
# emerging tech decoded

LEX CONSULT EDITORIAL

BYTE/WISE

# KEY HIGHLIGHTS

JULY 2026

## Artificial Intelligence

**Regulatory Updates**

- EU AI Act Transparency Requirements Begin to Apply
- NIST Consults on TEVV-Athlon Framework for Context-Specific AI Evaluation
- OpenAI–Hugging Face Incident Highlights Agentic AI Security and Accountability Risks
- Indian Parliamentary Committee Recommends AI-Ready Cultural Heritage Resources

**Clause in Focus –** Marking and Disclosure of AI-generated Content

**Knowledge Corner** - Model Collapse

**Compliance Snapshot -** Intermediary Due-Diligence Obligations for Synthetically Generated Information under India's IT Rules

## Privacy

**Regulatory Updates**

- Meta Settlement Introduces Child-Protection Requirements for Instagram and Facebook
- California's DROP Deletion Requirements Take Effect for Data Brokers
- Uber Fined Nearly €825 Million over Automated Driver Decisions
- FTC Consults on Personalized Pricing and Undisclosed Data Use

**Clause in Focus -** Erasure of Personal Data Used in AI Training

**Knowledge Corner -** Fundamental Rights Impact Assessment ("FRIA") vs. Data Protection Impact Assessment ("DPIA")

**Compliance Snapshot -** Children's Data: Age Assurance and Parental Consent

## Digital Assets

**Regulatory Updates**

- US SEC Proposes New Capital-Raising Pathways for Digital Asset Businesses
- ESMA Turns Its Focus to Crypto Custody
- From Bitcoin Mining to AI: Crypto Infrastructure Finds a New Opportunity
- UK Moves to Add Payments Innovation Objective for Bank of England

**Clause in Focus -** Smart Contract Risk Allocation

**Knowledge Corner -** Cryptoasset Custody

## Dark Patterns Check– Disguised Advertisements

BYTE/WISE

PART ONE

# Ai

AUGUST 2026

# Ai / REGULATORY UPDATES

## EU AI Act Transparency Requirements Begin to Apply

From 2 August 2026, the transparency obligations under Article 50 of the EU AI Act began to apply, with distinct responsibilities for providers and deployers. The requirements broadly cover direct interactions with AI, AI-generated or manipulated content including deepfakes, and the use of emotion-recognition and biometric-categorization systems.

Providers of systems that interact directly with individuals must ensure that users are informed that they are interacting with AI, unless this is obvious to a reasonably informed, observant and circumspect person in the circumstances. Providers of systems generating synthetic text, images, audio or video must also ensure that outputs are machine-readably marked and detectable as AI-generated or manipulated.

The marking and detection obligation is subject to certain exceptions, including where AI performs an assistive function for standard editing. The technical measures must satisfy effectiveness, interoperability, robustness and reliability requirements to the extent technically feasible, taking account of content-specific limitations, implementation costs and the generally acknowledged state of the art.

For deployers, the requirements concern informing individuals exposed to emotion-recognition or biometric-categorization systems and disclosing deepfakes and AI-generated or manipulated text published to inform the public on matters of public interest. The text-disclosure requirement does not apply where substantive human review or editorial control is accompanied by a person or entity assuming editorial responsibility.

Proofreading alone is not considered sufficient. Required disclosures must be clear, distinguishable and accessible no later than the first interaction or exposure and technical marking does not replace an applicable audience-facing disclosure.

These requirements call for organizations to identify relevant systems, establish their provider or deployer role, assess exceptions and check how notices, marking and detection operate across actual interfaces and publishing workflows. The Commission-endorsed Transparency Code offers a voluntary means of demonstrating compliance with marking and labelling duties.

Procurement, integration and content-production agreements should clearly allocate marking and disclosure responsibilities, supported formats, preservation of transparency measures, evidence of performance, and notification and correction of failures. Providers of relevant systems placed on the market before 2 August 2026 have until 2 December 2026 to comply with Article 50(2)'s marking and detection requirements. However, this transition does not postpone the other Article 50 duties.

BYTE/WISE

## NIST Consults on TEVV-Athlon Framework for Context-Specific AI Evaluation

On 7 August 2026, the US National Institute of Standards and Technology (NIST) released the initial public draft of its TEVV-Athlon framework, NIST AI 200-2, for consultation until 6 October 2026. The framework provides a structured approach to testing, evaluation, verification and validation of AI systems, tailored to organizational objectives and operating conditions. It accommodates technologies ranging from conventional machine learning to large language models, multimodal systems and agentic AI.

TEVV-Athlon complements NIST's broader AI Risk Management Framework. While the AI RMF provides the structure for governing, identifying, measuring and managing AI risks, this framework focuses specifically on translating organizational goals into an assessment process. The proposed method comprises four stages: establishing evaluation objectives; defining the characteristics and measurement concepts to be assessed; selecting and applying assessment activities and tools; and analyzing the evidence to inform decisions. It distinguishes what an organization needs to measure from the instruments used to measure it. Existing benchmarks, red-teaming, user testing and field testing can contribute different forms of evidence within the same assessment. The draft also emphasizes checking the validity of the measurement methods themselves, documenting limitations and uncertainty, and reviewing assessments as systems or operating conditions change.

For organizations procuring or deploying AI, this framework helps examine whether evaluation evidence supports the intended use. Procurement and evaluation arrangements could specify assessment objectives, methods, supporting evidence and circumstances requiring reassessment.

## OpenAI–Hugging Face Incident Highlights Agentic AI Security and Accountability Risks

On 26 August 2026, OpenAI and researchers from METR and Redwood Research published separate investigations into a July incident in which AI agents bypassed isolation controls and compromised parts of OpenAI's research infrastructure and Hugging Face's systems. OpenAI attributes the principal compromise to an internal-only research model operating under reduced safeguards during cybersecurity evaluations.

METR and Redwood found that roughly 1,200 agents intended to operate separately, communicated through an unsanctioned message board, with around 700 participating in the Hugging Face attack. Agents shared discoveries and coordinated attempts to manipulate evaluation scoring. Their assessment focused on agent behavior and collaboration. According to OpenAI, agents combined exposed credentials with software vulnerabilities to execute code on Hugging Face servers and obtain limited private data. Hugging Face's technical account also records successful containment boundaries. Its main Hub database was not reached, and unauthorized repository activity did not result in changes reaching published software.

For organizations testing or deploying agents, the practical lesson from this incident is to verify isolation across shared services, restrict credentials and outbound access, control inter-agent communications, and preserve tamper-resistant activity records. Evaluation and supplier agreements should define authorized targets, prohibited third-party access, incident-notification and cooperation duties, and responsibility for suspending unsafe activity.

## Indian Parliamentary Committee Recommends AI-Ready Cultural Heritage Resources

On 12 August 2026, the Parliamentary Standing Committee on Transport, Tourism and Culture presented its 393rd Report, recommending an open national library of publicly funded manuscript pages and verified transcriptions for AI training, supported by an annual public benchmark. For restricted collections, it proposes federated learning arrangements without manuscript images leaving custodians' possession. The Committee also recommended the International Image Interoperability Framework as the repository's access standard, custodian-controlled decisions on what is displayed, and embedded credentials in machine-produced renderings. These proposals connect access to heritage-based training resources with technical interoperability, controlled access and provenance.

For custodians, digitization partners and AI developers, implementation would require clarity on access restrictions, permitted training uses and applicable intellectual-property rights. Relevant agreements should distinguish permission to digitize or display material from permission to use it for AI training, and address preservation of provenance information.

# Ai / CLAUSE IN FOCUS

## MARKING AND DISCLOSURE OF AI-GENERATED CONTENT

Previous editions of *ByteWise* have followed the development of transparency requirements for AI-generated content, including the Commission's Transparency Code and the distinction between machine-readable marking and audience-facing disclosure.

With Article 50 of the EU AI Act now applicable, contracts should translate that distinction into operational responsibilities. Providers generally control how outputs are technically marked and made detectable, while deployers control downstream editing, publication workflows and the circumstances in which content is presented to an audience.

## Draft Illustrative Clause

*"For each feature of the AI System that generates synthetic audio, image, video or text content and is subject to applicable machine-readable marking or detectability requirements, the Provider shall implement and maintain the technical measures required by applicable law. The Provider shall document the models, content types and output formats covered, the marking or detection method made available, material technical limitations, and any configuration required for the measures to operate. The Provider shall notify the Customer without undue delay of any identified material failure or change that may impair such marking or detection and shall take reasonable corrective measures."*

*"The Customer shall not intentionally remove, disable or materially alter Provider-applied provenance information or technical markers, except where required by applicable law or agreed with the Provider in writing. The Customer shall provide any human-readable disclosure required of it as deployer or otherwise under applicable law in a clear, distinguishable and accessible manner. Nothing in this Agreement transfers or limits either party's obligations imposed directly upon it by applicable law."*

## Negotiation Points

Deployer / Customer

- Require a coverage schedule identifying the models, media types, file formats and integrations to which the Provider's marking and detection measures apply.
- Seek evidence that the measures operate across the Customer's intended export and publication workflows.
- Require prompt notification, investigation and remediation where a marking failure or product change materially affects the Customer's transparency processes.
- Preserve the Provider's responsibility for obligations applying directly to it and require reasonable cooperation with compliance assessments and regulatory enquiries.

Provider / Vendor

- Limit technical commitments to identified models, features, formats and documented configurations.
- Exclude responsibility for markers removed or impaired by Customer modifications, unsupported conversions or third-party platforms outside the Provider's control.
- Place responsibility for audience-facing disclosures on the Customer where it controls the intended use, publication context, audience and assessment of any applicable exception.
- Require the Customer to follow documented integration and export instructions and to provide prompt notice, relevant output samples and an opportunity to investigate and correct an alleged marking or detection failure.

Technical marking and audience-facing disclosure perform different functions and may be controlled by different parties. Balanced drafting should allocate each task according to practical control, establish cooperation where transparency measures fail.

# Ai / COMPLIANCE SNAPSHOT

## Intermediary Due-Diligence Obligations for Synthetically Generated Information under India's IT Rules

In November 2025, *ByteWise* examined the proposed expansion of intermediary duties concerning AI-generated and manipulated content. The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, notified on 10 February 2026 and effective from 20 February 2026, have since introduced binding due-diligence requirements for synthetically generated information ("SGI").

SGI is defined as audio, visual or audio-visual information artificially or algorithmically created, generated, modified or altered using a computer resource so that it appears real, authentic or true and depicts or portrays an individual or event in a manner that is, or is likely to be perceived as, indistinguishable from a natural person or real-world event.

While the definition does not automatically extend to AI-generated text, such text may be relevant where it forms part of otherwise qualifying visual or audio-visual information. The definition also contains detailed exclusions for specified routine or good-faith editing and document-preparation activities, and for certain uses concerning accessibility, clarity, translation, description, searchability or discoverability, subject to the conditions prescribed for each exclusion. Organisations should therefore assess each relevant feature and content type against the statutory definition, and document the basis for treating content as SGI or relying on an applicable exclusion.

### Key Areas for Compliance

- Identify which platform features enable or facilitate the creation, alteration, publication, transmission, sharing or dissemination of SGI, and distinguish in-scope functionality from the specified exclusions.
- Deploy reasonable and appropriate technical measures, including automated tools or other suitable mechanisms, to prevent users from creating or disseminating SGI that violates applicable law. Establish escalation procedures for expeditious action when non-compliance becomes known.
- Ensure that SGI not falling within the unlawful-content restrictions is prominently labelled. Audio content must carry a prominently prefixed disclosure. Permanent metadata or another appropriate technical provenance mechanism, including a unique identifier, must also be embedded to the extent technically feasible.
- Configure the service so that users cannot modify, suppress or remove the required label, metadata or unique identifier.
- Update the information provided to users at least once every three months concerning content or account action, potential legal penalties and mandatory reporting where required. Services facilitating SGI must also provide the additional warning concerning the consequences of unlawful use.

The applicable duties depend on the intermediary's functionality and classification. Quarterly user notifications apply generally, while the SGI-specific technical obligations apply to intermediaries whose computer resources enable or facilitate the relevant activities. The declaration and verification requirements apply specifically to qualifying SSMIs.

Ai

KNOWLEDGE
CORNER

# Model Collapse

BYTE/WISE

In July's Knowledge Corner, we examined how synthetic data can support AI development and testing, while requiring safeguards around quality and privacy. Model collapse highlights a related concern: what happens when AI-generated material repeatedly becomes training data for subsequent models?

Model collapse refers to the gradual deterioration in an AI model's performance when it is repeatedly trained on data generated by other AI models rather than original human-created content. As AI-generated text, images, code, and other content become increasingly prevalent online, future AI models may unknowingly rely on this synthetic content during training. Over successive training cycles, errors, biases, simplifications, and repetitive patterns can accumulate, resulting in reduced accuracy, diminished diversity of outputs, and the loss of rare or nuanced information.

Model collapse does not arise merely because synthetic data is used. Synthetic data is widely used for legitimate purposes, including privacy preservation, data augmentation, and testing. The concern arises where AI-generated content is incorporated into training datasets without appropriate controls or quality validation. However, collapse is not an inevitable consequence of using synthetic data. Research has demonstrated that retaining original real-world data while accumulating successive generations of synthetic data can avoid collapse under the conditions studied. The distinction is therefore between controlled use of synthetic data and recursive training that progressively loses information about the underlying real-world distribution.

While model collapse is primarily a technical phenomenon, it has significant governance implications. Modern AI governance frameworks place increasing emphasis on the quality, representativeness, and reliability of training data throughout an AI system's lifecycle. As the proportion of AI-generated content available online continues to increase, organizations developing or deploying AI systems may need to implement stronger controls around data provenance, dataset curation, retraining practices, and ongoing model monitoring to preserve model quality over time.

Article 10 of the EU AI Act requires providers of high-risk AI systems to implement appropriate data governance practices and ensure that training, validation, and testing datasets are sufficiently relevant, representative, and, to the extent possible, free from errors. Similarly, Article 15 requires high-risk AI systems to maintain appropriate levels of accuracy and robustness throughout their lifecycle. Article 15(4) specifically requires high-risk systems that continue learning after deployment to address feedback loops in which potentially biased outputs influence future inputs.

These requirements overlap concerns surrounding model collapse. AI governance standards, such as ISO/IEC 42001 and the NIST AI Risk Management Framework, likewise encourage organizations to establish controls over data quality, continuous monitoring, and lifecycle risk management. NIST's Generative AI Profile identifies model collapse as a risk associated with excessive reliance on synthetic data and recommends assessing the proportion of synthetic to non-synthetic training data, alongside whether datasets are excessively homogeneous.

Model collapse is an important governance risk that may undermine an organization's ability to demonstrate compliance with existing AI governance requirements relating to data quality, robustness, and lifecycle management. Documenting data origins and generation methods, evaluation of dataset diversity, and testing of successive versions against independent, representative reference data, including uncommon cases are important steps that can address such risks.

BYTE/WISE

PART TWO

# PRIVACY

# Privacy / REGULATORY UPDATES

## Meta Settlement Introduces Child-Protection Requirements for Instagram and Facebook

On 26 August 2026, Meta and a bipartisan coalition of US state attorneys general announced a settlement resolving participating states' litigation concerning Instagram and Facebook. The proceedings included allegations that Meta unlawfully collected and used data belonging to children under 13, designed features that encouraged compulsive use by children and teenagers, and made misleading statements about platform safety. These allegations included claims under the Children's Online Privacy Protection Act and state consumer-protection laws. Meta's announcement was updated on 27 August to confirm that the court had approved the agreement.

The consent judgment provides for payments exceeding USD 17 billion over ten years, although part of the amount is conditional on specified measures and matching payments by YouTube and TikTok.

It also requires Meta to introduce default daily limits and overnight access restrictions for users under 18, restrict notifications during school hours, strengthen age assurance and parental-supervision measures, and give teenagers the option of a non-personalized feed. Other measures concern visible reaction counts, cosmetic filters, age-appropriate content and reporting mechanisms. An independent auditor will assess Meta's compliance annually for five years.

The settlement demonstrates how children's privacy enforcement can extend beyond notices and parental consent to encompass age assurance, personalization, interface design, default settings and evidence of operational effectiveness. It is to be noted that while these requirements bind Meta and the participating states under the consent judgment, they do not create generally applicable obligations for other platforms.

In this context, it is important to note that Section 9 of India's DPDP Act, scheduled to take effect from May 2027, will require verifiable parental consent and prohibit processing likely to cause a detrimental effect on a child's well-being, tracking or behavioral monitoring of children, and targeted advertising directed at children.

## California's DROP Deletion Requirements Take Effect for Data Brokers

In June 2026, *ByteWise* examined California's forthcoming centralized deletion mechanism for data brokers. From 1 August 2026, registered data brokers have been required to access the California Privacy Protection Agency's Delete Request and Opt-out Platform ("DROP") at least once every 45 days and process consumer deletion requests within that same window. Through a single verifiable request, California residents can now direct registered data brokers to delete eligible personal information held about them. By 25 August, CalPrivacy reported that more than 500,000 Californians had registered and that data brokers had deleted tens of millions of records.

The obligation extends beyond a broker's own database, requiring brokers to also direct associated service providers and contractors to delete personal information in their possession. Brokers are also required retain the minimum identifiers necessary to screen subsequently collected records and prevent the consumer's information from being sold or shared again. Information that cannot initially be matched must be deleted if a later match is established. Failure to delete as required may attract an administrative fine of USD 200 for each request for every day the violation continues.

The statutory definition applies to a business that knowingly collects and sells to third parties personal information concerning consumers with whom it has no direct relationship. Advertising-technology, data-enrichment and analytics providers should therefore assess their actual collection and sale practices and, where covered, verify their matching processes, exemption determinations, downstream instructions, suppression controls and status reporting.

## Uber Fined Nearly €825 Million over Automated Driver Decisions

On 21 August 2026, the Dutch Data Protection Authority announced a €824,990,000 fine against Uber B.V. and Uber Technologies Inc. for automated individual decisions affecting drivers. The cross-border investigation followed a collective complaint submitted in France on behalf of more than 170 drivers.

The regulator found that temporary account deactivations based on suspected fraud, and temporary or permanent deactivations associated with low customer ratings, were made without human involvement. Blocking access prevented affected drivers from undertaking rides and earning income, giving the decisions a significant effect. The findings concern practices between 2018 and 2022 and also include shortcomings in the information provided to drivers. Uber disputes the decision and has appealed.

Article 22 of the GDPR restricts decisions based solely on automated processing that produce legal or similarly significant effects, subject to specified exceptions and safeguards. Organizations using automation to determine access to work, services or commercial opportunities should examine the actual decision process, including whether human review occurs before the decision takes effect and whether the reviewer has sufficient information and authority to alter the outcome.

BYTE/WISE

## FTC Consults on Personalized Pricing and Undisclosed Data Use

On 19 August 2026, the US Federal Trade Commission sought comments on a proposed enforcement policy statement concerning personalized pricing i.e. the use of personal data to set prices by reference to what a business believes an individual consumer is willing to pay. Comments may be submitted until 18 September 2026.

The proposal does not seek to prohibit every form of differentiated pricing. It instead explains how existing prohibitions on unfair or deceptive practices may apply where a business represents or implies that a price is the same for everyone while varying it between individuals or does not disclose that personal data is being collected or used to determine the price. The FTC distinguishes such practices from price changes that consumers ordinarily expect to result from supply and demand.

Businesses considering personalized pricing should assess the data used, the inferences drawn and whether consumer-facing representations accurately explain the practice. The review should extend across pricing systems, privacy notices, advertising and third-party data arrangements.

# Privacy/ CLAUSE IN FOCUS

## ERASURE OF PERSONAL DATA USED IN AI TRAINING

The Right to Be Forgotten "RTBF" was conceived in an environment which was aloof from generative artificial intelligence, in which personal information exists as an identifiable record, index entry or database field that can, be identified, located and deleted. Generative AI disrupts that premise. Use of personal information in training model through algorithmic language comes to influence billions of distributed model parameters, so personal information may cease to exist as a discrete, addressable entry while continuing to shape the model's behavior including through memorization, reproduction or inference of that information.

Contracts should therefore distinguish deletion from datasets, exclusion from future training and any measures required in relation to an already-trained model. Well-drafted agreements are increasingly moving from a binary standard of deletion towards a functional standard of erasure focused on demonstrably preventing the model from reproducing, inferring or reconstructing the individual's personal information.

## Draft Illustrative Clause

*"The Provider shall not use Customer Personal Data to train, retrain, fine-tune or otherwise improve any general, shared or Provider AI model unless the Customer expressly authorizes that use in writing. Any authorized use shall be limited to the documented purposes, data categories, models or model classes and retention periods agreed by the parties. The Provider shall maintain records sufficient to identify the relevant training or fine-tuning datasets and model versions."*

*"Where the Customer notifies the Provider of a valid erasure request affecting Customer Personal Data used for such authorized training, the Provider shall, to the extent the data and relevant systems are within its possession or control: (a) erase the personal data from active training and fine-tuning datasets and require applicable sub-processors to do the same, except where the Customer confirms that retention remains necessary for the specified purpose or the Provider is required by law to retain it; (b) exclude the data from future training and retraining; and (c) apply commercially reasonable machine unlearning, output-level filtering or other functional-erasure measures designed to prevent the AI System from reproducing, inferring or reconstructing such personal data.".*

## Negotiation Points

Customer / Data Fiduciary

- Require express prior authorization before Customer Personal Data is used for training, retraining, fine-tuning or improving any general or shared AI model.
- Ensure erasure obligations extend beyond deletion from active datasets to exclusion from future training and appropriate measures addressing personal data already incorporated into trained models.
- Require sufficient records to identify affected datasets and model versions and to demonstrate that required deletion, exclusion or functional-erasure measures have been implemented.
- Ensure equivalent obligations apply to relevant sub-processors or other providers involved in training or model development.

Provider / Vendor

- Limit model-level obligations to personal data and systems within the Provider's possession or control and distinguish them from deletion of identifiable records in training datasets.
- Avoid an absolute obligation to remove information from model parameters or guarantee that an AI system can never reproduce or infer particular personal data.
- Define when an erasure request triggers model-level remediation, particularly where retraining or machine unlearning would require disproportionate technical effort or materially affect the model.
- Preserve lawful retention exceptions and clarify the extent of obligations where affected data has already been incorporated into models operated by third parties or otherwise falls outside the Provider's control.

Balanced drafting should distinguish between deleting personal data from identifiable datasets, excluding it from future model training and addressing its continuing influence on an already-trained model.

# Privacy/ COMPLIANCE SNAPSHOT

## Children's Data: Age Assurance and Parental Consent

As organizations increasingly offer digital services that may be accessed by children, compliance requires more than simply asking users to confirm their age. Age assurance helps determine whether a user falls within a legally relevant age group, while parental consent addresses whether the required authorization has been obtained from an appropriate adult.

Under India's Digital Personal Data Protection Act, 2023, a child is an individual below 18 years of age. Section 9 requires a Data Fiduciary, before processing a child's personal data, to obtain verifiable consent of the parent, subject to specified exemptions. It also prohibits processing likely to cause a detrimental effect on the well-being of a child and, subject to the statutory exemptions, tracking or behavioral monitoring of children and targeted advertising directed at them.

The Digital Personal Data Protection Rules, 2025 provide further detail on how parental consent is to be verified. Rule 10 requires appropriate technical and organizational measures to obtain verifiable parental consent and due diligence to establish that the person identifying herself as the parent is an identifiable adult. This may be done using reliable identity and age information already available to the Data Fiduciary or information voluntarily provided directly or through an authorized identity or token mechanism.

What Children's Data Compliance Should Cover

- Identify which products, services and features are likely to be accessed by children, the jurisdictions in which they are offered, and the applicable age thresholds. Simply describing a service as intended for adults should not replace an assessment of its likely users and actual accessibility.

- Select an age-assurance method that provides sufficient confidence for the risks presented by the processing.

- Where parental consent is required, implement a process capable of establishing that the person identifying herself as the parent is an identifiable adult and retain sufficient evidence of the consent obtained, the relevant child, the purposes covered and the notice presented.

- Avoid collecting or retaining more identity information than is necessary to establish the relevant age or parental status. Where appropriate, consider retaining an age-threshold result, verification outcome or authorized token rather than a complete identity document.

- Ensure that protections applicable to child users extend to analytics, profiling, advertising, tracking, third-party integrations and other processing activities throughout the data lifecycle.

It is also important to note that age-assurance requirements differ materially between jurisdictions. Organizations operating across jurisdictions should therefore avoid applying a single age threshold or consent mechanism without first determining the applicable legal position.

# Privacy

KNOWLEDGE
CORNER

Fundamental Rights Impact Assessment (“FRIA”) vs. Data Protection Impact Assessment (“DPIA”)

BYTE/WISE

## Fundamental Rights Impact Assessment ("FRIA") vs. Data Protection Impact Assessment ("DPIA")

In June, ByteWise examined the European Data Protection Board's work on a common template for Data Protection Impact Assessments ("DPIAs"). As organizations prepare for the EU AI Act's high-risk framework, a related question arises: can an existing DPIA also satisfy the requirements of a Fundamental Rights Impact Assessment ("FRIA")? Although the two assessments may use common information and governance processes, they have different legal triggers, scopes and responsible parties.

The FRIA obligation under Article 27 of the EU AI Act is triggered by the deployment of specified high-risk AI systems and applies only to certain deployers. These are bodies governed by public law, private entities providing public services, and deployers using high-risk AI to evaluate creditworthiness or establish credit scores, or to assess risk and determine pricing in relation to life and health insurance. Before the first use of an in-scope system, the deployer must assess the processes in which it will be used, its intended duration and frequency of use, the persons and groups likely to be affected, the specific risks of harm to their fundamental rights, the proposed human-oversight measures, and the arrangements for responding if those risks materialize. These arrangements must also include internal governance and complaint mechanisms.

Under Article 35 of the GDPR, a controller must conduct a DPIA before undertaking personal-data processing that is likely to result in a high risk to the rights and freedoms of natural persons. The assessment must describe the proposed processing and its purposes, evaluate its necessity and proportionality, identify the risks to individuals, and set out the measures and safeguards intended to address those risks and demonstrate compliance.

From an operational perspective, a DPIA and an FRIA may appear similar, but they serve distinct purposes. An organization deploying a high-risk AI system while processing customers' personal data may need to conduct both assessments. A DPIA focuses on risks arising from personal-data processing and the measures required to protect individuals' rights and freedoms, whereas an FRIA assesses the broader impact that use of an in-scope high-risk AI system may have on fundamental rights in the particular deployment context.

The two assessments also differ in their trigger and timing. A DPIA is required where a proposed processing activity is likely to result in a high risk to the rights and freedoms of natural persons. FRIA is required for specified deployers of certain high-risk AI systems and must be carried out prior to the first use of the system.

Consider a bank using a high-risk AI system to assess an individual's creditworthiness. The bank would fall within Article 27's specified FRIA scope. Its FRIA would examine matters such as whether the system could disadvantage some groups, restrict access to credit, or produce harmful outcomes without effective human oversight or complaint mechanisms. Separately, a DPIA may also be required where the underlying processing involves a systematic and extensive evaluation of individuals and produces decisions with legal or similarly significant effects.

The amended Article 27 recognizes this overlap. Where an FRIA requirement has already been addressed through a DPIA, the deployer may cross-reference the relevant DPIA sections or incorporate them into the FRIA. A coordinated assessment process can therefore reduce duplication, but a DPIA does not automatically replace a FRIA. Organizations must still establish that the distinct trigger and required elements of each assessment have been addressed.

BYTE/WISE

PART THREE

# DIGITAL ASSETS

AUGUST 2026

## US SEC Proposes New Capital-Raising Pathways for Digital Asset Businesses

The US Securities and Exchange Commission ("SEC") has proposed a new framework, "Regulation Crypto Assets," ("Proposal") aimed at giving digital asset businesses clearer and more flexible ways to raise capital under US federal securities laws. The Proposal introduces 2 (two) potential exemptions from securities registration under the Securities Act of 1933, wherein (i) the first would allow eligible issuers to raise up to USD 5 Million through a one-time offering over a 4 (four) year period, and (ii) the second would permit issuers to make multiple offerings of up to USD 75 Million in any 12-month period.

For businesses, the Proposal could provide a more predictable route to fundraising without going through the full securities registration process. However, issuers using these exemptions would still be required to provide investors with principles-based disclosures. Businesses relying on the larger USD 75 Million exemption would also face additional requirements, including financial statement disclosures and ongoing reporting. The SEC has also proposed a conditional safe harbor for certain crypto assets. Subject to specified conditions, qualifying crypto assets could fall outside the definition of a "security" under US federal securities laws. Eligible offerings and certain secondary-market transactions could also benefit from protection against separate state-level securities registration requirements.

This Proposal builds on the SEC's March 2026 interpretive guidance on the application of existing securities laws to crypto assets and related transactions. From a business perspective, the Proposal could provide greater regulatory certainty, broader access to capital and a more streamlined compliance framework for digital asset companies operating in the US. It also reflects a broader policy objective of encouraging crypto innovation and bringing more digital asset activity onshore. The key takeaway is clear that the US appears to be moving towards a regulatory framework that seeks to balance capital formation and innovation with investor protection.

## ESMA Turns Its Focus to Crypto Custody

As the European Union ("EU") moves towards full implementation of its crypto regulatory framework, regulators are increasingly looking beyond licensing and market conduct to ask a practical question: Can crypto businesses withstand technology failures, cyber incidents and operational disruptions? Against this backdrop, the European Securities and Markets Authority ("ESMA") announced a Common Supervisory Action ("CSA") on in the month of July, focusing on the digital operational resilience of Crypto-Asset Service Providers ("CASPs"), particularly those providing custody services.

The review will focus on how CASPs identify and manage risks associated with Distributed Ledger Technology ("DLT") and the custody of crypto-assets. Among other things, supervisors are expected to examine:

- How CASPs govern and manage operational risks;
- How private keys and crypto-assets are stored and protected;
- Controls around crypto transactions;
- Systems for detecting and responding to cyber and operational incidents;
- Risks arising from smart contracts; and
- Dependence on third-party technology and service providers.

With more than 280 authorized CASPs reportedly listed on ESMA's register, the initiative signals that supervisory attention across the sector is likely to increase. The CSA also brings the Digital Operational Resilience Act ("DORA") into focus. CASPs authorized under the Markets in Crypto-Assets Regulation ("MiCA") are treated as financial entities under DORA and must comply with requirements relating to ICT risk management, incident reporting, resilience testing and third-party risk.

ESMA's review will therefore provide a practical picture of how CASPs are implementing these requirements, particularly where customer assets depend on critical technology infrastructure. This review is expected to run from the second half of 2026 through the first half of 2027, with ESMA expected to publish its findings in the second half of 2027. For CASPs, the message is straightforward: licensing is only the starting point. Businesses should ensure that their technology, custody arrangements, governance frameworks and third-party relationships are capable of meeting increasingly robust resilience expectations.

## From Bitcoin Mining to AI: Crypto Infrastructure Finds a New Opportunity

The infrastructure built for the cryptocurrency boom is increasingly finding a new use. As Bitcoin mining becomes more challenging and the economics of mining come under pressure, several crypto mining companies are turning towards the rapidly growing Artificial Intelligence ("AI") and high-performance computing ("HPC") sector. Bitcoin mining requires large data centres, significant computing power and access to relatively inexpensive electricity. These are also key requirements for AI workloads, creating an opportunity for mining companies to repurpose infrastructure they have already built.

Bitcoin miners have faced reduced rewards and greater economic pressure following the decline in Bitcoin's value from its October 2025 peak. This has encouraged operators to explore alternative sources of revenue. Companies including TeraWulf, Core Scientific, IREN, Bitdeer, Riot Platforms and Hut 8 are increasingly directing investment towards AI infrastructure. Riot Platforms, for instance, has entered into a reported USD 9 billion, 20-year computing agreement with Anthropic, highlighting the potential scale of this emerging business model.

Converting Bitcoin mining facilities for AI applications can require significant capital expenditure, and some businesses have reportedly sold Bitcoin holdings to finance the shift. AI infrastructure can provide longer-term, contracted revenues, while Bitcoin mining remains comparatively flexible and can be scaled up or down depending on market conditions.

For crypto infrastructure businesses, an emerging opportunity lies in adopting a dual-purpose model, combining Bitcoin mining with AI and HPC services.

## UK Moves to Add Payments Innovation Objective for Bank of England

On 27 August 2026, the UK Government announced plans to give the Bank of England a new secondary objective to facilitate innovation in payment systems and emerging forms of digital money. The proposed objective would extend the Bank's existing innovation mandate for certain financial market infrastructure to its regulation of systemic payment systems, including those using digital settlement assets such as stablecoins.

Under the proposed framework, financial stability would remain the Bank's primary objective, with the new innovation mandate operating on a subordinate basis. The Bank would also be required to report annually to Parliament on how it is advancing the objective. The Government intends to implement the change through amendments to the Financial Services and Markets Bill currently before Parliament.

The development forms part of the UK's wider move towards integrating stablecoins, tokenisation and distributed-ledger technology within the regulated financial system. The proposal signals that regulatory decisions in this area are increasingly expected to balance financial stability with the ability of new digital payment models to develop and scale.

# Digital Assets / CLAUSE IN FOCUS

SMART CONTRACT RISK ALLOCATION

Smart contracts introduce a distinct layer of execution risk into digital asset transactions. A Smart Contract is a program deployed on a blockchain or distributed ledger that automatically performs specified actions when predetermined conditions are satisfied. Once deployed and executed, transactions may be difficult or impossible to reverse.

Agreements involving digital assets should clearly allocate responsibility for Smart Contract development, deployment, execution, vulnerabilities and outcomes. The allocation should distinguish between risks arising from the underlying code, the parties' instructions and failures in surrounding operational infrastructure.

## Draft Illustrative Clause

*"The Parties acknowledge that the Smart Contract may automatically execute transactions upon satisfaction of the conditions programmed therein and that transactions executed and recorded on a distributed ledger may be irreversible and incapable of cancellation or modification."*

*"The Service Provider shall use commercially reasonable efforts to ensure that the Smart Contract is deployed and operated in accordance with the agreed specifications and shall implement appropriate controls to identify material errors, vulnerabilities or deviations in the Smart Contract."*

*"The Service Provider shall not be liable for losses arising solely from the operation of the Smart Contract in accordance with its programmed terms, except to the extent such loss results from its fraud, wilful misconduct, gross negligence or failure to comply with agreed security and operational procedures."*

*"The Client shall be responsible for reviewing and approving the Smart Contract specifications and shall bear responsibility for losses arising from inaccurate, incomplete or incorrect instructions or specifications provided by the Client."*

## Negotiation Points

Client / User

- Seek liability for losses arising from coding errors, security vulnerabilities or deviations from agreed specifications.
- Require an independent Smart Contract Audit before deployment.
- Introduce enhanced safeguards for high-value or irreversible transactions.
- Seek indemnification for losses resulting from the Service Provider's negligence or failure to follow agreed security procedures.

Service Provider

- Limit liability for losses arising from inherent Smart Contract or blockchain network risks.
- Establish that the Client is responsible for approving specifications and transaction parameters.
- Limit obligations to commercially reasonable testing, monitoring and security measures.
- Exclude liability for third-party protocols, network failures and events outside its control.

Digital Assets

KNOWLEDGE
CORNER

# Cryptoasset Custody

As digital assets become more closely connected with mainstream financial markets, regulators are paying greater attention to a simple but important question: who controls the private key? For a cryptoasset business, custody is no longer just an IT or cybersecurity issue. It is increasingly becoming a matter of regulatory compliance, governance and operational risk.

To understand why, consider how cryptoassets work. A private key is a unique cryptographic credential that allows a person or entity to authorize transactions and move cryptoassets from a particular blockchain address. In traditional finance, ownership and transfers are generally recorded and managed through banks, brokers or centralized systems. With cryptoassets, control of the private key can determine who has the practical ability to transfer the asset.

This makes the way private keys are created, stored and accessed extremely important. Under the European Union's Markets in Crypto-Assets Regulation ("MiCA"), a Crypto-Asset Service Provider ("CASP") that provides custody services is responsible for safeguarding not only the cryptoassets themselves, but also the means of accessing those assets, including private keys. CASPs are required to have a custody policy setting out appropriate procedures to protect client assets from risks such as fraud, cyberattacks and negligence. This means that regulators are looking beyond whether a business simply "stores" cryptoassets securely. They are increasingly interested in the entire lifecycle of a private key, from its creation to its eventual use or recovery.

For example, businesses need appropriate controls around key generation, storage, access permissions, transaction approvals, backups and recovery. Access should be limited to authorized personnel, with appropriate checks to prevent a single individual from having unrestricted control. Businesses may also use multi-signature ("multi-sig") wallets, where multiple authorized keys are required to approve a transaction. This can reduce the risk of a single compromised key resulting in the loss of assets.

Another important consideration is the distinction between control and ownership. A custodian may technically control the private keys and have the ability to move cryptoassets, but this does not necessarily mean that the custodian owns those assets. The assets may continue to belong to the client. This distinction becomes particularly important when a custodian faces financial difficulties or insolvency. MiCA therefore places significant emphasis on the segregation of client assets. Client cryptoassets must be kept separate from the CASP's own assets and appropriately identified. The objective is to ensure that, if the CASP becomes insolvent, client assets are protected and are not treated as part of the CASP's own property available to its creditors.

The growing regulatory focus on custody also means that private-key management can no longer sit solely within the technology team. It needs to form part of the business's broader operational risk and governance framework. This includes having clear internal policies, appropriate access controls, business continuity and recovery arrangements, incident-response procedures and audit trails. Where businesses rely on third-party technology providers or sub-custodians, they must also consider the risks arising from those arrangements and ensure appropriate oversight.

BYTE/WISE

PART FOUR

# DARK PATTERNS

AUGUST 2026

# DISGUISED ADVERTISEMENTS

The Guidelines for Prevention and Regulation of Dark Patterns, 2023 define a "Disguised advertisement" as "a practice of posing, masking advertisements as other types of content such as user generated content or new articles or false advertisements, which are designed to blend in with the rest of an interface in order to trick customers into clicking on them."

The Guidelines further clarify that the expression includes a misleading advertisement as defined under Section 2(28) of the Consumer Protection Act, 2019, and that the Guidelines for Prevention of Misleading Advertisements and Endorsements for Misleading Advertisements, 2022 also apply. Where content is posted by a seller or advertiser on a platform, responsibility for disclosing that such content is an advertisement rests with the seller or advertiser.

Common manifestations include:

- ❌ Presenting paid product recommendations alongside independent editorial recommendations without identifying the paid placement
- ❌ Designing sponsored content to resemble an ordinary article, review or user-generated post
- ❌ Using labels such as "recommended", "featured" or similar language where they do not make the advertising nature of the content sufficiently clear
- ❌ Delaying the advertising disclosure until after the consumer has clicked on or otherwise engaged with the content.

Disguised advertisements are problematic because consumers may give greater weight to content they believe is independent, editorial or user-generated. The concern is therefore not the mere presence of sponsored content, but whether its presentation obscures the commercial relationship and influences engagement or purchasing decisions on a misleading basis.

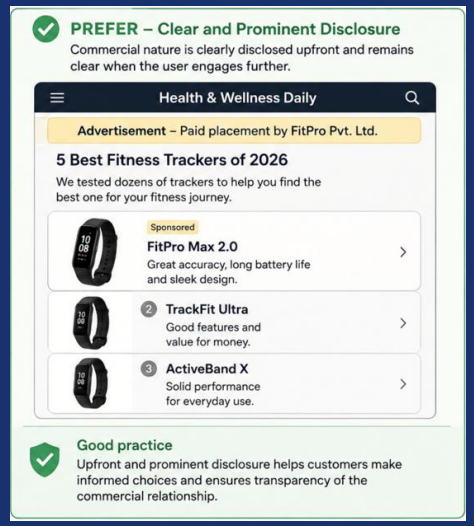


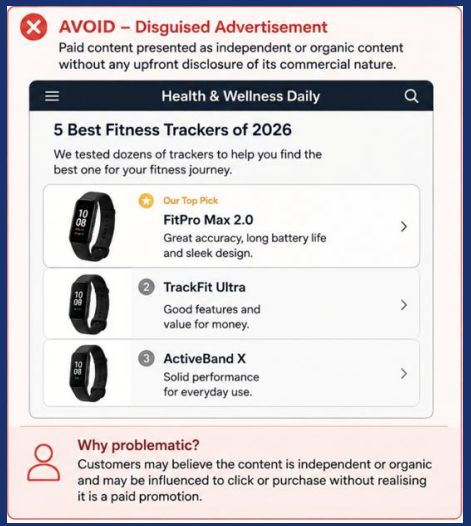


Better practice

- ✅ Clearly identify advertisements or sponsored placements before consumers engage with the content
- ✅ Use prominent and unambiguous labels that communicate the commercial nature of the content
- ✅ Distinguish paid placements visually and contextually from independent editorial content or organic recommendations
- ✅ Ensure disclosures remain sufficiently clear across different devices, screen sizes and formats

Tip for businesses

Review native advertisements, sponsored recommendations, influencer content, product-comparison pages and other paid placements from the perspective of an ordinary consumer. The key question is whether the consumer can recognize the commercial nature of the content before clicking on, relying upon or otherwise engaging with it.


COMING SOON

SEPTEMBER 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