personal data.

Drivers

- The accelerated development and adoption of AI capabilities means that the likelihood of sensitive data being ingested by language models is near absolute.
- Modern privacy regulations have raised the bar for data breach notification. When personal data is impacted, disclosure to a supervisory authority within days of discovery is often required.
- In the U.S., the regulatory environment for breach notification is fragmented and overlapping. There are a variety of federal mandates, such as SEC and HIPAA disclosure rules, which are constantly undergoing modernization. Each state also has its own breach notification laws, which are regularly amended and updated to meet consumer expectations. For example, Texas and New York have amended their laws in the past few years. In California, the interplay between breach disclosure and the California Consumer Privacy Act (CCPA) has brought the risk of private right of action to the forefront.
- Regulatory evolution illustrates the need for organizational commitment and a reallocation of resources.
- Organizations must constantly align the technical and operational elements of incident response (IR) with new legal and regulatory requirements.
- Many organizations still need to elevate their capacity to disclose a data breach to regulators and potentially affected individuals in an accelerated time frame.

Obstacles

- Establishing and testing a data breach program is an expense without an immediate return.
- Data breach service retainers are commonly available but due to the variability of the breaches and the data involved, each breach scenario is unique.
- Even with a strong program, the time to discover an incident is measured in months – although it is improving over time.
- Tensions between general counsel and chief information security officers (CISOs) could negatively impact the organization's breach handling capacity.
- Regulatory breaches requiring disclosure of the misuse of AI are evolving with the introduction of AI governance laws which number in the hundreds.
- The lack of vendor transparency regarding the use of client data for model training is an omnipresent risk for all organizations.
- Data breach response requires a rare combination of technical acumen, forensic capabilities and legal know-how.

User Recommendations

- Partner with your general counsel from the beginning; reportable breaches carry liability, which is commonly the purview of your legal department.
- Assess whether an incident will trigger regulatory actions such as those defined under reportable privacy or AI violations.
- Record and maintain the details for all incidents (not just violations), as some jurisdictions require stringent record-keeping.
- Assemble diverse teams as data breach response requires multidisciplinary functions, documented policies and procedures, defined RACI roles.
- Conduct regular simulated drills, such as tabletop exercises and update your procedures accordingly.
- Introduce obligations in vendor contracts that mandate transparency on data use, subprocessors, and model training beyond the enterprise tenant.
- Conduct periodic gap assessments against evolving breach notification laws.

Sample Vendors

BigID; BreachRx; Canopy Software; OneTrust; RadarFirst; Resolver; Securiti

Gartner Recommended Reading

[Toolkit: Cybersecurity Incident Response Plan](#)

[Toolkit: Security Incident Response Roundtable Scenario for Privacy](#)

Subject Rights Requests

Analysis By: Stefan Dumitrescu

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Subject rights requests (SRRs) are a set of entitlements by which individuals exercise their legal rights over their personal data, such as requesting access, correction, deletion or receiving a copy of their information. Organizations handling personal data must address SRRs within a defined time frame set by applicable privacy regulations. Doing so at scale requires varying levels of automation as part of a modern privacy management program.

Why This Is Important

More than 75% of the world's population now has access to fundamental privacy rights under statutory regimes. For cybersecurity leaders, subject rights management is not only a compliance requirement but also a prerequisite for building customer trust. Vendors in the SRR market provide partial to full automation of the fulfillment life cycle, from initial capture request to final response, either as a stand-alone capability or as part of the broader privacy user-experience ecosystem.

Business Impact

The impact of poor or delayed responses to SRRs is threefold:

- Regulators can impose fines for failure to comply or for delayed compliance, and rulings often mandate prompt execution of requests.

- Trust between the organization and its customers erodes, which may lead individuals to request deletion or limit processing of their data.
- Longer response times damage customer sentiment, especially when it takes 90 days to receive a copy of their data.

Drivers

- The strategic importance of security and responsible data handling continues to grow as meaningful customer control over personal data becomes a core determinant of trust. Without that control, confidence in the organization's broader governance and reliability erodes.
- Privacy is now seen as a conviction-based buying factor, with customers making decisions based on how organizations process personal data, accelerating innovation in data-protection technologies.
- Large multinationals embedding privacy principles across products, AI systems and supply chains signal a structural shift toward transparency, accountability and ethical data governance. Without trust, customers limit data sharing and consent, and without both, data-driven strategies collapse.
- High-profile breaches and AI misuse have eroded fundamental customer trust, increasing demand for stronger security controls and demonstrable accountability. Regulators are showing decreasing tolerance for repeated failures and are moving more quickly from investigation to enforcement and fines.
- As organizations rely on personal data to drive personalization and AI, a strong privacy UX across SRRs, consent and preference management and tracker consent becomes a key driver of trust, compliance and sustainable data access.

Obstacles

- Given the unpredictable demand for SRRs, cybersecurity leaders may delay adopting scalable infrastructure and automation.
- Locating and retrieving data across multiple systems and formats becomes a technical challenge for organizations with decentralized storage, including cloud-based environments, SaaS sprawl, AI systems and legacy databases.
- Automation is essential for addressing SRRs at scale, but it requires moving away from manual methods that no longer scale and recognizing that no technology can resolve foundational gaps in data visibility, ownership and governance. Maintaining sufficient capacity to respond within regulatory time frames and meet customer expectations remains critical.
- The growth of unstructured content and AI-generated outputs increases the complexity of locating, reviewing and validating data.
- Budget pressure, unclear ownership of privacy responsibilities and shortages of privacy-engineering expertise collectively hinder program maturity.

User Recommendations

- Ensure request fulfillment follows a repeatable, AI-enabled, cost-aligned process capable of scaling across jurisdictions.
- Prioritize customer loyalty and support regulatory compliance by delivering a transparent, user-centric experience for SRR fulfillment.
- Make SRRs an integral part of the privacy UX ecosystem by delivering transparent, user-centric, end-to-end experiences that reinforce trust and informed decision making.
- Implement centralized discovery and mapping across cloud, SaaS, legacy and AI-generated data sources.
- Establish, monitor and report foundational SRR metrics by tracking response times, cost per request and throughput capacity against legal timelines.
- Propose an automation level for SRR fulfillment aligned with projected spending, and identify areas within the response chain that would benefit from improvement.
- Embed a culture of privacy consciousness across all customer-facing functions to continuously optimize workflows and reduce operational risk.

Sample Vendors

DataGrail; IntraEdge (Truyo); MineOS; Netzima (Icaria); OneTrust; Osano; Securit; Transcend

Gartner Recommended Reading

[Best Practices for Automating Subject Rights Requests](#)

[State of Privacy: Global Overview of Regulatory Evolution by Region](#)

Consent and Preference Management

Analysis By: Tia Zervas

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Consent and preference management platforms consolidate end users' choices on how their personal data should be handled. The intent is to enhance transparency by extending control to users to determine what data about them is processed, by whom, and for what purpose. These choices are synchronized across systems — old and new, with on-premises and cloud-based repositories. Organizations can target engaged customers and respect their choices with minimum manual overhead.

Why This Is Important

The EU's General Data Protection Regulation, the California Consumer Privacy Act, and Brazil's General Data Protection Law are the standard for organizations to follow to adhere to privacy protections for personal data collected. Consent and preference management (CPM) platforms empower organizations to comply with new laws, preserve and extend essential capabilities, and demonstrate to customers and stakeholders that they care about privacy.

Business Impact

- CPM platforms help address the technical and operational challenges of protecting organizations from compliance violations, while maintaining the ability to use customer data for business purposes.
- CPM platforms can show end users that organizations value their privacy. Compliance is essential for meeting regulatory requirements and is a fundamental expectation of end users. A lack of compliance can be more detrimental than the advantages gained from adhering to it.

Drivers

- **Regulatory oversight:** Regulators scrutinize organizations' use of personal data for marketing and advertising. CPM platforms serve as the hub in marketing and advertising technology stacks to process the appropriate consents and maintain audit trails for compliance.
- **Evolving laws and variations in legislation:** The proliferation of consumer privacy laws across countries and regions makes tracking and complying with local requirements a complex but essential task. CPM platforms address specific requirements, such as auditing websites for compliance, enforcing consent choices, and supporting data subject requests.
- **Safety regulations for minors:** New child safety regulations now require both age verification and explicit parental consent before processing minors' data. CPM platforms are adapting to facilitate dual-verification workflows to support these heightened requirements.
- **Demand for personalization:** As organizations prioritize one-to-one personalized experiences, challenges arise from a lack of a unified data strategy. The absence of a unified customer data strategy presents significant challenges to achieving this at scale. CPM platforms address these hurdles by enhancing preference centers, enabling progressive consent, and supporting advanced profiling capabilities.
- **Disparate data across technology:** Managing consent and preferences across multiple platforms is resource-intensive, especially with overlapping CPM solutions. Some CPM platforms address this challenge by serving as the primary source of truth of digital users' consent and preference selections. This function is critical in jurisdictions like the U.S., where implicit consent is still prevalent.
- **Consumer expectations:** Consumers are increasingly aware of their data rights and demand transparency, control over, and expect a consistent experience across channels. Poorly designed consent flow banners and dialogues degrade the user experience, underscoring the need for improved design enabled by certain CPM platforms.

Obstacles

- **Evolving global laws and best practices:** Organizations must adapt to region- and country-specific privacy and compliance legislations. CPM platforms tend to oversell their ability to simplify managing consent options, often downplaying the complexity of managing an organization's internal and external databases.
- **Lack of UX design support:** Forcing too many privacy choices on consumers degrades UX and leads to high opt-out and abandonment rates. Yet, having too few choices limits the ability to tailor experiences. To strike the right balance requires cross-functional, collaborative activities throughout the organization.
- **Complex technology architectures and the rise of AI:** Accelerated digital transformation efforts have propelled organizations to rethink how technology solutions work together and how data flows throughout the ecosystem. AI adoption has further complicated the subject of consent since removing data from a model is practically unfeasible if consent is revoked. Adopters need to factor in the number and type of connections – both native and customized (such as APIs and extraction, transformation, and loading) – that are needed to effectively use a CPM platform.

User Recommendations

- Prioritize consent management policies and initiatives as critical for all functions. Establish consent as the foreground to maintain a consistent privacy UX initiative throughout the enterprise.
- Avoid “dark patterns” or deceptive language for consent dialogues that attempt to influence users’ choices.
- Use a “telescoping” approach to disclosures and preference dialogues that allow users to go as deep as they choose into specific details. Offer consistent, easy access to preference settings that users can view and change on demand to ensure a privacy-by-default approach.
- Compare and assess CPM offerings against your organization's highest-priority privacy protection and integration requirements and internal costs.
- Build your own CPM platform wherein the current CPM market cannot effectively connect and integrate with legacy internal tools.
- Take a modular approach to adoption and avoid excessively broad project scopes. Anticipate sufficient time to resolve unforeseen complications in these projects.

Sample Vendors

BigID; Didomi; Ketch; OneTrust; Osano; PossibleNOW; Syrenis; Transcend.io

Gartner Recommended Reading

[Market Guide for Consent and Preference Management](#)

Digital Ethics

Analysis By: Bart Willemsen, Frank Buytendijk

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Digital ethics consists of combined systems of values and moral principles to guide organizational approaches regarding the architecture, intent, and consequence of facilitated and/or executed interactions between people, organizations, and things. It applies to areas such as AI, data and analytics, and online/social media presence and behavior.

Why This Is Important

Digital ethics concerns most, if not all stakeholders in an organization, especially around privacy, bias controls, and censorship. The voice of society gets louder for responsible positioning, for example on the use of AI. The hype around (Gen)AI agents further increased public concerns about transparency, disinformation, content lineage, and intellectual property. Organizations are trying to mitigate ethical risks, while several governments simultaneously regulate AI technology.

Business Impact

Digital ethics strengthens an organization's positive influence and reputation among customers, employees, partners, and society. Areas of business impact include innovation, product development, customer engagement, strategy, and go-to-market. Intention is key: If ethics is merely addressed to "polish" reputation, it becomes disingenuous. The goal to be an ethical organization serves all parties, including society at large, and leads to enhanced trust, loyalty, and performance.

Drivers

- The media frequently feature stories about misinformation yet they are not immune to spreading it themselves. Coverage of how AI affects, or sometimes even is connected to ending, lives is increasing. Board members and other executives share concerns about the unintended consequences of innovative technology use.
- Ethics has often been an afterthought and are thus insufficiently effective. Especially with the amplification of effects at scale by AI technology usage, the time to respond in hindsight, to failure, is no more. The ethical debate has to take place both before and during technology development and adoption. AI ethics aims to establish responsible use of AI and to harness AI's growing powers.
- Ongoing hype regarding agentic technology increases the urgency of ethical and legal issues surrounding the veracity and ownership of data. The potential impact of an inaccurate, misleading, or insensitive output fuels ethical concerns.
- Government commissions and industry consortia have developed guidelines for mitigating the risks and managing safe and ethical AI use (see, for example, the EU's [AI Act](#) and the U.N.'s [resolution on artificial intelligence](#)).
- Organizations increasingly declare their digital and AI ethics principles, frameworks, and guidelines, and many are in the process of going from declaration to execution. However, several seem to be simply "ethics washing" their activities or diverting from previous moves toward ethics goals.
- Around the world, education institutions have added digital ethics courses and launched programs to address the ethical, policy, and legal challenges new digital technologies pose.
- Digital ethics is expanding to address concerns about intellectual property, working conditions, and a growing environmental footprint, particularly in terms of energy and water consumption spurred by very large AI foundation models.

Obstacles

- Organizations often struggle to operationalize digital ethics because of its ambiguous, pluralist, and contextual nature, with positive effects often being only indirectly tangible.
- Diverging ethical views and expectations among regions and cultures challenge organizations that operate internationally. Consensus among internal and external stakeholders, such as customers, can be elusive.
- Technology can support digital ethics to some extent. For example, it can help to protect privacy, improve transparency or mitigate bias. However, it is far from perfect and needs to mature further. Many of the tools designed to enforce ethical principles are still experimental, prone to errors and inconsistencies, and often lack the robustness needed to function reliably at scale.
- Digital ethics is too often reactive, mistakenly interpreted as only compliance, reduced to a checklist, confined to technical support for privacy protection, and/or viewed only as explainable AI.

User Recommendations

- Elevate the conversation by focusing on digital ethics as a source of societal and business value and an approach to reducing uncertainty and stagnation in innovation, rather than focusing only on compliance and risk.
- Turn awareness into action by identifying specific digital ethics issues, as well as existing or emerging regulations or guidelines from governments and industry associations.
- Discuss ethical dilemmas from diverse points of moral reasoning. Anticipate and account for ethical consequences.
- Ensure digital ethics is leading and not following the adoption of new, transformative technology, such as AI. Address digital ethics upfront, by design, to create methods that identify and resolve ethical dilemmas as early as possible.
- Organize training in ethics and run workshops to create ethical awareness within all digital and AI initiatives. These moves should include governance processes for clear accountability and building trust that AI and other technology are applied responsibly.

Gartner Recommended Reading

[How to Manage Digital Ethical Dilemmas](#)

[Tool: Use These Real-World Cases to Master Digital Ethics](#)

[Use These 6 AI Ethics Principles to Develop and Run AI](#)

[Key Trust and Ethical Implications of Generative AI](#)

Climbing the Slope

Data Discovery

Analysis By: Apurva Singh, Michael Hoeck

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Mature mainstream

Definition:

Data discovery is a foundational technology that identifies and analyzes structured and unstructured enterprise data. Powered by AI, pattern recognition and machine learning (ML), it automates data classification, categorization and sensitivity labeling. Acting as a core intelligence engine, it embeds into broader security, privacy, compliance and governance platforms, feeding essential insights to downstream data protection and storage management controls.

Why This Is Important

Data discovery helps mitigate data sprawl across hybrid, cloud and SaaS environments. By establishing visibility into dispersed data, it enables security and compliance teams to identify and protect sensitive, regulated information. Security teams use discovery insights to mitigate data-related risks and enforce security controls while infrastructure teams identify redundant, outdated or trivial (ROT) data to optimize storage, and data teams curate high-quality inputs for AI pipelines.

Business Impact

- Identify sensitive and regulated data to enforce compliance, privacy and security policies.
- Drive data life cycle management by linking retention, deletion and minimization policies to discovery insights.
- Identify business risk by surfacing ROT data, granular sensitive-data exposure, and content and location context for downstream governance and security actions.
- Improve AI initiatives by identifying and flagging sensitive or high-value datasets for further security and governance action.

Drivers

- Increasing data volumes and enterprises' use of disparate applications and storage platforms are driving organizations to expand capabilities to mitigate data vulnerabilities. Unmanaged, these vulnerabilities lead to severe business consequences, including data breaches and exfiltration, operational interruption, reputational damage, intellectual property loss, and regulatory fines, alongside the mounting financial burden of unchecked data sprawl.
- There is a need to reduce the blast radius of a cyberattack in terms of its access to sensitive information through data discovery coupled with data classification and categorization, security controls, defensible deletion and minimization efforts.
- Organizations need data discovery outputs to inform downstream access, governance and protection decisions based on data sensitivity and categorization.
- Growing adoption of GenAI and AI governance increases the need to identify sensitive data, especially across unstructured repositories.
- Data discovery is increasingly critical in preparing enterprise data for use in AI pipelines. By identifying, contextualizing and curating high-quality input data across sources, it fuels training datasets, large and small language models (LLMs/SLMs), and broader ML initiatives.
- Expanding global compliance requirements (e.g., General Data Protection Regulation, California Privacy Rights Act, Health Insurance Portability and Accounting Act and financial services compliance) require comprehensive and accurate identification of sensitive and regulated data.

Obstacles

- For their broad set of use cases, capabilities and benefits, funding for data discovery technologies may require a collaborative effort across departments including security, privacy, compliance, legal, data and IT teams.
- There is frequently no clear owner accountable for data discovery outcomes. While IT or security teams typically operate the tools, they may not have the authority to act on the findings. This disconnect between visibility and downstream action limits risk reduction and creates accountability gaps across data domains.
- Organizations often struggle to translate data discovery insights into prioritized deletion, remediation or security actions when policies, owners and execution tools are fragmented.

- Comprehensive discovery can be constrained by the scalability of legacy platforms and the integration effort required to connect siloed or hidden enterprise data environments.
- Advanced AI/ML discovery techniques may lack transparency, reducing stakeholder trust and adoption.

User Recommendations

- Adopt automated data discovery tools capable of enterprisewide scanning to surface data risks and subsequent security, privacy, compliance and IT governance activities.
- Use data-discovery-generated inventories and risk profiles to strengthen downstream data classification, data loss prevention, data security posture management and governance controls.
- Improve data life cycle management principles by using discovery insights to prioritize retention, deletion and minimization policies and actions.
- Ensure data discovery tools are capable of detecting sensitive content across diverse structured and unstructured data types.
- Develop clear data visualization maps using discovery insights to communicate data risks, ownership and value to relevant stakeholders.

Sample Vendors

AvePoint; Bedrock Data; BigID; Concentric AI; Cyera; EncompaaS; Netwrix; RecordPoint; Sentra; Varonis

Gartner Recommended Reading

[Buyers' Guide for Privacy Management Tools](#)

[4 Key Practices to Keep Sensitive Data Safe From GenAI](#)

[2026 Strategic Roadmap for Storage](#)

Decentralized Identity

Analysis By: Michael Kelley, Akif Khan, Arthur Mickoleit

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Definition:

Decentralized identity (DCI) democratizes digital identity by decentralizing both the storage and the use of identity data. The primary benefits of DCI are privacy, anonymity, and user autonomy. DCI components include an identity trust fabric, a digital wallet, which is tied to an entity (user), verifiable credentials (VCs), which represent identity attributes used to prove identity claims, and interfaces for issuers and verifiers of those credentials.

Why This Is Important

DCI can establish trust in the identity of a person, which can help solve problems related to identity verification, account takeovers, privacy, and security. It is driven by a decentralized approach to identity data in the form of VCs, which represent proofs for claims about identity attributes, like employment status, citizenship, or authorizations for access. DCI, compared with existing approaches, represents magnitudes of improvement in terms of efficiency, cost, privacy, and assurance.

Business Impact

Users gain control of their identities and data, and service providers (SPs) gain higher trust, speed to value, and lower exposure to risk. SPs collect huge amounts of identity data about users for every interaction to increase assurance to an acceptable level. DCI can provide trust, security, privacy, convenience, and portability of identity data for end users without needing centralized data, thereby reducing the risks of data breaches, account takeovers, and privacy compliance violations.

Drivers

- **Vendor investments in DCI:** In addition to observing influential vendors (such as IBM, Microsoft, and Ping Identity) making significant investments in DCI, Gartner has tracked more than 80 startups or established vendors of DCI technologies and DCI components.
- **Government activity:** Public sectors are increasingly shaping DCI trends. The EU has ratified [eIDAS 2.0](#), which mandates the provisioning of identity wallets (used in DCI) to citizens by the end of 2026. Other national, regional, and local authorities are exploring and investing in DCI use cases across public and private sectors. Examples include Finland, Buenos Aires in Argentina, and the Basque Country region in Spain.
- **Regulations:** Countries continue to formalize requirements for user privacy, establishing regulations for collecting and securing large amounts of user data. DCI complies with privacy regulations through decentralizing user data. In addition, basic use cases for know your customer (KYC) and anti-money-laundering are being developed for DCI use cases.
- **Client and overall market interest in DCI:** Interest is increasing due to the momentum of companies beginning to use DCI approaches to enable new digital business opportunities while maintaining client privacy.
- **Standards:** Standards are maturing, led by entities such as the World Wide Web Consortium (W3C), Trust Over IP, the OpenWallet Foundation, and OpenID for Verifiable Credentials (OID4VC) to create a consistent approach to DCI.
- **User experience (UX):** Asking users to repeatedly go through identity verification (IDV) and affirmation processes for every new online interaction with an SP is a broken model. Significant friction can be removed from UX if users could verify once and then assert their identity to each new SP, as needed, using an identity wallet with full control over their identity data. DCI can make high-trust IDV available to a larger number of SPs without having to invest in discrete IDV tools.

Obstacles

- **Authority of issuers:** There is no global standard for ensuring that an organization has the authority to issue a VC (such as only an accredited facility issuing educational credentials).
- **Infrastructure standardization:** A uniform standard for adding issuers of VCs, as well as verifiers, for any DCI network does not yet exist.
- **Interoperability:** Most development is taking place in pockets, and standards continue to mature slowly.
- **Technical challenge:** Concerns exist about performance, scalability, maturity, and wallet standards.
- **Regulations:** More work is required for how verifiable claims can be used in regulated use cases, such as driver's licenses and other government records, or KYC and AML.
- **Security:** Identity wallets must enforce strong authentication controls to ensure that only the person whose identity attributes are held within the wallet can make identity assertions. Wallet recovery processes must also be secure.
- **Market competition:** Commercial wallets like Apple and Google.

User Recommendations

- Explore use cases by identifying use cases that are expensive, complex, and time-consuming in the real world, which will benefit from a VC approach.
- Follow government progress on DCI for citizen ID for “bootstrap” opportunities, as well as the various open-source initiatives (like those from Walt.ID and SpruceID).
- Leverage platforms that were developed to help clients build and test use cases for DCI.
- Track new developments in emerging standards, like Trust DID Web (did:tdw) and Key Event Receipt Infrastructure (KERI), which may help address some of the technical challenges for decentralized trust infrastructure and registries.
- Observe the development of organization credentials taking place with verifiable legal entity identifiers from the Global Legal Entity Identifier Foundation for the potential of establishing authoritative VC issuers.
- Be prepared for early disruption in the DCI market. Gartner expects some chaos with mergers and acquisitions, as well as vendors exiting the business.

Sample Vendors

1Kosmos; Avast; IBM; IdRamp; Interac; Lissi; Microsoft; Ping Identity; Scytáles

Gartner Recommended Reading

[Enhance Trust With Decentralized Identity and Verifiable Credentials](#)

[Begin Portable Digital Identity Evaluation With Use Cases](#)

[Top Trend in Government: Digital Identity Ecosystems](#)

[Market Guide for Decentralized Identity](#)

Format-Preserving Encryption

Analysis By: Joerg Fritsch, Sarah Almond

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Format-preserving encryption (FPE) protects data at rest and in use and when accessed through applications while maintaining the original data length and format. It's primarily used to protect fields in databases and selected semistructured and unstructured data types. FPE is an important anonymization technique to support data protection, compliance and privacy requirements. It can also reduce the risk of data residency, hacking or insider threats by controlling access to data.

Why This Is Important

FPE can be used to protect data at the point of ingestion, storage in a database or access through data pipelines. It is deployed to protect sensitive data stored or processed across a variety of databases. It is a robust access control, and, when applied, will protect data wherever it resides or is accessed.

Business Impact

FPE is widely accepted to protect data without having to extensively modify databases or applications. It provides a strong method to prevent unauthorized access to data on-premises and in cloud service platforms (CSPs), and to mitigate security threats. FPE can also help organizations meet data protection regulations, privacy regulations and data residency requirements to protect personal, health, credit card and financial data, and to adhere to data breach disclosure regulations.

Drivers

- There are widely accepted standards from National Institute of Standards and Technology Special Publication (NIST SP) 800-38 and 800-38G Revision 1 that specify and approve the Format-Preserving Encryption, Mode 1 (FF1) for FPE. This is currently the only NIST-approved mode; modes FF2 and FF3 have suffered security flaws.
- Adoption is increasing due to the fast-growing need to provide data protection and data residency restrictions and the increasing number of privacy laws across the globe.
- Organizations increasingly want to analyze data while keeping it anonymized. But some staff will need access to clear text, which is driving the need for FPE to provide business-friendly access controls that leverage augmented data cataloging.
- The ability to integrate FPE with data masking and test data management as part of a data security platform is also increasing its adoption to support the requirements for data residency, compliance and privacy, and test and development.

Obstacles

- FPE maps clear text data to encrypted data with the same dataset, making the **data from small domain sizes** (for example, a dataset with one million U.S. SSNs or less) vulnerable to offline brute force attacks and dictionary attacks.
- Native database functionality — for example, indexing or ordering — is either lost or requires workarounds that can leak statistical data. Without these workarounds, a database will need to decrypt a whole column in order to sort or filter ranges.
- FPE presents challenges for encryption key management and encryption key rotation because reencrypting the data is much more disruptive than with traditional encryption due to potential collisions (of so called “tweaks”) and the need to reencrypt the whole dataset. Rewrapping a data encryption key (DEK) is not possible for FPE data.

User Recommendations

- Review how FPE interacts with applications. Establish whether it can be used to control which user identities are allowed to see the data in clear text. Try to leverage data catalogs and metadata management to support adaptive business access requirements.
- Evaluate the impact of FPE on the performance and functionality of applications accessing the database.
- Identify any impacts on application and database functionality, such as search and sorting.
- Evaluate encryption key versioning and lazy reencryption where only new data is encrypted with the “rotated” encryption keys while older records remain encrypted with the previous key.

Sample Vendors

Baffle; Capital One; comfote; IronCore Labs; OpenText; PKWARE; Prime Factors; Protegrity; SecuPi; Thales

Gartner Recommended Reading

[Market Guide for Data Security Platforms](#)

[How to Improve Data Security With Cloud Key Management](#)

[Postquantum Cryptography: The Time to Prepare Is Now!](#)

[Create a Data Risk Assessment That Drives Business Decisions](#)

Age Assurance and Parental Consent Tracking

Analysis By: Bernard Woo

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Early mainstream

Definition:

Age assurance uses digital, biometric, and validation techniques to verify users' age group participation. It is used to prevent minors from accessing age-restricted content and digital social media platforms, and ensure compliance with many recent child-protection regulations. "Age assurance" is the umbrella term for age estimation and age verification. Many regulations include parental consent tracking, a requirement often overlooked in identity assurance solutions.

Why This Is Important

Driven by the U.S. Children's Online Privacy and Protection Act (COPPA), the UK's Online Safety Act, Australia's Online Safety Amendment Act, and similar laws, a new wave of age assurance regulations is emerging. Lawmakers are defining what content is in scope, determining valid age checks and enforcement and parental consent use cases, and aiming to limit minors' access to age-restricted content and services to protect their data and restrict targeted ads.

Business Impact

Online vendors of age-restricted content need fast, reliable age verification that doesn't disrupt the user experience. They must comply with evolving global laws while minimizing friction for eligible users to prevent abandonment. Solution providers must factor in the complexity of regulatory requirements. Managing parental consent is also required by many regulations, but not always supported by identity vendors.

Drivers

- **Parental concerns:** The rise of social media, increased availability of age-restricted products and services, as well as ever expanding amounts of tracking (and the potential of misuse and/or breach of the accumulated data) has parents concerned about their children's online activities.
- **Compliance needs for businesses:** Organizations offering these products and services must comply with new legislation to avoid legal repercussions, including significant fines.
- **Innovation in age assurance and consent:** The new wave of regulations is prompting technology vendors to develop innovative solutions that help businesses integrate age assurance services and comply with parental consent requirements.
- **Expansion of laws and growing market:** Regulatory considerations started with a focus on processing minors' personal data, with requirements modeled after the EU's GDPR. The goals were to limit targeted advertising to minors and safeguard their personal data, with the primary difference being the age at which an individual is considered a minor. New laws are now emerging to move beyond the processing of minors' personal data and focus on restricting access to content and/or platforms, adding more complexity to the list of the requirements. This new wave of legislation expands the demand for age assurance offerings and attracts new service providers.
- **Client needs:** Organizations are seeking assistance in understanding the relevant laws and implementing compliance measures to avoid the penalties outlined in the legislation(s).
- **Age-gated content management:** Many online vendors offer a mix of products and services with varying age restrictions. For instance, a movie streaming platform may need different age assurance methods for content rated G, PG, PG-13, R, and NC-17.

Obstacles

- **Lack of traditional credentials:** Unlike adults, minors often have no formal IDs, making age and parental consent verification challenging for sellers.
- **User experience versus assurance:** If not balanced with usability, age checks and parental consent can add friction, risking customer loss.
- **Cost considerations:** Risk-based analysis is needed, as age assurance adds costs and may affect business models.
- **Parental consent technology:** This technology continues to evolve, which may hinder effective compliance in the near term.
- **Law versus market readiness:** New regulations have been observed to have differing requirements and often outpace solution provider capabilities, leading to legal uncertainty and slow adoption.
- **Free speech concerns:** Lawsuits have overturned some age assurance rules, citing infringement on legal adult access and free speech rights.

User Recommendations

- Determine whether your services are affected by age assurance or parental consent laws on a product-by-product basis.
- Choose solution providers with user-friendly age assurance methods that minimize user disruption.
- Ensure that partners understand local child protection and age assurance regulations in their service areas.
- Solicit existing solution providers first to determine whether their roadmaps include age assurance capabilities and relevant integration experience, before onboarding new vendors.
- Collect only the minimum data needed to verify a user's age and avoid retaining a minor's personal data. Regulations limit the collection and storage of such data.
- Partner with reputable age assurance vendors and consult organizations like the U.S. NIST, the EU's eIDAS, or the Age Verification Providers Association for guidance.

Sample Vendors

AU10TIX; BlueCheck; FaceTec; Jumio; Persona; PRIVO; Veridas; Veriff; Yoti

Gartner Recommended Reading

[Privacy Trends — All Eyes on Protection of Minors' Privacy and Age Assurance](#)

[Magic Quadrant for Identity Verification](#)

[Critical Capabilities for Identity Verification](#)

[Market Guide for Decentralized Identity](#)

[Enhance Trust With Decentralized Identity and Verifiable Credentials](#)

Privacy by Design

Analysis By: Bart Willemsen, Bernard Woo

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Adolescent

Definition:

Privacy by design (PbD) is no single technology. It's an approach to prioritize embedding privacy safeguards into technical designs, workflows and organizational practices. It includes using privacy-enhancing technology and consistent data minimization. The list of principles is not carved in stone, yet PbD as good practice is globally applicable to the basis of any privacy program.

Why This Is Important

Privacy is one of the core pillars for organizations seeking to earn trust with their customers and increase revenue opportunities. Regulatory requirements worldwide demand caution regarding people's information usage. Organizations can expect to operate more efficiently by adopting PbD and embedding privacy considerations throughout their processing activities. As privacy programs mature, PbD becomes a more foundational and default element, and it greatly reduces data breach impacts.

Business Impact

Privacy must be built in. A proactive approach helps enhance consumer trust and prevent data misuse or unneeded impacts from data breaches. If violations or incidents do occur, PbD helps reduce harm such as fines or brand damage. All technology design must account for the protection of any personal data at the core to mitigate privacy risk, which is at unprecedented heights with the current data volumes processed.

Drivers

- Systems and processes must be designed so that the collection of sensitive data is transparent to the data subject. Some technology-focused ideas for implementing PbD include the following, all of them alongside purpose-based access controls (PBAC):
 - Reducing retention length and amount of personal data (data minimization)
 - Working on the original data (rather than on copies)
 - Applying anonymization or pseudonymization where possible
- The need persists to continuously evaluate the risks of reidentification and traceability, and include data location in the considerations for clarity on regulatory impact.
- Implementing PbD (in)directly brings positive program changes such as a privacy officer with reach, procurement activities for new IT services, and frequent privacy impact assessments.
- PbD and one of its subcomponents, privacy engineering, enable an approach to process and architecture that combines proven practices in design, deployment and governance. Properly implemented, PbD yields more accessible functionality that fulfills a wide range of privacy principles.
- PbD involves ongoing recalculation and rebalancing of the risk to the individual data owner while preserving optimum utility for personal data processing use cases. As a result, organizations can rely on the right data being available at the right time with optimized and justifiable data retention for a compliant operation; yet shift-left technology adoption seems to stagnate for now.
- Stakeholders benefit from reducing the data footprint and accompanying breach exposure risk reduction. Further, PbD allows consistent delivery to subjects upon a privacy promise as well as collaterally enhanced customer trust and engagement levels.

Obstacles

- Adoption of PbD has been hampered by a lack of consistent interpretation of industry-recognized principles, willingness to interpret its components in-depth and consistent regulatory framework support. Singular global implementation standards remain absent.
- Legislative requirements include “data protection by design and/or by default,” implying a PbD approach to all activities. Precedent-shaping rulings emerge only intermittently. Some vendors add statements like “product X was designed with PbD in mind,” with little reference material to support the claim.
- The benefit of PbD increases only when privacy is a genuine component of development processes.

User Recommendations

- Tackle PbD in manageable steps; a wholesale shift will be too much to handle. PbD is a cultural change about the processing of personal data. This pertains both to existing operations and to innovations.
- Include responsibilities flowing from PbD application across the vendor landscape as early as possible (for example, in a request for proposal [RFP]).
- Adjust existing operations through business process reengineering. Especially in innovative developments and new processes, begin by asking questions such as: Can we achieve our purpose by using less personal data? Can we end the personal data life cycle sooner? Can we provide the same functionality or customer experience without using the identifiable data? Can we adequately protect what we process? Do customers understand what we are processing about them and why?
- Identify use cases where privacy-enhancing technologies (PETs) can be adopted to support embedding of privacy into current and future operational activities.

Gartner Recommended Reading

[CISO Effectiveness: Drive Good Privacy Practices by Consistently Asking Key Questions](#)

[Use Consistent Privacy and AI Impact Assessments to Ensure Control Criteria](#)

[Tool: AI and Privacy Impact Assessment](#)

[Privacy Imperatives for Executive Leaders](#)

How to Protect Privacy in Third-Party and Supply Chain Engagements

Use-Case Guidance for Choosing Privacy-Enhancing Technologies

Privacy Impact Assessments

Analysis By: Bart Willemsen

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Privacy impact assessments (PIAs) enable stakeholders to shape business processes and manage privacy risk. Typically conducted before implementing new processing activities and/or major changes, the PIA considers process owner and description, types of data processed for specific purposes, and retention periods per purpose. Elaboration may include legal grounds, civil liberty studies, and adequacy of security measures to ensure a controlled personal data processing environment.

Why This Is Important

Privacy mandates that organizations develop foundational insight into what personal data they process, why they do so, and how it's protected. The ability to granularly demonstrate control over personal data across various repositories and silo types is often lacking. Yet, this intentional approach is vital to adequate deployment of privacy and security controls and AI project success, making PIAs arguably the most powerful privacy management tool and cornerstone of an effective privacy program.

Business Impact

When used frequently and consistently, PIAs:

- Improve compliance and control over personal data throughout the data life cycle
- Assist in prevention of (internal) data breaches and personal data misuse
- Identify risks to people and reveal suitable mitigating controls
- Provide a basis for responsible and transparent data management

Drivers

- For organizations that have few processing activities and limited change in those processes, spreadsheets are sufficient. But, once you start juggling more than 150 processing activities, the need for better workflow and automation fast becomes a necessity.
- Overstandardization limits the skills needed to conduct PIAs to a few people rather than making them part of an organization's data-handling fabric.
- PIA automation tools allow for (API-driven) triggers to initiate the assessment process, collecting the needed information at every step and tracking it through a predefined workflow all the way until a case is closed or flagged for remediation.
- When done well, the PIA sits at the center of connecting legal requirements and business process (re)engineering to practical operationalization in privacy by design and enablement of adequate security control application.
- The PIA results on a business-process level also dictate purpose-based access controls (PBAC). This makes it a prerequisite to any AI project success.
- The entire PIA process eases data governance initiatives to a more controlled state; yet, the current main drivers still primarily come from regulatory requirements. Additional frameworks, like the 2023 revamped ISO 29134, help.

Obstacles

- Often considered a tedious task because of poorly conceived manual workflows or, opposite, a one-size-fits-all mentality, there is PIA fatigue with employees who have faced them for a decade now.
- Business partners' view of a "checkbox mentality" does not help the quality of the PIA.
- Underestimation of a PIA's relevance and position, incomplete and inaccurate PIAs, or the failure to frequently keep them updated make the initial attempt ultimately futile.
- PIA automation tools are hard to tailor to an organization's needs in the absence of knowledgeable and trained staff. As a result, even an automated approach fails to fulfill the purpose for subsequent automation and alignment of the personal data life cycle governance or management activities that are ideally connected to the PIA.

User Recommendations

- Appoint and mandate business process owners with responsibility over their respective personal data processing activities, and actively involve them in optimizing the process for fluency and detail.
- Require PIAs to be conducted following defined triggers, and periodically revalidate high-risk processing. Triage the necessity for PIAs in change processes and the introduction of new processing activities.
- Include the PIA's results, especially from large projects, in the corporate risk register for monitoring and follow up.
- Extend the assessment's effectiveness to processing of personal data carried out by service providers by demanding they complete and periodically revise a full PIA.
- Use a centrally provisioned tool for consistently conducting PIAs where scale and risk justify it (for example, as an internal automated workflow process). For low volumes, a manual exercise may suffice.

Sample Vendors

DataGrail; OneTrust; RESPONSUM; Securiti; Smart Global Governance; Transcend; TrustArc

Gartner Recommended Reading

[Privacy Program FAQs: Building an Effective Privacy Program](#)

[Use Consistent Privacy and AI Impact Assessments to Ensure Control Criteria](#)

[Tool: AI and Privacy Impact Assessment](#)

Entering the Plateau

Personification

Analysis By: Andrew Frank

Benefit Rating: Moderate

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Definition:

Personification allows marketers to deliver targeted digital ads and experiences to individuals based on their inferred membership in a characteristic customer segment without collecting or processing personal data.

Why This Is Important

Personification is a key strategy for reconciling consumer privacy concerns with the benefits of personalized experiences. Ads and experiences based on inferred characteristics and behaviors, rather than explicit personal data, can lead to more relevant and engaging interactions. This improves customer satisfaction and conversion rates while offering superior scalability and optimization potential.

Business Impact

Marketers continue to resist an absolute dependency on hyperscalers' walled gardens for ad targeting and measurement, and are seeking alternatives to explicit user consent for site personalization. Persona modeling approaches provide the keys to reducing these dependencies while staying within the bounds of privacy requirements and expectations.

Drivers

- **Technical advances:** Breakthroughs in AI are boosting the effectiveness of customer segment inferences and optimization without personal data (often with synthetic data). Content productivity gains from GenAI are driving vendors and marketers toward personalized content strategies that require effective segmentation to scale.
- **Data restrictions:** Ongoing challenges to collecting and acting on first-party data are driving marketers to seek relevant connections with customers that don't require personal data, login and tracking consent.

- **Compliance barriers:** The cost of compliance with an expanding patchwork of inconsistent regional privacy laws escalates the urgency for a global privacy-safe solution to ad targeting and measurement.
- **Cost-efficiency:** Personification can reduce the costs associated with AI computation, data storage, processing and security. The cost of rendering user-level personalized content and dialogs can quickly become prohibitive for mass-market brands.
- **Monetization opportunities:** Businesses with extensive first-party data, such as retail, travel, publishing and telecommunications, face pressure to find ways to monetize their data while staying within the bounds of privacy regulations and consumer expectations. Persona-level offerings address this need.

Obstacles

- **Data quality and availability:** Although personification reduces the need for personal data, it still relies on accurate behavioral and contextual data to reliably infer and target segments. Incomplete or inaccurate data hinders adoption efforts. Gaps can be amplified by AI overfitting when a model treats noise, outliers or incomplete records as genuine, actionable patterns.
- **Walled garden supremacy:** The scale and concentration of behavioral data and compute power at Google, Meta and Amazon far outweigh the potential effectiveness of any brand or publisher's personification efforts. Many will opt to simply rely on these platforms for ad optimization.
- **Hidden bias and manipulation:** Biases may emerge as unintended side effects of segmentation schemes. Consumers may suspect unauthorized use of personal data. Machine customers acting as proxies for buyers may defy or manipulate personification efforts.
- **AI-induced obsolescence:** Dynamically inferred, context-based relevance and presentation may obviate the need for personification.

User Recommendations

- Experiment with different personification strategies across platforms to continuously evolve and refine your ad targeting and measurement approach.

- Focus data and analytics resources on customer segmentation strategies to continuously refine and optimize persona definitions based on observed behavioral, contextual and synthetic data.
- Incorporate established clustering approaches using geography and demographics, supplemented by public and commercial data sources.
- Reevaluate personalization strategies and designs to minimize personal data requirements while maximizing opportunities for needs discovery, contextual relevance and creative impact.
- Make privacy-first data collaboration and federated learning with partners a core element of your strategy to use AI for marketing and customer insights.

Sample Vendors

Adobe; Amazon; Google; LiveRamp; Publicis Groupe; The Trade Desk; WPP

Gartner Recommended Reading

[Top Marketing Trends for 2024: AI for Privacy, Personalization and Influence](#)

Privacy Management Tools

Analysis By: Bart Willemsen

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Definition:

Privacy management tools help organizations navigate compliance insights and check processing activities against regulatory requirements. They add structure and scalability to privacy processes and workflows, enhance insight into data flows and governance maturity, and monitor and track the privacy program's efficacy.

Why This Is Important

Privacy laws worldwide require organizations to gain insight into and maintain control over personal data processing actions. Privacy management tools are the foundational technology for privacy programs and are typically the first capabilities organizations adopt. As privacy regulations span city, provincial or state, national, and cross-border jurisdictions, their complexity at a granular level makes manual privacy management impractical.

Business Impact

Privacy management tools provide business leaders with oversight and accountability for the handling of personal data, while enabling transparency and control across the organization. The toolset typically includes impact assessments, records of processing activities, and compliance and risk management reporting, among others. Over time, point solutions increasingly integrate into broader platforms, supporting automated privacy UX, vendor risk management and data intelligence inventories.

Drivers

- In almost all cases, enacted and proposed privacy laws draw heavily on the GDPR, and commonly include principles such as data subject rights, explicit consent (including tracking consent) and timely breach notification. Regulatory change will continue, reinforcing privacy as a foundational legislative requirement.
- Individual awareness of and demand for privacy continue to increase, often driven by direct exposure to the widespread adoption of new technologies, most notably AI. In response to these pressures, organizations must adapt privacy programs to scale and perform effectively within constrained budgets, without exposing the business to regulatory fines or reputational damage.
- A privacy-first posture that emphasizes transparency and control over personal data processing activities benefits the business. This is particularly true at the intersection of AI adoption and emerging regulations, such as the EU AI Act, which builds on existing privacy requirements. Establishing a strong privacy foundation enables brand protection, customer trust and an ethical approach to the individuals behind personal data.
- The privacy landscape continues to grow in complexity. More than 80% of organizations worldwide now operate under modern privacy and data protection requirements (see [Predicts 2025: Privacy in the Age of AI and the Dawn of Quantum](#)). Even early-stage privacy programs require core capabilities, including records of processing activities (ROPAs), privacy impact and compliance assessments, key elements of the privacy user experience, and incident or data breach management.
- To establish a consistent and holistic privacy program, organizations require two complementary capability sets: privacy management and data-centric controls. Some organizations prioritize data-centric foundations, addressing discovery, classification, authorization and access controls, and end-of-life data management first.

Obstacles

- Over time, a “good enough” mindset can delay the adoption of privacy management tools or prevent organizations from implementing automated integrations needed to realize full value.
- Some organizations already use point solutions for privacy UX, data breach response or impact assessments, yet overlook the benefits of more mature, complementary capabilities available in integrated suites. Efforts to stitch together disconnected tools from multiple vendors frequently yield limited results, particularly when the ongoing manual workload remains high.
- Ongoing initiatives, including the “shift left” approach that embeds certain privacy capabilities directly into code, tend to be highly technical, which can slow adoption even when these capabilities promise meaningful short- and long-term benefits.
- When organizations do not face immediate sanctions or regulatory scrutiny, perceived pressure diminishes, leading some to adopt a wait-and-see posture and continue relying on fragmented, hard-to-scale practices.

User Recommendations

- Incorporate the requirements of a rapidly evolving privacy landscape into the organization’s data strategy by establishing a common baseline informed by applicable regulatory guidance and recognized privacy frameworks.
- Maintain a focus on cross-cutting privacy capabilities that apply across the organization, including privacy impact assessments (PIAs), ROPAs, consistent vendor risk management, policy and retention management, and a people-centric privacy UX.
- Adapt privacy capabilities to evolving business needs by leading with a cost-optimized, scalable set of privacy controls that support both regulatory compliance and business agility.
- Evaluate how well privacy management tools align with organizational standardization criteria, including support for certifications and frameworks such as EuroPriSe, ISO/IEC 27701, the NIST Privacy Framework and relevant cloud codes of conduct.

Sample Vendors

DataGrail; Ketch; OneTrust; Osano; Responsum; TrustArc

Gartner Recommended Reading

[CISO Effectiveness: Drive Good Privacy Practices by Consistently Asking Key Questions](#)

[State of Privacy – The European Union](#)

[State of Privacy: Global Overview of Regulatory Evolution by Region](#)

[State of Privacy – India](#)

Appendixes

See the previous Hype Cycle: [Hype Cycle for Privacy, 2025](#)

Hype Cycle Phases, Benefit Ratings and Maturity Levels

Table 2: Hype Cycle Phases

(Enlarged table in Appendix)

Phase	Definition
<i>Innovation Trigger</i>	A breakthrough, public demonstration, product launch or other event generates significant media and industry interest.
<i>Peak of Inflated Expectations</i>	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the innovation is pushed to its limits. The only enterprises making money are conference organizers and content publishers.
<i>Trough of Disillusionment</i>	Because the innovation does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
<i>Slope of Enlightenment</i>	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the innovation's applicability, risks and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
<i>Plateau of Productivity</i>	The real-world benefits of the innovation are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
<i>Years to Mainstream Adoption</i>	The time required for the innovation to reach the Plateau of Productivity.

Source: Gartner

Table 3: Benefit Ratings

Benefit Rating	Definition
<i>Transformational</i>	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
<i>High</i>	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
<i>Moderate</i>	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
<i>Low</i>	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner

Table 4: Maturity Levels

(Enlarged table in Appendix)

Maturity Levels	Status	Products/Vendors
<i>Embryonic</i>	In labs	None
<i>Emerging</i>	Commercialization by vendors Pilots and deployments by industry leaders	First generation High price Much customization
<i>Adolescent</i>	Maturing technology capabilities and process understanding Uptake beyond early adopters	Second generation Less customization
<i>Early mainstream</i>	Proven technology Vendors, technology and adoption rapidly evolving	Third generation More out-of-box methodologies
<i>Mature mainstream</i>	Robust technology Not much evolution in vendors or technology	Several dominant vendors
<i>Legacy</i>	Not appropriate for new developments Cost of migration constrains replacement	Maintenance revenue focus
<i>Obsolete</i>	Rarely used	Used/resale market only

Source: Gartner

Evidence

2026 Gartner CISO Role-Based Survey: The purpose of this survey was to evaluate changes in cybersecurity budget allocation, assess the evolving threat landscape, and understand organizations' approach to managing nth-party risks and post quantum cryptography (PQC). The survey also included questions specifically on the CISO role — including shifts in their ownership and accountability for key cybersecurity activities, influence on board-level decisions, and the impact of stress and burnout. The online survey was administered from 18 March 2026 through 7 May 2026 and included 297 respondents from North America (n=134), EMEA (n=91), APAC (n=40), Latin America (n=32). Respondents were required to be the most senior leader of the cybersecurity function (CISO or equivalent) and came from a wide variety of industries, including banking/investment services (n=76), technology products or services (n=58), manufacturing (n=35) and others. Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.

Document Revision History

[Hype Cycle for Privacy, 2025 - 28 July 2025](#)
[Hype Cycle for Privacy, 2024 - 29 July 2024](#)
[Hype Cycle for Privacy, 2023 - 24 July 2023](#)
[Hype Cycle for Privacy, 2022 - 2 August 2022](#)
[Hype Cycle for Privacy, 2021 - 13 July 2021](#)
[Hype Cycle for Privacy, 2020 - 23 July 2020](#)
[Hype Cycle for Privacy, 2019 - 11 July 2019](#)
[Hype Cycle for Privacy, 2018 - 17 July 2018](#)
[Hype Cycle for Privacy, 2017 - 20 July 2017](#)
[Hype Cycle for Privacy, 2016 - 20 July 2016](#)
[Hype Cycle for Privacy, 2015 - 24 July 2015](#)
[Hype Cycle for Privacy, 2014 - 18 July 2014](#)
[Hype Cycle for Privacy, 2013 - 31 July 2013](#)
[Hype Cycle for Privacy, 2012 - 25 July 2012](#)

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[Understanding Gartner's Hype Cycles](#)

[Tool: Create Your Own Hype Cycle With Gartner's Hype Cycle Builder](#)

[State of Privacy: The Privacy Tech Driving a New Age of Data Wealth](#)

[Market Guide for Consent and Preference Management](#)

[Adopt 5 Key Privacy Lessons When Preparing for AI Regulations](#)

[State of Privacy: Global Overview of Regulatory Evolution by Region](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's Business and Technology Insights Organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner insights may address legal and financial issues, Gartner does not provide legal or investment advice and its insights should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its insights is produced independently by its Business and Technology Insights Organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner insights may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

Table 1: Priority Matrix for Privacy, 2026

Benefit ↓	Years to Mainstream Adoption			
	Less Than 2 Years ↓	2 to 5 Years ↓	5 to 10 Years ↓	More Than 10 Years ↓
Transformational	Sovereign Cloud	Decentralized Identity Homomorphic Encryption	Influence AI	
High	Data Breach Response Digital Ethics Privacy Management Tools Subject Rights Requests	Age Assurance and Parental Consent Tracking AI TRiSM Confidential Computing Privacy Impact Assessments	Federated Machine Learning	
Moderate	Format-Preserving Encryption Personification	5G Network Security AI Security Posture Management Consent and Preference Management Data Clean Rooms Data Discovery Data Security Governance Privacy by Design Synthetic Data		
Low		Differential Privacy		

Source: Gartner (June 2026)

Table 2: Hype Cycle Phases

Phase	Definition
<i>Innovation Trigger</i>	A breakthrough, public demonstration, product launch or other event generates significant media and industry interest.
<i>Peak of Inflated Expectations</i>	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the innovation is pushed to its limits. The only enterprises making money are conference organizers and content publishers.
<i>Trough of Disillusionment</i>	Because the innovation does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
<i>Slope of Enlightenment</i>	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the innovation's applicability, risks and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
<i>Plateau of Productivity</i>	The real-world benefits of the innovation are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
<i>Years to Mainstream Adoption</i>	The time required for the innovation to reach the Plateau of Productivity.

Source: Gartner

Table 3: Benefit Ratings

Benefit Rating	Definition
<i>Transformational</i>	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
<i>High</i>	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
<i>Moderate</i>	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
<i>Low</i>	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner

Table 4: Maturity Levels

Maturity Levels	Status	Products/Vendors
<i>Embryonic</i>	In labs	None
<i>Emerging</i>	Commercialization by vendors Pilots and deployments by industry leaders	First generation High price Much customization
<i>Adolescent</i>	Maturing technology capabilities and process understanding Uptake beyond early adopters	Second generation Less customization
<i>Early mainstream</i>	Proven technology Vendors, technology and adoption rapidly evolving	Third generation More out-of-box methodologies
<i>Mature mainstream</i>	Robust technology Not much evolution in vendors or technology	Several dominant vendors
<i>Legacy</i>	Not appropriate for new developments Cost of migration constrains replacement	Maintenance revenue focus
<i>Obsolete</i>	Rarely used	Used/resale market only

Source: Gartner