Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report

Response to Consultation

ProcessUnity, Inc.

1. **Do you agree with the benefits and risks of AI adoption by financial institutions described in this report? Are there any substantive benefits or risks not covered?**

   ProcessUnity agrees with the risk description in Section 1.3 and suggests two points of additional emphasis.

   First, on unmonitored AI use: in ProcessUnity's experience, the largest and fastest-growing source of unmonitored AI in financial institutions is not employee-initiated tooling but AI capabilities delivered inside third-party products the institution has already procured. Providers increasingly add generative and agentic AI features to existing products between contract cycles, so an institution's AI footprint can expand materially without any new procurement event to trigger inventory capture or risk review.

   Second, on third-party dependencies: the relevant unit of analysis is frequently not the contracted vendor but the chain beneath it - the concentration and correlation risks described in Section 1.3 arise at these lower layers and are invisible to institutions whose third-party visibility stops at the direct counterparty.

2. **Are the sound practices sufficiently comprehensive and clear to enable financial institutions' responsible AI adoption?**

   ProcessUnity considers the sound practices broadly comprehensive and offers four clarifications that would improve implement ability.

   2.1 - AI inventories should expressly capture vendor-embedded AI (Sound Practice 3). The documentation practices described alongside Sound Practice 3 contemplate leveraging existing processes such as model inventories and third-party or systems registries, and the accompanying case study describes centralized AI inventories supported by dedicated software, including agent-level tracking. ProcessUnity supports this and requests that the final report make explicit that a complete AI use case inventory includes AI delivered through third-party products and services.

   2.2 - Continuous monitoring for material AI third-party relationships (Sound Practices 9 and 12). The Consultation Report observes in Section 1.3 that third-party providers may change the underlying models powering their AI systems without adequate notice, leaving

institutions unaware of changes that could affect performance or risk outcomes. ProcessUnity's 2026 research indicates that 64% of large organizations take an average of four months to complete a single vendor assessment and that 60% of organizations experience vendor response timelines of four months or longer. Where the dependency can change faster than a periodic assessment cycle completes, point-in-time assessment alone cannot provide the visibility the Consultation Report contemplates. The final report could helpfully clarify that, for material AI third-party relationships, the following are consistent with and support Sound Practice 12: (a) continuous monitoring of the third party's risk posture between formal assessments; (b) contractual change-notification terms covering underlying model changes, material capability changes, and subcontractor or model-provider substitutions; and (c) event-triggered reassessment upon such notifications or upon detection of performance or security anomalies.

2.3 - Standardized, attested, shareable assessments (Sound Practice 12). AI adoption concentrates assessment demand onto a small set of providers - the same cloud, hardware, and foundation-model firms identified in Section 1.3. Under a purely bilateral model, thousands of institutions independently issue bespoke questionnaires to the same providers, duplicating effort without improving risk insight; ProcessUnity's research indicates that 27% of vendors do not respond to assessments at all. The final report could recognize that use of industry-standard assessment instruments (for example, the Shared Assessments SIG or the Cloud Security Alliance Questionnaire) and of shared, attested assessment mechanisms - industry utilities, exchanges, or comparable arrangements - is consistent with Sound Practice 12, provided institutions retain accountability for their own risk decisions and supplement shared artifacts with institution-specific due diligence proportionate to materiality and risk.

2.4 - Dependency registers spanning the AI supply chain (Sound Practice 12). Consistent with the concentration risks described in Section 1.3, the final report could encourage institutions to maintain AI dependency registers spanning the full AI supply chain, including nth-party model and infrastructure dependencies, structured consistently with the FSB's 2023 toolkit on third-party risk management and oversight and with register-of-information approaches under frameworks such as the EU's Digital Operational Resilience Act. Standardized registers would also, over time, give authorities a practical empirical basis for monitoring system-level concentration, should they wish to draw on them.

3. Do the sound practices strike an appropriate balance between managing risks relating to all forms of AI, and addressing some of the risks relating to emerging and new complex forms of AI, such as GenAI and agentic AI?

ProcessUnity considers the balance appropriate and requests one confirmation relating to oversight at scale. The Consultation Report acknowledges in Section 1.3 that real-time human monitoring of agent decisions becomes impractical as use scales, and that effective monitoring and detection may in some cases require augmentation with another AI agent or other forms of AI; the executive summary similarly contemplates the use of AI tools within cyber and ICT risk management under Sound Practice 11.

The same logic applies to third-party risk management. An institution with thousands of third parties, each capable of shipping AI-driven changes between assessment cycles, cannot achieve the visibility contemplated by Sound Practice 12 through manual evidence review

alone. AI-assisted assessment - automated review of certifications, policies, and control evidence; validation of questionnaire responses; and continuous risk scoring across the portfolio - is how the coverage and cadence described under Question 2 become operationally achievable. Human accountability is retained throughout: analysts set thresholds, review exceptions, approve determinations, and remain responsible for risk decisions, with audit trails preserved.

ProcessUnity requests that the final report confirm that the use of AI to assist third-party assessment and continuous monitoring, subject to human accountability, defined escalation thresholds, and auditability, is consistent with Sound Practices 10 and 12. This would align the treatment of AI in the second line's third-party function with the report's existing treatment of AI in cyber defense and in agent monitoring.

4. **Are the sound practices sufficiently flexible to accommodate and address newer types of AI and responsible adoption over time?**

   Yes, provided two anchors are preserved. First, proportionality: in the third-party context, proportionality is operationalized through risk-based tiering, with assessment depth, evidence requirements, and monitoring cadence scaled to the materiality of the AI use case and the criticality of the relationship, consistent with Sound Practice 5. Second, interoperability: aligning terminology and expectations with existing third-party frameworks - the FSB's 2023 third-party toolkit, DORA registers of information, and comparable national guidance - allows a single, well-structured assessment and register to serve multiple regimes, and allows the sound practices to accommodate new AI forms without requiring a parallel, AI-specific assessment apparatus.

5. **Do the case studies in this report sufficiently highlight how different types of financial institutions can benefit from responsible AI adoption? Are there additional case studies for inclusion in the report? If so, please provide such case studies, particularly for nonbanks.**

   They are high level but we are not providing any additional use cases at this time as we need to discuss with the FSB and our clients as to whether they would allow an anonymized use case submittal.

6. **Do the case studies in this report provide actionable insights for financial institutions in their responsible AI adoption?**

   They are very high level but adequate.

7. **Are the definitions in the glossary clear and aligned with industry sound practices, including recent developments in AI?**

   ProcessUnity suggests two additions to support consistent inventory and assessment scoping across institutions:

   • Vendor-embedded AI: AI capabilities delivered within a third-party product or service procured by a financial institution, including capabilities added or materially changed by the provider after initial procurement.

• AI supply chain: The set of third-party and nth-party dependencies through which an AI system is developed, hosted, and operated, including cloud infrastructure, hardware, foundation models, data providers, and deployment platforms.

8. Are there any comments you would like to make on this report? Please fill in.

ProcessUnity appreciates the FSB's work on this consultation and supports publication of the final sound practices in October 2026. Because most financial institutions will adopt AI through third parties, the sound practices will succeed in practice largely on the strength of third-party visibility, assessment scalability, and supply-chain transparency. ProcessUnity respectfully requests the clarifications described above and would welcome continued dialogue with the FSB on the third-party dimensions of responsible AI adoption.