yubico

5 Identity Controls to Include in Your AI Threat Mitigation Plan

A working checklist for banking security and risk teams preparing for the ECB's 31 October deadline.

On 7 July 2026, the European Central Bank directed eurozone banks to have plans in place by 31 October to address AI-enabled cyber threats that could disrupt financial services. This checklist focuses on identity—consistently one of the fastest-moving, fastest-to-fix pieces of that plan, and one of the first places AI is changing attacker economics.

Use it as a working document: check off what's already covered, flag what needs an owner, and treat unchecked items as your near-term action list.

The Checklist

1. Phishing-resistant authentication for privileged and customer-facing access

Why AI raises the stakes:
AI-generated phishing is more convincing, more personalized, and produced at far greater scale than manual campaigns. Any authentication method relying on a human recognizing a fake—codes, links, push approvals—is now a weaker control than it was two years ago.

- [ ] Privileged accounts (admins, engineers, treasury/payments staff) require hardware-backed, phishing-resistant MFA (FIDO2/WebAuthn), not SMS or app-based OTPs.
- [ ] High-risk customer flows (large transfers, account recovery, new payee setup) use step-up authentication that can't be relayed or approved under social-engineering pressure.
- [ ] MFA coverage has no exceptions for "low-risk" or legacy systems—these are frequently where attackers pivot first.

2. Deepfake-resistant verification for helpdesk and call-center processes

Why AI raises the stakes:
Voice-cloning tools can now defeat call-center identity checks that rely on voice recognition, security questions, or callback verification—the exact channel used for account recovery and high-value transaction approval.

- [ ] Helpdesk and call-center identity verification does not rely solely on voice, security questions, or caller ID.
- [ ] Account recovery and payment-change requests require a step that can't be completed by an AI-generated voice or message alone (e.g., hardware key, in-app confirmation on a trusted device).
- [ ] Staff are trained on current deepfake-enabled social engineering patterns, not just legacy phishing awareness content.

3. Continuous identity and access risk monitoring

Why AI raises the stakes:
AI is shrinking the window between a stolen credential and its use. Point-in-time access reviews (quarterly, annually) are too slow to catch AI-accelerated credential abuse before damage is done.

- [ ] Anomalous authentication patterns (impossible travel, new device + high-risk action, off-hours privileged access) trigger automated alerts, not just quarterly review.
- [ ] Session and token risk is monitored continuously, not only at initial login.
- [ ] Ownership for reviewing and acting on identity risk alerts is assigned to a named team, not left ambiguous across security and IT.

4. Segmentation and just-in-time access for critical systems

Why AI raises the stakes:
AI speeds up exploitation timelines, making standing access to critical systems a major liability. Removing persistent permissions for both employees and supply chain vendors reduces the blast radius, and directly addresses DORA's strict mandates for Third-Party ICT Risk Management.

- [ ] Access to core banking, payments, and customer data systems is time-bound or approval-gated rather than standing/always-on.
- [ ] Privileged access is segmented so that compromise of one system or one identity doesn't cascade into adjacent critical systems.
- [ ] Third-party and vendor access to critical systems follows the same just-in-time principle as internal privileged access.

5. Incident response playbooks updated for AI-accelerated timelines

Why AI raises the stakes:
Response plans built around days-to-weeks detection and containment windows assume attacker timelines that AI has already compressed. A plan that isn't rehearsed against faster timelines is a plan that arrives late.

- [ ] Detection-to-containment SLAs have been reviewed and tightened in the last 12 months to reflect AI-accelerated exploitation timelines.
- [ ] Incident response has been tabletop-tested against an AI-enabled scenario (mass credential phishing, deepfake-enabled fraud, or automated exploit chaining), not only traditional breach scenarios.
- [ ] Escalation paths and decision authority are documented and known in advance—not decided for the first time during an active incident.

Where to start

Controls #1 and #2 are the fastest to close, and carry the highest-leverage: phishing-resistant hardware authentication like the YubiKey can be deployed immediately to solve critical identity gaps without replacing existing IAM infrastructure.

Talk to our team today about mapping this checklist to your environment before the 31 October 2026 deadline.

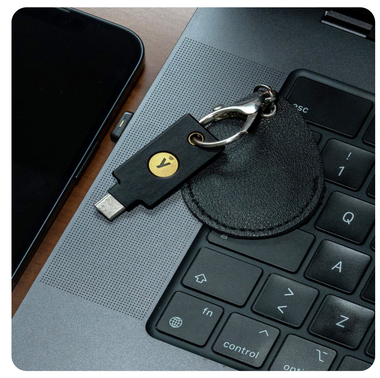

Contact Us
yubi.co/contact