ABA MEMBER INSTITUTION QUESTIONNAIRE

AI-Assisted Vulnerability Detection & Remediation

Assessment Questions for Core Service Providers

Prepared by the American Bankers Association (ABA) Cybersecurity Working Group · June 2026

PURPOSE & INSTRUCTIONS

This questionnaire helps banks assess core service providers' preparedness for risks arising from frontier artificial intelligence (AI) models (e.g., Anthropic Claude Mythos Preview/Project Glasswing) that can identify and exploit software vulnerabilities.[1] ABA member banks may use this questionnaire to request information directly from core service providers. It is organized into 11 sections and references existing frameworks and standards; consolidated questions include labeled sub-items for reference. **ABA will not collect responses**; that is the responsibility of banks in alignment with their agreements with core service providers. Given the sensitivity of some questions, responses may be provided by service providers under non-disclosure agreements (NDAs) and appropriate restrictions. ABA encourages core service providers to answer these questions in sufficient detail for banks to assess risk and meet cybersecurity and third-party risk management expectations and align to recognized frameworks, standards and regulatory requirements.

INDEPENDENT REVIEW AND OVERSIGHT STATEMENT

Internal Review Protocol

Information Security and Enterprise Risk Management independently review all questionnaire responses as a formal component of the Bank's second-line third-party risk oversight.

Accountability and Governance

This validation ensures risk assessments align with internal safety and soundness standards while meeting regulatory expectations for institutional accountability.

Oversight and Quality Assurance Statement

Independent Review

Information Security and Enterprise Risk Management independently review all responses as part of second-line oversight to ensure institutional accountability.

Response Quality & Remediation

Submissions must be specific. Vague or boilerplate responses may trigger additional due diligence, a risk rating reassessment, or formal contractual remediation.

[1] In April 2026, ABA staff worked with cybersecurity, third-party risk management, and AI professionals from ABA member companies and partner associations to develop this common questionnaire for core service providers. ABA also vetted the draft with several leading core service providers and incorporated feedback, as appropriate. It is intended for use as a third-party risk management tool and does not constitute legal advice. Regulatory citations are provided for reference; institutions should consult applicable guidance and counsel regarding specific compliance obligations. The US Treasury Department published in February 2026 the AI lexicon which includes helpful definitions for AI terms.

PROVIDER INFORMATION DISCLOSURE FRAMEWORK

To ensure a secure, consistent, and accurate assessment process, providers are encouraged to respond to this questionnaire in alignment with a formalized, tiered disclosure model. This framework balances the Bank's need for robust risk verification with the provider's mandate to safeguard operational environment integrity:

- Standardized Framework Alignment: Responses should emphasize programmatic alignment and control compliance with major industry standards (e.g., NIST, FFIEC, ISO) and control framework (e.g., Cyber Risk Institute (CRI) AI Risk Management Framework). Providers should utilize standardized response language and abstract away specific technical configurations, proprietary tooling names, or raw internal environment metrics.
- Tiered & Client-Specific Data: Highly sensitive technical artifacts—such as architecture/segmentation diagrams, specific vulnerability findings, or localized remediation timelines—are recognized as NDA-restricted. This information should be evaluated on a case-by-case basis, restricted strictly to the scope of services provided to the specific evaluating institution, and provided when contractually required.
- Non-Disclosable Materials: Certain internal operational and governance data—including End-of-Life (EOL) asset inventories, SOC staffing metrics, and internal reporting or Board of Directors materials—are considered strictly confidential and excluded from external distribution.
- Governance & Gating: All evidence packages and questionnaire responses should undergo a mandatory internal review and sanitization process by the provider prior to release to ensure compliance with data protection policies.

LEGAL DISCLAIMER: USE AND GOVERNANCE

Use and Scope

This questionnaire is for risk assessment, due diligence, and governance purposes only, evaluating security posture, resilience, and regulatory alignment.

No Contractual Obligation

Submission and subsequent dialogue are for informational purposes. Responses do not independently create, amend, or supersede any contractual obligations or warranties.

Incorporation by Reference

Contents are legally binding only if expressly incorporated into a signed written agreement. In case of conflict, the terms of the executed contract prevail.

FRAMEWORK REFERENCES

- *Center for Internet Security Controls (CIS)*
- *Cybersecurity Infrastructure and Security Agency (CISA) Known Exploited Vulnerabilities Catalog (KEV)*
- *Cloud Security Controls Matrix (CCM)*
- *Cyber Risk Institute (CRI) AI Risk Management Framework (FS AI RMF)*
- *Federal Financial Institutions Examination Council IT Handbook (FFIEC)*
- *International Organization for Standardization (ISO)*
- *NIST Cybersecurity Framework (CSF)*
- *NIST Special Publications (SPs)*

CONTROL OBJECTIVES

Internal Controls = IC	External Vendor Controls = EVC	Both = B

SECTION 1 FRONTIER AI PARTICIPATION & SCOPE

Objective: Confirm participation in, or evaluation of, frontier AI security programs and ensure testing activities are conducted in a controlled, auditable manner.

Control Objective: B

Frameworks: *FFIEC IT Handbook (Management, Information Security); CCMv4.0; CSF v1.1: ID.RA-2*

1.A	Are you — or any of your critical subcontractors, cloud providers, or development partners — currently participating in, or actively evaluating, Project Glasswing, Claude Mythos, or a comparable frontier AI vulnerability discovery initiative? *If yes, identify which specific products, platforms, Application Programming Interface (APIs), or infrastructure are within scope. If no, describe your evaluation timeline.*
1.B	What specific controls are in place to ensure that all AI-assisted cybersecurity testing — including vulnerability discovery — is conducted in an authorized, auditable manner that prevents unintended disruption to Bank-facing production or staging environments?
1.C	Please identify the AI models used in connection with security testing, vulnerability discovery, and remediation support. In your response, please indicate: (i) whether any such models are developed or hosted outside the United States (and if so where), (ii) whether any models are open-source, community-developed, or otherwise not subject to a single commercial provider's governance controls, and (iii) how model origin, licensing, and jurisdictional considerations are assessed as part of your security and third-party risk management processes. *Responses may be provided at a high level and need not include proprietary model details or training data.*

SECTION 2 AI-ASSISTED VULNERABILITY DISCOVERY & PROACTIVE TESTING

Objective: Assess the breadth of AI-assisted testing, the integrity of discovery outcomes, and the operational safeguards that prevent testing from creating new risks.

Control Objective: B

Frameworks: *CCMv4.0; CIS Controls v8: 7.1; CSF v1.1: ID.RA-2; FFIEC VI.B.3; NIST SP 800-218*

2.A	How are you integrating AI-assisted discovery — including automated code review, penetration testing, automated software testing techniques where a program is bombarded with invalid, unexpected, or random data (fuzzing), and Software Composition Analysis (SCA) — into your environment? Have Bank-facing APIs, hosted platforms, or core banking applications been specifically scanned using tools such as Claude Mythos or equivalent frontier models? *Include scope, frequency, and high-level characterization regarding High-level (default) - approved brand statement or previously unknown vulnerabilities (without disclosure of exploit details or sensitive technical artifacts).*
2.B	When AI-assisted testing identifies high-risk or zero-day vulnerabilities, what is your process for validating those findings to eliminate false positives and to ensure remediation and customer notification both occur within defined timeframes?

2.C	What specific controls prevent AI-assisted testing activities from causing service instability, unauthorized data exposure, or other operational disruptions within Bank environments?

SECTION 3 SURGE READINESS & ADVERSARIAL AI THREAT ENVIRONMENT

Objective: Verify that the vendor has sufficient capacity, automation, and tested escalation procedures to remain effective when AI accelerates both discovery volume and adversarial exploitation speed.

Control Objective: IC

Frameworks: *CCMv4.0; FFIEC VI.B.3; CSF v1.1: RS.RP-1; CIS Controls v8: 7*

3.A	What specific processes, staffing, tooling, and automation have you implemented to handle a material surge in vulnerability findings? How do you prioritize simultaneous critical issues across multiple customer environments when AI-driven exploitation may shrink traditional remediation windows?
3.B	Have you conducted tabletop exercises or simulations specifically modeling AI-accelerated exploitation or mass vulnerability disclosure? How do you measure whether your current remediation capacity is sufficient to counter adversaries also using AI-assisted tools? *Indicate frequency of exercises and whether findings resulted in program changes.*

SECTION 4 REMEDIATION SERVICE LEVEL AGREEMENTS, PRIORITIZATION & GOVERNANCE

Objective: Establish clear, enforceable performance benchmarks, risk-based decision-making criteria, and board-level accountability for unresolved vulnerabilities.

Control Objective: B

Frameworks: *CCMv4.0; CRI Profile v2.0; CSF v1.1: RS.MI-3; FFIEC VI.C.1 (Service Management); CISA KEV Catalog*

4.A	What are your defined remediation Service Level Agreements (SLAs) for each risk tier (Critical, High, Medium, Low)? What is the formal escalation and emergency remediation process when a vulnerability poses a material threat to core banking services? *Include the authority level required to accept delayed remediation and how exceptions are tracked.*
4.B	Beyond base CVSS scores, how do you incorporate contextual risk factors — such as CISA KEV catalog status, internet exposure, authentication requirements, data sensitivity, and financial institution customer impact — into your remediation prioritization?
4.C	What is the formal governance and reporting process for escalating material, unresolved critical vulnerabilities to executive management or the Board of Directors?

SECTION 5 LEGACY ASSETS, EXTENDED SUPPORT & EMERGENCY PATCHING

Objective: Manage the heightened risk posed by older systems in an AI-accelerated threat environment, with clear remediation guarantees and a published modernization roadmap.

Control Objective: B

Frameworks: *CCMv4.0; CRI Profile v2.0; CSF v1.1: PR.MA-1; FFIEC VI.B.3*

5.A	Identify any systems, applications, operating systems, databases, middleware, APIs, or components supporting Bank services that are currently under extended support, limited support, or end-of-life (EOL) status. Provide the roadmap for their replacement or retirement.

5.B	For legacy assets, what is the guaranteed turnaround time for emergency security hotfixes once a zero-day or critical vulnerability is disclosed? If a patch cannot be deployed within 48 hours, what specific compensating controls (e.g., isolation, Web Application Firewall (WAF rules), privileged access restrictions) are automatically triggered? *Specify whether compensating controls include increased logging, endpoint detection logic, and manual monitoring.*
5.C	What is the formal process for notifying the Bank when a legacy component creates elevated remediation or availability risk, and how is this risk reflected in your internal governance and operational controls?

SECTION 6 SUPPLY CHAIN, SBOM & DEVELOPMENT PIPELINE INTEGRITY

Objective: Ensure full visibility into software dependencies and subcontractors and validate that the development pipeline is hardened against AI-driven supply chain attacks.
Control Objective: B
Frameworks: *CCMv4.0: CCC-04; CRI Profile v2.0: ID.AM-02; CSF v1.1: ID.AM-2; FFIEC VII.D; NIST SP 800-218 (SSDF)*

6.A	Do you maintain a Software Bill of Materials (SBOM) or Software Composition Analysis (SCA) output for Bank-facing services? Will you provide dependency summaries or vulnerability exposure attestations to the Bank upon request? What is your guaranteed timeframe for assessing Bank-specific exposure following a newly disclosed third-party vulnerability?
6.B	How do you monitor open-source dependencies for malicious packages, typosquatting, compromised maintainers, and abandoned projects? What technical controls — including code signing, artifact integrity checks, and controlled build pipelines — are in place to protect software supply chain integrity?
6.C	What controls — including multi-factor authentication (MFA), branch protection, Privileged Access Management (PAM), and segregation of duties — are enforced across your Continuous Integration / Continuous Delivery (CI/CD) pipeline? How do you detect unauthorized code changes such as malicious commits (a deliberate, unauthorized change introduced into a software codebase that is intended to compromise a system's security, introduce malware, or steal sensitive data) or compromised developer credentials? Can you verify that 100% of production deployments are traceable to approved change tickets and authorized personnel?

SECTION 7 AI DATA GOVERNANCE, PRIVACY & MODEL TRAINING

Objective: Prevent Bank-proprietary data from being ingested by AI models and ensure lifecycle governance over AI-generated security artifacts.
Control Objective: B
Frameworks: *CCMv4.0: DSP-18; CRI Profile v2.0: GV.RM-04; CSF v1.1: ID.RM-2; FFIEC II.C.13; GLBA; FS AI RMF*

7.A	Are Bank datasets — including transaction data, configurations, logs, and vulnerability findings — excluded from AI model training or fine-tuning by default? What contractual restrictions prevent Bank data from being used by, or disclosed to, AI providers or other third parties? *Include reference to specific contractual provisions or data processing agreements.*
7.B	Where are AI prompts, outputs, embeddings, and testing artifacts processed and stored? Who — internally and via third parties — has access to this data, and how is that access monitored, logged, and restricted?

7.C	How do you classify AI-generated security findings, and what specific retention and destruction rules apply to these artifacts to ensure they are not retained beyond operational necessity?
7.D	What AI governance framework governs your use of AI tools (e.g., *FS AI RMF*, ISO 42001, etc.), and how does it specifically address model risk, privacy obligations Gramm-Leach-Bliley Act (GLBA), and alignment with Payment Card Industry Data Security Standard (PCI, NIST), and applicable Digital Operational Resilience Act (DORA) requirements?
7.E	Does any AI-assisted security testing, prompt processing, or artifact storage occur outside the United States? What legal, contractual, and technical controls are in place to address cross-border data transfer, regulatory access, and sovereign risk considerations?

SECTION 8 CUSTOMER NOTIFICATION, REGULATORY SYNCHRONIZATION & INTERIM MITIGATION

Objective: Ensure the Bank receives timely, actionable notifications that support its own regulatory reporting obligations, along with clear technical and executive guidance.

Control Objective: B

Frameworks: *CCMv4.0: DSP-18; CRI Profile v2.0: RS.CO-02; CSF v1.1: RS.CO-2; FFIEC III.D; 12 CFR Parts 53/225/304 (OCC/Fed/FDIC Computer-Security Incident Notification Rule)*

8.A	What specific materiality thresholds trigger a customer notification (e.g., critical zero-days, confirmed compromise, data exposure, service disruption of 4+ hours)? What are your guaranteed notification timelines, and do you maintain Bank-designated emergency contacts for after-hours escalation? *Regulatory baseline: Under 12 CFR Parts 53/225/304, bank service providers must notify each affected banking organization customer as soon as possible upon determining they have experienced a computer-security incident that has caused, or is reasonably likely to cause, a material service disruption or degradation for four or more hours. Upon receiving such notice, the banking organization must then notify its primary federal regulator as soon as possible and no later than 36 hours after determining a notification incident has occurred. Your notification timeline must support that downstream obligation.*
8.B	Will your notifications be delivered in a timeframe that allows the Bank to meet its own regulatory reporting obligations? Will they include specific details such as affected systems, interim controls in place, recommended customer actions, and an expected remediation timeline?
8.C	Following a material event, will you provide formal written incident summaries, root-cause analysis (RCA), and verified evidence of remediation suitable for the Bank's audit files, examiner review, IT Steering Committee, and enterprise risk management (ERM) reporting?
8.D	When immediate remediation is not feasible, will you provide specific, actionable compensating controls (e.g., firewall rules, detection logic, access restrictions) along with a methodology for verifying their effectiveness until a permanent fix is deployed? Will advisories clearly delineate which actions require vendor intervention versus Bank-side configuration, and include plain-English risk summaries suitable for executive management and regulatory examiners?

SECTION 9 ARCHITECTURE, SEGMENTATION & INCIDENT RESPONSE

Objective: Evaluate resilience against successful breaches through strong isolation controls, tested recovery capabilities, and defined performance benchmarks.

Note: Requested responses can be provided in summary form or a high-level description.

Control Objective: B

Frameworks: *CCMv4.0; CRI Profile v2.0; CSF v1.1: PR.AC; FFIEC VII.B; NIST SP 800-61*

9.A	How are customer environments logically or physically segregated to prevent a compromise in one environment from affecting others? What architectural controls — including micro-segmentation, tokenization, and least-privilege access — limit lateral movement, and what is the maximum foreseeable 'blast radius' if a critical system component or privileged account is breached?
9.B	How frequently is penetration testing performed to verify the effectiveness of segmentation controls? Are Bank backups physically or logically segregated from the production environment, and can you confirm they are encrypted, immutable, and regularly tested for full restoration?
9.C	Does your Incident Response (IR) plan specifically address modern threat vectors — including AI-assisted attacks, rapid vulnerability exploitation, and supply-chain compromise? How frequently is the plan tested, and do exercises simulate critical scenarios such as core banking disruption, zero-day vulnerabilities, and ransomware?
9.D	What are the defined return-to-operation (RTOs) and recovery-point objective (RPOs) for services provided to the Bank? In a multi-client impact scenario, how are recovery priorities determined, and what manual workarounds are available to maintain Bank operations? Please provide current escalation procedures and contact information.
9.E	What changes have been made to your monitoring, detection, or response tools or processes to adapt to the evolving threats posed by AI-enabled vulnerability discovery and attacks? Are you employing or modifying any deception tools or techniques to better protect your environments from AI-enabled attacks?
9.F	What are the current Key Performance Indicators (KPIs) for Mean Time to Detect (MTTD) and Mean Time to Contain (MTTC), and how have these targets been adjusted to counter AI-driven threats?
9.G	Does your SOC have adequate manual procedures to protect the Bank's data if your automated monitoring technology suffers a catastrophic failure?
9.H	Is your Security Operations Center fully staffed for active threat hunting and detection 24/7/365, or does coverage transition to on-call during weekends/nights?

SECTION 10 SECURE DEVELOPMENT LIFECYCLE & AI USE IN DEVELOPMENT

Objective: Ensure developer use of AI tools is governed, Bank data cannot leak through unauthorized AI inputs, and AI-generated code is validated before deployment.

Control Objective: B

Frameworks: *NIST SP 800-218 (SSDF); CCMv4.0; CSF v1.1: PR.IP-2; FFIEC Development, Acquisition, and Maintenance*

10.A	Do you maintain written policies governing developer use of public or third-party AI tools for code generation, debugging, or vulnerability research? Do these practices align with the NIST Secure Software Development Framework (SSDF / SP 800-218) or a comparable standard?
10.B	What specific technical controls prevent source code, credentials, or Bank-specific proprietary information from being entered into unauthorized or public AI models?
10.C	Does your SDLC require mandatory human review of all AI-generated code before merge and deployment? What testing methods — such as SCA, logic checks, licensing scans, and dependency audits — are applied to validate AI-assisted commits?
10.D	Does your software/solution include the use of an AI model?
10.E	Has this model been built and tested against a framework such as the Open Worldwide Application Security Project (OWASP), Top 10 Large Language Model (LLM) and GenAI Application Risks?

10.F	Are updates to this model regression tested against a framework such as the OWASP Top 10 LLM and GenAI Application Risks?

SECTION 11 METRICS, CONTRACTUAL ACCOUNTABILITY & FORWARD-LOOKING PREPAREDNESS

Objective: Confirm governance visibility, enforceable contractual protection, and a credible strategic roadmap for the AI-accelerated threat environment.

Control Objective: B

Frameworks: *CCMv4.0: GRC-02; CRI Profile v2.0; CSF v1.1: ID.GV; FFIEC II.C.13; FFIEC Outsourcing Technology Services*

11.A	What vulnerability management metrics — including discovery rate, remediation rate, average time-to-remediate, overdue findings, and exceptions — are reported to senior management and the Board? Do you differentiate AI-assisted security findings from traditional scanning results in your reporting?
11.B	Do your agreements with the Bank include comprehensive provisions for security incident notification, audit rights, indemnification, and service-level remedies? Are there limitations of liability that would restrict the Bank's recovery in cases involving gross negligence, data breaches, or failure to meet security SLAs? *Identify whether AI-related risks (e.g., AI-generated code defects, unauthorized AI tool usage, AI-assisted testing failures) are explicitly covered or excluded from warranty and indemnity provisions.*
11.C	What material changes have been made to your cybersecurity program in direct response to frontier models such as Claude Mythos or Project Glasswing? What changes are planned over the next 6, 12, and 18 months — including staffing, tooling, AI governance, and customer communication processes? *Describe your strategy for maintaining detection and remediation velocity as adversaries increasingly adopt AI-assisted exploitation.*
11.D	Identify any information, access, or advanced coordination from the Bank that would improve your ability to deliver timely notifications, execute joint incident response, or conduct coordinated vulnerability testing. The Bank is prepared to discuss mutually beneficial arrangements that strengthen both parties' security posture without compromising its own regulatory obligations.
11.E	Who within your organization has formal accountability for AI-assisted cybersecurity decisions (e.g., vulnerability discovery tooling, AI-generated findings, remediation prioritization), and how is this accountability reflected in governance structures, escalation paths, and executive or board reporting?
11.F	Do you maintain contractual terms that impose AI governance obligations on downstream vendors and subcontractors?
11.G	What is the frequency and rigor of vendor due diligence reviews as they relate specifically to AI usage, AI governance, and AI risk management practices within the vendor ecosystem?

REQUESTED EVIDENCE PACKAGE

Please provide the following documentation where available. Materials may be submitted under NDA. If a document does not exist, indicate the reason and your plan to address the gap.

Category	Requested Documentation / Evidence
Governance & Policy	AI Cybersecurity Governance Policy (or board-approved summary)

Category	Requested Documentation / Evidence
	AI Acceptable Use Policy governing developer and security team use of LLMs and generative tools
	Subcontractor Oversight Program summary, including AI tool governance requirements for downstream vendors
Vulnerability & Technical Security	Vulnerability Management Policy, including the SLA schedule for Critical, High, Medium, and Low risk tiers
	Secure SDLC Policy or summary of controls integrated into the CI/CD pipeline
	Software Bill of Materials (SBOM) or software composition summary for Bank-used products
	Summary of material security changes implemented in response to Claude Mythos, Project Glasswing, or equivalent AI-assisted discovery findings
Assurance & Testing	Executive Summary of the most recent annual Penetration Test
	Most recent SOC 2 Type II report or equivalent independent assurance report (e.g., ISO 27001, SIG Lite)
Resilience & Incident Management	Incident Response Plan summary, including specific procedures for customer notification following a material cybersecurity event
	Business Continuity Plan (BCP) and Disaster Recovery (DR) strategy summary
	Emergency Escalation Contact List for security and executive leadership