emerging tech decoded

LEX CONSULT EDITORIAL

BYTE/WISE

KEY HIGHLIGHTS

JAN 2026

Artificial Intelligence

Regulatory Updates

- Office of the Principal Scientific Adviser, India releases Techno-Legal White Paper on Safe & Trusted AI
- Singapore releases Governance Framework for Agentic AI Systems
- China publishes Draft Provisional Measures on the Administration of Human-Like Interactive Artificial Intelligence Services
- UK's FCA Launches Mills Review on Advanced AI Risks in Retail Financial Services
- Grok Triggers Cross-Jurisdictional Scrutiny Over Synthetic Images and Deepfake Risks

Clause in Focus - Serious Incident Handling & Reporting

Knowledge Corner - Incident vs. Serious Incident

Compliance Snapshot - AI Vendor & Third-Party Risk Controls

Privacy

Regulatory Updates

- EU Adequacy Decisions (Brazil and UK)
- DPDP Implementation Accelerates Across Financial And Healthcare Sectors amid Signals of Shortened Compliance Timeline
- California's Key CCPA and CPRA Regulatory Obligations Come into Effect
- France's Data Protection Authority Imposes High-Profile GDPR Fines

Clause in Focus - Audit & Inspection Rights

Knowledge Corner - Reasonable Security Safeguards

Compliance Snapshot - Handling Data Subject Requests

Digital Assets

Regulatory Updates

- FIU-IND's FY 2024–25 Annual Report: Signals for Digital Asset Regulation in India
- Europe's Stablecoin Moment: Regulation First, Adoption Later?
- U.S. Senate Advances Crypto Market Structure Bill: Key Regulatory and Legal Implications

Clause in Focus - Custody Termination & Asset Migration Clauses

Knowledge Corner - Privacy Coins

Dark Patterns Check – Drip Pricing

BYTE/WISE

PART ONE

Ai

JAN 2026

Ai / REGULATORY UPDATES

Office of the Principal Scientific Adviser, India releases Techno-Legal White Paper on Safe & Trusted AI

In early January 2026, India's Office of the Principal Scientific Adviser released a *Techno-Legal White Paper on Safe & Trusted AI*, proposing a shift toward more operationally grounded AI governance. The paper situates itself alongside earlier government initiatives such as the India AI Mission's Safe & Trusted AI pillar and India's AI governance guidelines but moves beyond high-level principles to focus on how legal accountability can be translated into enforceable and auditable technical controls across the AI lifecycle.

A central feature of the paper is its proposed institutional architecture, including an AI Governance Group (AIGG), a Technology and Policy Expert Committee (TPEC), an AI Safety Institute (AISI), and an AI Incident Database to support monitoring and evidence-based refinement of controls.

In addition, the key proposals of this paper include:

- Compliance-by-design: AI regulation must be embedded directly into system design and workflows, spanning data collection, model training, inference, and deployment.

- Lifecycle-based governance: AI risk is addressed across the full lifecycle, from data sourcing to inference and agentic operation, recognizing that compliance failures can arise upstream or downstream.

- Agentic AI as a distinct risk category: Autonomous and tool-using AI systems are treated as a separate governance stage, requiring controls such as authorization, behavioral logging, and human override mechanisms.

- Anchoring in existing Indian law: AI governance is operationalized through current legal frameworks, including the IT Act, DPDP Act, BNS 2023, and sectoral regulations, rather than a standalone AI statute.

- Risk and scale-proportionate obligations: Compliance expectations are explicitly tied to deployment scale and impact, echoing the DPDP Act's Significant Data Fiduciary approach.

- Demonstrable accountability: Accountability through documentation, logs, traceability, and audit trails across the AI value chain.

- Deepfakes as a pipeline-level risk: Deepfakes are addressed as an ecosystem issue involving generation, distribution, platforms, and infrastructure.

These proposals point toward a governance model in which AI compliance is likely to be evaluated through demonstrable oversight structures, documented risk controls, and incident readiness across the system lifecycle.

Singapore releases Governance Framework for Agentic AI Systems

In January 2026, Singapore's Infocomm Media Development Authority (IMDA) released an updated governance framework addressing the deployment of agentic AI systems, focusing on AI applications capable of autonomous decision-making, task execution, and adaptive behavior. The framework builds on Singapore's earlier Model AI Governance Framework but responds to the growing use of AI agents in enterprise workflows, customer-facing services, and operational decision-making, where autonomy and delegation of tasks raise distinct accountability and risk-management concerns. Unlike the EU's rights- and risk-classification model or China's prescriptive, behavior-restriction approach, Singapore frames AI governance primarily through an enterprise risk management lens, emphasizing internal controls, auditability, and supervisory accountability.

The framework places particular weight on clear allocation of responsibility for AI agent behavior, including delegated tasks and downstream actions, and sets out expectations around lifecycle governance such as use-case scoping, human oversight, change management, and post-deployment monitoring. It emphasizes operational controls for agentic behavior, including defined action boundaries or "significant checkpoints" requiring human approval for high-risk or irreversible actions, supported by robust monitoring, logging, and regular audits of human oversight to mitigate automation bias and alert fatigue.

China publishes Draft Provisional Measures on the Administration of Human-Like Interactive Artificial Intelligence Services

In late December 2025, China's Cyberspace Administration of China (CAC) released for public consultation the *Draft Provisional Measures on the Administration of Human-Like Interactive Artificial Intelligence Services*, targeting AI systems that simulate human personality traits or engage users through emotionally interactive, conversational, or companion-style interfaces. The draft measures are grounded in China's existing legal framework and apply specifically to AI services that engage users through human-like emotional or interactive behavior addressing specific regulatory concern around psychological influence, user dependency, and behavioral manipulation arising from increasingly human-like AI systems deployed in consumer-facing contexts.

The proposed measures impose concrete obligations on providers of human-like interactive AI services, such as safety responsibility throughout the product lifecycle, establishing systems for algorithmic review, data security, personal information protection, emergency response planning, and assessing users' emotional states and dependence so that necessary interventions can be made if extreme emotions or addiction are detected.

Providers must also warn users they are interacting with AI, set reminders after prolonged use (e.g., two hours), prohibit services designed to replace real social interaction or induce dependency, and ensure services do not generate content that threatens national security, undermines social order, promotes violence, obscenity, gambling, or harms users' physical or mental health.

UK's FCA Launches Mills Review on Advanced AI Risks in Retail Financial Services

In January 2026, the UK Financial Conduct Authority (FCA) launched the *Mills Review* to examine how advanced AI could reshape retail financial services, with a particular focus on risks to competition, consumer outcomes, and market integrity. While the review does not propose AI-specific rules, the FCA has explicitly framed advanced AI as a potential source of regulatory risk, including the concentration of market power among data-rich firms, AI-driven distortions in competition, and new forms of consumer detriment arising from opaque or behavior-shaping automated systems. The FCA has also highlighted concerns around algorithmic bias, reduced consumer agency, and the use of AI in facilitating fraud or financial crime, alongside the possibility that rapid AI deployment could outpace existing supervisory tools. The FCA has indicated that the review's findings will be reported to its board in mid-2026, with any implications for supervisory strategy or regulatory intervention to follow.

Grok Triggers Cross-Jurisdictional Scrutiny Over Synthetic Images and Deepfake Risks

In January 2026, Elon Musk–backed AI system Grok became the focus of regulatory scrutiny across multiple jurisdictions following concerns over its ability to generate and circulate synthetic and manipulated images, including sexually explicit, abusive, or misleading content involving real individuals. Regulators and data protection authorities have raised questions about whether such outputs engage existing legal frameworks governing biometric data, non-consensual sexual imagery, and deceptive or harmful AI-generated content, particularly where images are realistic enough to enable identification or impersonation.

In India, the Ministry of Electronics and Information Technology (MeitY) issued a notice under the IT Act 2000 and the IT Rules 2021, to the platform directing the removal of obscene and sexually explicit Grok-generated content and requiring an Action Taken Report within 72 hours, later extended to 7 January 2026. Following the MeitY directive, the platform removed approximately 3,500 pieces of Grok-generated obscene content and deleted over 600 associated user accounts, and assured authorities of compliance with Indian law going forward. This enforcement action also illustrates the practical relevance of India's recently proposed amendments addressing deepfakes and harmful synthetic content, discussed in earlier Bytewise editions, which seek to clarify intermediary accountability and takedown obligations for AI-generated harms even in the absence of AI-specific legislation.

Brazil's data protection authority examined whether certain synthetic images generated by Grok could qualify as biometric data under data protection law, and indicated that non-consensual, realistic synthetic images capable of identifying individuals may engage data protection obligations and constitute unlawful processing where generated or disseminated without a valid legal basis. Following this, UK policymakers have also accelerated efforts to criminalize sexual deepfakes and tighten platform responsibilities for AI-generated abuse. These developments indicate increasing regulatory willingness to treat synthetic media risks as a cross-cutting AI governance issue, spanning privacy, content safety, and consumer protection, with regulatory scrutiny extending beyond model capability to the adequacy of safeguards, downstream misuse controls, and post-deployment governance.

Ai / CLAUSE IN FOCUS

SERIOUS INCIDENT HANDLING & REPORTING

As AI systems move from assistive tools to operational decision engines, incident response is no longer a purely technical workflow but a core compliance and liability control. Under emerging AI governance frameworks, particularly the EU AI Act, serious incidents are treated as compliance-relevant events, with providers bearing primary responsibility for regulatory reporting and post-market oversight, while deployers are expected to detect and escalate incidents arising in real-world use and to cooperate without delay to enable timely reporting, investigation, and corrective action.

The EU AI Act requires providers of high-risk AI systems to report 'Serious Incidents' (further elaborated in our knowledge Corner) to competent authorities without undue delay and, in any event, no later than 15 (fifteen) days after becoming aware of the incident, with shorter timelines applying where death or serious harm is involved. Deployers are required to inform providers without delay and to cooperate to enable compliance with statutory reporting obligations.

Draft Illustrative Clause

"The Deployer shall notify the Provider without undue delay upon becoming aware of any incident or malfunction that has resulted in, or is reasonably likely to result in, a serious incident under applicable law. The Deployer shall provide such information and assistance as is reasonably necessary to enable the Provider to assess the incident and comply with any applicable regulatory or statutory reporting obligations within prescribed timelines. The parties shall cooperate in good faith to implement and document corrective and preventive measures proportionate to the risk."

Negotiation Points

Deployer

- Ensure that notification and cooperation obligations are triggered only where the incident meets the statutory thresholds such as "serious incident" under the EU AI Act.
- Limit cooperation and information-sharing duties to matters within the Deployer's control, particularly where incidents arise from provider-side defects or undisclosed system limitations.
- Ensure cooperation obligations remain proportionate and do not evolve into audit or inspection rights.
- Clarify that incident notification or cooperation does not, by itself, constitute an admission of fault, liability, or non-compliance.

Provider

- Require notification without undue delay to enable timely assessment and compliance with mandatory regulatory reporting timelines.
- Seek sufficient information and cooperation to assess the nature, impact, and regulatory significance of the incident.
- Preserve coordinated regulatory engagement where the Provider bears primary statutory reporting responsibility.
- Avoid contractual formulations that could be interpreted as shifting regulatory liability to the Provider solely due to its reporting role.

Balanced drafting should align incident reporting and cooperation duties with statutory requirements and real deployment responsibility, enabling regulatory compliance without expanding liability beyond what the law requires.

Ai / COMPLIANCE SNAPSHOT

AI Vendor & Third-Party Risk Controls

As organizations increasingly rely on third-party AI providers, system integrators, and data vendors, regulatory attention is shifting from individual system compliance to the governance of AI supply chains, particularly where high-risk AI systems are involved. Under the EU AI Act, statutory obligations relating to risk management, post-market monitoring, serious incident handling, and documentation cannot be outsourced through procurement arrangements.

While AI vendors most commonly act as providers, deployers remain accountable for how AI systems are selected, integrated, and used in practice, and must ensure that vendor arrangements support compliance with statutory requirements, including serious incident escalation and reporting obligations. Against this backdrop, deployers should assess whether their vendor arrangements are structured to support compliance with these statutory obligations across the AI system lifecycle.

Key Compliance Readiness Questions

- Have party roles been clearly mapped as provider, deployer, distributor, or ancillary service provider under the EU AI Act framework, particularly for high-risk AI systems?
- Do contracts require vendors to provide information, documentation, and cooperation necessary to support statutory obligations such as post-market monitoring and serious incident reporting?
- Are escalation and incident-notification pathways aligned across vendors to avoid delays in identifying, classifying, or reporting serious incidents?
- Are changes to models, datasets, or system functionality by vendors subject to notification or approval where they may affect risk classification, performance, or compliance status?
- Do termination or exit arrangements preserve access to logs, documentation, and records required to demonstrate compliance with regulatory obligations?

Effective management of AI vendor and third-party risk requires integrating AI-specific controls into existing vendor governance frameworks. This includes risk-based due diligence at onboarding, contractual alignment with statutory responsibilities under the EU AI Act, and ongoing monitoring rather than one-time compliance checks. Treating vendor oversight as a continuous compliance function is particularly critical where high-risk AI systems are deployed, as gaps in third-party cooperation can directly undermine compliance with serious incident reporting, accountability, and supervisory expectations.

Ai

KNOWLEDGE
CORNER

Incident vs. Serious Incident

BYTE/WISE

Defined Terms (EU AI Act)

Under the EU AI Act, a '*serious incident*' is defined as an incident or malfunctioning of an AI system that directly or indirectly leads to any of the following:(a) the death of a person, or serious harm to a person's health; (b) a serious and irreversible disruption of the management or operation of critical infrastructure; (c) the infringement of obligations under Union law intended to protect fundamental rights; (d) serious harm to property or the environment.

Explanation

In AI governance, not every malfunction, error, or unexpected output qualifies as a regulatory event. Most AI systems generate incidents in the ordinary course of operation, ranging from inaccurate outputs and performance degradation to bias signals or system instability. These incidents may require internal remediation, model tuning, or operational fixes, but they do not automatically trigger statutory reporting or external notification obligations. A model producing inaccurate outputs, biased recommendations, or unstable performance may constitute an incident, but it becomes a serious incident when those failures materially affect individuals, public services, or legally protected interests. The same technical failure may therefore be non-reportable in one context and a serious incident in another, depending on how and where the system is deployed.

For example, an AI-based recruitment tool that occasionally ranks candidates inaccurately due to data quality issues presents an incident that should be corrected through internal review and model tuning. If, however, the same system systematically excludes qualified applicants in a manner that materially affects access to employment, the issue may qualify as a serious incident requiring escalation and regulatory reporting, as it involves a potential infringement of fundamental rights. The distinction matters because regulatory obligations attach only once this seriousness threshold is crossed. Ordinary incidents are generally managed through internal monitoring, testing, and corrective controls as part of day-to-day AI governance. Serious incidents activate time-bound escalation and reporting duties, structured cooperation between deployers and providers, and heightened scrutiny of post-deployment risk management.

For organizations, the governance challenge lies in distinguishing between internal best practices and legally mandated duties. Good governance may involve addressing a wide range of model errors and risks as a matter of quality or ethics, but regulatory requirements under the EU AI Act attach only once an issue qualifies as a serious incident under law. Systems and teams must therefore be able to identify when an operational issue crosses from desirable internal management into required regulatory escalation. Clear internal thresholds, documented escalation workflows, and coordination between technical, legal, and business functions help ensure that statutory obligations are met without treating every incident as a reportable event.

BYTE/WISE

PART TWO

PRIVACY

JAN 2026

Privacy / REGULATORY UPDATES

EU Adequacy Decisions (Brazil and UK)

- EU–Brazil Finalize Data Transfer Adequacy Arrangement

  In January 2026, the European Union and Brazil finalized an adequacy arrangement recognizing Brazil's data protection framework as providing a level of protection essentially equivalent to that under EU law. This enables personal data to flow from the EU to Brazil without additional transfer mechanisms such as Standard Contractual Clauses or binding corporate rules, reducing compliance friction for EU-based organizations operating in or transferring data to Brazil.

- EU–UK Data Transfer Adequacy Decisions Renewed

  On 19 December 2025, the European Commission renewed the EU–UK adequacy decisions first adopted in 2021, confirming that personal data may continue to flow freely from the EU to the UK until 27 December 2031. The renewal followed a six-month extension granted in June 2025 to assess UK legal developments, including the Data (Use and Access) Act, and was supported by a favorable opinion from the European Data Protection Board and approval by Member States. A further review is scheduled after four years, underscoring that continued adequacy remains subject to ongoing regulatory monitoring.

DPDP Implementation Accelerates Across Financial And Healthcare Sectors amid Signals of Shortened Compliance Timeline

As the Digital Personal Data Protection Act, 2023 (DPDP Act) and the Digital Personal Data Protection Rules, 2025 move through phased implementation, regulators and industry bodies have signaled heightened compliance expectations for financial services and healthcare entities, given the volume and sensitivity of personal data processed in these sectors. Banks, fintechs, insurers, hospitals, and health service providers are stepping up investments in privacy governance by strengthening consent management processes, implementing reasonable security safeguards, formalizing data principal rights-handling mechanisms, and improving breach detection and reporting workflows. Many organizations are also revisiting processor arrangements and internal escalation frameworks to align systems and documentation with statutory requirements relating to consent withdrawal, correction and erasure of personal data, and grievance redressal.

Recent reports indicate that the Ministry of Electronics and Information Technology (MeitY) is considering shortening the DPDP compliance window to 12 months, particularly for large or significant data fiduciaries, signaling an expectation of earlier operational readiness than initially anticipated. Organizations should therefore avoid assuming the availability of the full originally anticipated 18-month compliance period and instead prioritize early implementation of DPDP-compliant governance and operational controls.

California's Key CCPA and CPRA Regulatory Obligations Come into Effect

Amendments to the California Consumer Privacy Act (CCPA) regulations, adopted under the California Privacy Rights Act (CPRA), came into effect on 1 January 2026, introducing new operational obligations for covered businesses.

The updated framework introduces expanded requirements for conducting privacy risk assessments and mandatory cybersecurity audits in relation to processing activities that present heightened risks to consumers, reflecting a shift toward more structured accountability under California privacy law. It also expands governance and transparency obligations around the use of automated decision-making technologies, including profiling and automated processing, and strengthens requirements for recognizing and honoring consumer opt-out preferences. Corresponding updates to privacy disclosures and internal compliance processes are required to ensure that these obligations are operationalized in practice.

Businesses should assess whether their data practices trigger the new audit or risk assessment requirements, map any use of automated decision-making, and update internal governance, technical controls, and privacy notices to align with the revised regulations.

France's Data Protection Authority Imposes High-Profile GDPR Fines

In January 2026, France's data protection authority, the CNIL, issued significant GDPR sanctions underscoring its continued enforcement focus on core privacy obligations. On 13 January 2026, the CNIL imposed combined fines of €42 million on telecom entities Free Mobile (€27 million) and Free (€15 million) following a 2024 data breach affecting approximately 24 million subscribers, including exposure of sensitive information such as bank account details. The authority cited failures relating to security of processing and breach management under the GDPR.

Alongside this high-profile action, the CNIL also continued enforcement activity into late 2025 in relation to unlawful data use for advertising purposes. In December 2025 (publicly reported in January 2026), it imposed a €3.5 million fine on a company for transferring customer contact data to a social media platform for targeted advertising without a valid legal basis.

Privacy/ CLAUSE IN FOCUS

AUDIT AND INSPECTION RIGHTS

Audit and inspection rights are a core contractual mechanism for giving effect to accountability obligations under modern data protection laws. Under the GDPR, audit rights flow directly from the accountability principle (Article 5(2)), the obligation to implement and oversee appropriate technical and organizational measures (Article 24), and the mandatory requirement under Article 28(3)(h) for data processing agreements to permit and contribute to audits or inspections.

The DPDP Act does not prescribe equivalent mandatory audit language in processor contracts. However, Sections 8 and 9 of the DPDP Act require data fiduciaries to ensure that personal data is processed in accordance with the Act and to remain responsible for compliance even where processing is outsourced.

Draft Illustrative Clause

"The Service Provider shall make available to the Customer all information reasonably necessary to demonstrate compliance with applicable data protection laws and the obligations set out in this Agreement, and shall allow for and contribute to audits, including inspections, conducted by the Customer or another auditor mandated by the Customer. Where any non-compliance is identified, the Service Provider shall take appropriate remedial measures without undue delay."

Negotiation Points

Customer/Data Fiduciary

- Ensure audit and inspection rights are sufficiently broad to satisfy statutory accountability obligations, including the ability to conduct audits following personal data breaches, supervisory inquiries, or material compliance concerns.
- Ensure audit scope covers sub-processors, cross-border data flows, and security safeguards.
- Retain flexibility to rely on independent third-party audits where on-site inspections are impractical.
- Require timely remediation of identified non-compliance in line with regulatory expectations

Service Provider/Data Processor

- Limit audit scope and frequency to what is reasonably necessary to demonstrate compliance with applicable statutory data protection obligations, rather than broader commercial or operational oversight.
- Ensure audits are subject to reasonable procedural safeguards, including advance notice, and confidentiality protections.
- Push for alternative mechanisms such as recognized third-party certifications or audit reports (such as ISO-aligned assessments) as the primary means of demonstrating compliance, with on-site audits reserved for exceptional cases.
- Clarify that remediation of identified non-compliance is without prejudice to any admission of fault, liability, or breach, and does not expand the Service Provider's obligations beyond applicable statutory and contractual requirements.

Balanced negotiation should recognize that audit and inspection rights are intended to support regulatory accountability rather than operate as a general commercial control mechanism, and should be exercised in a manner that is proportionate, minimizes operational disruption, and preserves the confidentiality of the Service Provider's information, including its security mechanisms, technical architecture, and other clients' confidential data.

BYTE/WISE

Privacy/ COMPLIANCE SNAPSHOT

Handling Data Subject Requests

Compliance with data subject and data principal rights is a core statutory obligation under modern data protection laws, and failures in handling such requests are treated as standalone compliance violations. These rights include, access to information about processing, correction and erasure of personal data, grievance redressal, and nomination. Organizations acting are expected to operationalize clear, timely, and accountable mechanisms for receiving, assessing, and responding to rights requests.

Under the GDPR, Articles 12 to 22 require controllers to facilitate and respond to data subject requests within prescribed timelines, generally one month, with processors obligated under Article 28 to support controllers through contractual and operational cooperation.

The DPDP Act adopts a similar model under Sections 11 to 14, read with the DPDP Rules, 2025, requiring data fiduciaries to provide accessible mechanisms and to respond within defined timelines, including a 90-day outer limit for grievance redressal. In both frameworks, primary responsibility rests with the controller/data fiduciary, even where requests are received downstream by processors or sub-processors, who are expected to promptly route requests rather than respond independently.

Controllers/data fiduciaries must also take reasonable steps to verify the identity of requestors to prevent unauthorized disclosure, while the DPDP framework recognizes corresponding responsibilities on data principals to provide accurate information and use designated mechanisms.

Quick Internal Compliance Check

- Are clear and accessible mechanisms in place for data subjects/data principals to submit rights requests?
- Are responsibilities clearly assigned internally for receiving, authenticating, assessing, and responding to rights requests within statutory timelines?
- Are reasonable identity-verification steps built into the rights-handling process to prevent unauthorized disclosure of personal data?
- Are processors and sub-processors contractually required to promptly forward rights requests to the appropriate controller/fiduciary?
- Are statutory timelines tracked and monitored, including any permitted extensions or exemptions relied upon under applicable law?
- Are decisions on rights requests, including reliance on exemptions or refusals, documented and retained to demonstrate compliance if required?

Privacy

KNOWLEDGE
CORNER

Reasonable Security Safeguards

BYTE/WISE

Reasonable Security Safeguards

Security safeguards sit at the heart of modern data protection regimes, setting the minimum standards for how organizations are expected to protect personal data from unauthorized access, misuse, loss, or disclosure. Across jurisdictions, data protection laws increasingly rely on risk-based security obligations rather than prescriptive technical rules, reflecting the diversity of data processing activities and threat environments.

Under the GDPR, security obligations are anchored in Article 32, which requires controllers and processors to implement "appropriate technical and organizational measures" to ensure a level of security appropriate to the risk. Factors such as the nature of the data, potential impact on individuals, state of the art, and cost of implementation are expressly relevant.

The DPDP Act similarly requires data fiduciaries to implement reasonable security safeguards to prevent personal data breaches.

Explanation

Although the GDPR and the DPDP Act use different statutory language, they ask the same core question: has the organization taken sensible and proportionate steps to protect personal data, given the risks involved?

Neither regime expects perfect or absolute security; instead, both focus on whether security decisions are informed, proportionate, and capable of addressing foreseeable personal data breaches.

For example, frameworks such as ISO/IEC 27001 (including its updated versions), the NIST Cybersecurity Framework, and comparable national or sectoral standards provide a common language for identifying risks, selecting proportionate controls, assigning responsibility, and reviewing effectiveness over time. These frameworks are widely regarded as aligning with data protection requirements across jurisdictions. In the Indian context, the Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules, 2011 expressly recognized compliance with ISO/IEC 27001 as satisfying the requirement of "reasonable security practices and procedures", subject to periodic audit and certification.

However, what is considered reasonable is not one-size-fits-all and depends on factors such as organizational size, the nature and volume of personal data processed, and the potential harm to individuals. Smaller or early-stage organizations are not expected to implement the same level of formalization as large enterprises, but they are expected to make conscious, documented security decisions and to review their safeguards as risks evolve. Regular assessment and adaptation, rather than one-time implementation, is therefore central to meeting statutory security obligations.

To demonstrate compliance with reasonable security safeguards, organizations are generally expected to adopt a combination of technical and organizational measures that reflect informed risk assessment and ongoing oversight. This typically includes documented information security policies and incident response procedures, access controls and encryption calibrated to the sensitivity of the data, periodic testing and review of safeguards, and contractual alignment to ensure that processors and vendors are subject to equivalent security obligations.

BYTE/WISE

PART THREE

DIGITAL ASSETS

JAN 2026

BYTE/WISE

FIU-IND's FY 2024–25 Annual Report: Signals for Digital Asset Regulation in India

The Financial Intelligence Unit - India (FIU-IND) has released its annual report for FY 2024-25 (Report), offering important insights into India's evolving approach to regulating virtual digital assets (VDAs). The Report highlights two clear regulatory priorities: (i) expanding oversight of offshore virtual digital asset service providers (VDASPs) operating in India, and (ii) strengthening data-driven enforcement through enhanced inter-agency coordination.

A key trend identified by FIU-IND is the intensification of enforcement against offshore VDASPs for non-registration and deficiencies in AML/CFT controls. Following its initial actions in December 2023, FIU-IND issued show cause notices to 25 offshore VDASPs in October 2025. Enforcement measures have included monetary penalties, corrective directions, and platform blocking orders. Penalties have varied significantly, reflecting FIU-IND's willingness to calibrate sanctions based on the level of transparency and cooperation demonstrated by the regulated entity. The registration framework for VDASPs has also become more rigorous. Recent procedural changes require in-person interactions, live demonstrations of compliance systems, cybersecurity audits, and third-party certifications. The emphasis has clearly shifted from policy documentation to demonstrable operational readiness, with regulators assessing real-time KYC, transaction monitoring, blockchain analytics, and travel rule compliance capabilities.

Another notable development by FIU-IND is enhanced coordination across regulators and law enforcement agencies. FIU-IND has entered into memoranda of understanding with domestic regulators, including the Reserve Bank of India, and expanded international cooperation through global FIU networks. Registered VDASPs are now also required to engage with cybercrime reporting frameworks, reflecting growing concern around retail user protection.

The Report also underscores FIU-IND's move toward data-led supervision. Analysis of suspicious transaction reports filed by VDASPs is being used to identify emerging risk typologies and support faster, coordinated enforcement actions.

Looking ahead to 2026, FIU-IND's approach suggests a more mature, integrated regulatory environment. VDASPs, particularly offshore players, will need to embed robust compliance cultures, data transparency, and reporting discipline to sustainably operate and expand within India's digital assets ecosystem.

Europe's Stablecoin Moment: Regulation First, Adoption Later?

When Europe rolled out the Markets in Crypto-Assets (MiCA) regulation in June 2024, it was widely seen as a milestone. For the first time, a major economy had introduced a comprehensive rulebook for crypto assets, including stablecoins. MiCA promised clarity, consumer protection, and regulatory consistency across the EU, something the industry had long asked for.

Yet nearly two years on, Europe's stablecoin market tells a different story. Dollar-pegged stablecoins continue to dominate globally, with U.S. based issuers accounting for the vast majority of circulation. European alternatives remain niche. Even the most prominent euro-denominated stablecoin has a market presence that barely registers when compared to USD backed peers.

One reason lies in how Europe approached the problem. MiCA was built on a "compliance-first" philosophy, prioritizing risk management and consumer

safeguards from day one. While that approach brings credibility, it has also made stablecoin issuance capital-intensive and commercially challenging. Reserve requirements, in particular, limit issuers' ability to generate returns, dampening incentives to scale. Regulation, on its own, rarely creates innovation, it usually follows it.

Meanwhile, the policy focus on a potential digital euro has added to the uncertainty. While central bank digital currencies may have a role in preserving monetary sovereignty, they are not designed to compete with stablecoins on speed, programmability, or global reach. Stablecoins thrive because they solve real problems like cross-border payments, on-chain settlement, and integration with decentralized finance. A digital euro, positioned as "digital cash," addresses a very different use case.

That said, Europe may still have a second chance. A consortium of major European banks have announced plans to launch a euro-denominated stablecoin, combining regulatory credibility with institutional scale. If executed well, backed by real use cases beyond crypto trading, this could finally give Europe a meaningful foothold in the stablecoin economy.

U.S. Senate Advances Crypto Market Structure Bill: Key Regulatory and Legal Implications

U.S. senators have released draft legislation proposing a comprehensive market structure framework for digital assets, marking a potentially significant shift in how cryptocurrencies are regulated at the federal level. If enacted, the bill would address long-standing jurisdictional ambiguity between U.S. financial regulators and provide statutory clarity that has, until now, largely been shaped through enforcement actions and litigation. At the core of the proposal is a functional classification regime for crypto tokens, seeking to delineate when a digital asset constitutes a security, a commodity, or falls outside both categories. This clarification is particularly relevant in light of recent judicial scrutiny of the Securities and Exchange Commission's (SEC) expansive interpretation of securities laws in the crypto context. The bill would also formally allocate primary oversight of spot crypto markets to the Commodity Futures Trading Commission (CFTC), reflecting industry preference and reinforcing the CFTC's role beyond derivatives markets.

The draft legislation also revisits the federal stablecoin framework enacted last year. Responding to concerns raised by the banking sector, the bill would prohibit crypto intermediaries from paying interest solely for the passive holding of dollar-pegged stablecoins, while permitting activity-based rewards tied to payments or customer engagement. To mitigate consumer protection concerns, the SEC and CFTC would be required to jointly prescribe disclosure standards governing such incentive arrangements. From a legislative process perspective, the bill will be considered by the Senate Banking Committee, alongside parallel efforts by the Senate Agriculture Committee. While the House of Representatives passed its own market structure bill in mid-2025, prior negotiations stalled over anti-money-laundering obligations and the treatment of decentralized finance protocols.

With the 2026 midterm elections approaching, the bill's legislative timeline remains uncertain. Nevertheless, the introduction of this draft signals a meaningful attempt by Congress to transition crypto regulation from an enforcement-led model to a rules-based statutory framework, an outcome with significant implications for compliance strategy, regulatory risk, and future litigation in the digital assets sector.

Digital Assets / CLAUSE IN FOCUS

LIEN CLAUSES IN CRYPTOCURRENCY CUSTODIAN AGREEMENTS:

In traditional financial services, lien and set-off rights are commonly used to protect custodians and intermediaries against unpaid fees and losses. Crypto custodians seek similar protections, but the risks they face are significantly different. Blockchain transactions are irreversible, digital assets are highly volatile, and custodians may be required to freeze or segregate assets due to regulatory or law-enforcement actions. In addition, custodians often incur operational and third-party costs linked to specific clients. A lien clause allows custodians to manage these risks and secure their dues through a practical self-help mechanism, without having to immediately resort to litigation or regulatory escalation.

Draft Illustrative Clause

Lien

1.1 The Custodian shall have a continuing general lien and right of retention over all Digital Assets held for the Client (including any residual, suspended, restricted, or quarantined Digital Assets) as security for the payment and discharge of all bona fide, due, and payable obligations of the Client arising under this Custody Agreement. Pending satisfaction of such obligations, the Custodian may, to the extent reasonably necessary, decline to act on any withdrawal or transfer instruction issued by the Client in respect of the Digital Assets subject to such lien.

1.2 Where any such obligation remains unpaid following written notice and the expiry of a reasonable cure period, the Custodian may enforce its lien by applying or realising the relevant Digital Assets, including by sale or other commercially reasonable means, and applying the net proceeds towards the discharge of the outstanding obligations in accordance with applicable law.

1.3 Any enforcement of rights under this Clause shall be conducted in a commercially reasonable manner, and the Custodian shall account to the Client for any surplus proceeds remaining after satisfaction of the relevant obligations.

1.4 The rights set out in this Clause are in addition to, and not in substitution for, any other rights or remedies available to the Custodian under law or contract.

Negotiation Points

Custodians:

- Breadth of Digital Assets Covered: Including residual or quarantined digital assets ensures protection even where assets are temporarily restricted due to compliance actions.
- Right to Suspend Instructions: The ability to pause withdrawals of Digital Assets prevents asset flight before liabilities crystallize.
- Enforcement Without Court Delay: Given market volatility, custodians prefer contractual enforcement rights rather than waiting for court orders.
- Cumulative Remedies: Preserving lien rights alongside indemnities, set-off, and termination rights strengthens overall risk coverage.

Customers:

- Limiting the liability: Narrow the lien to *actual, due, and undisputed* amounts, excluding contingent or speculative claims.
- Notice and Cure Periods: Require prior written notice and a reasonable opportunity to cure before enforcement of lien.
- Commercial Reasonableness: Explicit standards for sale or disposal of digital assets help prevent distressed or opportunistic liquidation.
- Surplus Protection: Clear obligations on the custodian to return excess proceeds after satisfaction of liabilities.
- Segregation Clarity: Confirm that lien rights do not undermine asset segregation or beneficial ownership principles.

Digital Assets

KNOWLEDGE
CORNER

Privacy Coins

BYTE/WISE

Privacy Coins

Privacy Coins: Balancing Anonymity and Accountability

Blockchain technology is often associated with radical transparency. On most public blockchains, transactions and wallet balances are permanently recorded and visible to anyone, enabling trust without intermediaries and supporting fraud detection, asset tracing, and enforcement investigations. This transparency is a defining feature of mainstream cryptocurrencies like Bitcoin and Ether. Privacy coins, however, are a notable exception.

What are privacy coins?

Privacy coins are cryptocurrencies designed to enhance user anonymity and reduce transaction traceability. They function more like digital cash, while an initial transaction (such as acquiring the coins through an exchange) may be recorded, subsequent transfers are intentionally difficult to trace. Although exchanges typically conduct identity verification at the entry point, the built-in privacy features of these coins limit visibility into how funds move thereafter. However, privacy does not mean absolute anonymity. With advanced tools and techniques, investigators may still track activity, though doing so is far more complex than on transparent blockchains.

How do privacy coins work?

Privacy coins use cryptographic techniques to obscure transaction details. Common methods include stealth addresses (which generate a new address for each transaction), ring signatures (which mix multiple users' transactions together), and zero-knowledge proofs, which validate transactions without revealing underlying data. Different privacy coins adopt different combinations of these tools.

Key privacy coins

The most well-known privacy coins include Monero (XMR), Zcash (ZEC), and Dash (DASH). Monero offers privacy by default, Zcash allows users to choose between transparent and shielded transactions, and Dash provides optional privacy features layered onto a Bitcoin-derived model.

Use cases and concerns

Privacy coins are often associated with illicit activity, but evidence suggests most criminal actors still prefer Bitcoin due to its liquidity and ease of conversion into fiat currency. Legitimate use cases include protecting sensitive financial information, reducing exposure to hacking, and preserving financial autonomy in jurisdictions with heavy surveillance or capital controls.

Regulatory outlook

While privacy coins are legal in some jurisdictions, they are facing increasing regulatory scrutiny. Several countries have restricted or banned them, and many exchanges have delisted privacy-focused tokens. The future of privacy coins will likely depend on whether regulators and market participants can strike a workable balance between privacy, transparency, and financial integrity.

BYTE/WISE

PART FOUR

DARK PATTERNS

JAN 2026

BYTE/WISE

Drip Pricing

The CCPA Guidelines for Prevention and Regulation of Dark Patterns, 2023 define “Drip Pricing” as the practice of advertising a product or service at an initial price while revealing additional mandatory charges only at later stages of the transaction, such as during checkout or payment. These charges may include convenience fees, platform fees, handling charges, or other non-optional costs that materially alter the final price.

Drip pricing is explicitly recognized as a dark pattern and an unfair trade practice under the Consumer Protection Act, 2019, as it misleads users about the true cost of a transaction and distorts price comparison and decision-making.

Common manifestations include:

- ❌ Displaying a low headline price while adding mandatory fees at checkout
- ❌ Revealing platform, convenience, or handling charges only at the payment stage
- ❌ Separating unavoidable taxes or fees from the advertised price without clear upfront disclosure
- ❌ Incrementally adding charges across multiple screens to normalize the final price

Drip pricing is problematic not because businesses may charge legitimate fees, but because staggered disclosure undermines price transparency, impairs informed consent, and exploits behavioral biases by anchoring users to an artificially low initial price.

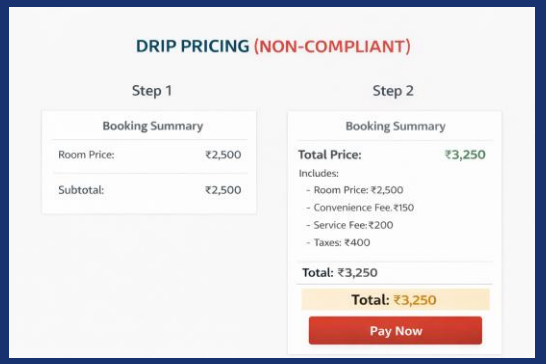

Better practice

- ✅ Disclose the total payable price upfront, including all mandatory charges
- ✅ Clearly distinguish optional add-ons from unavoidable fees
- ✅ Ensure headline prices reflect the minimum amount a user must pay to complete the transaction
- ✅ Present taxes and statutory charges transparently and consistently across the user journey

Tip for businesses

Audit pricing and checkout flows to identify where mandatory charges are introduced late in the transaction. Rework interfaces so that users see the full, final price before committing to a purchase, and document pricing logic to demonstrate that cost disclosures are clear, timely, and non-misleading.

JAN 2026

COMING SOON

FEB ISSUE

ISSUE EDITORS

Partner: Naresh Pareek
Consultant: Shravan Kalluri

CONTRIBUTORS

Abhishek Nair
Jash Doshi

PUBLISHING SUPPORT

Vivek Yadav

309-10 Madhava,
C5, E Block, BKC,
Bandra East,
Mumbai 400 051