Cyber Resilience in the Age of AI-Driven Warfare

Guidance for Bank CIOs Responding to the ECB Letter of 7 July 2026

By Snehal Antani

Co-Founder & CEO, Horizon3.ai

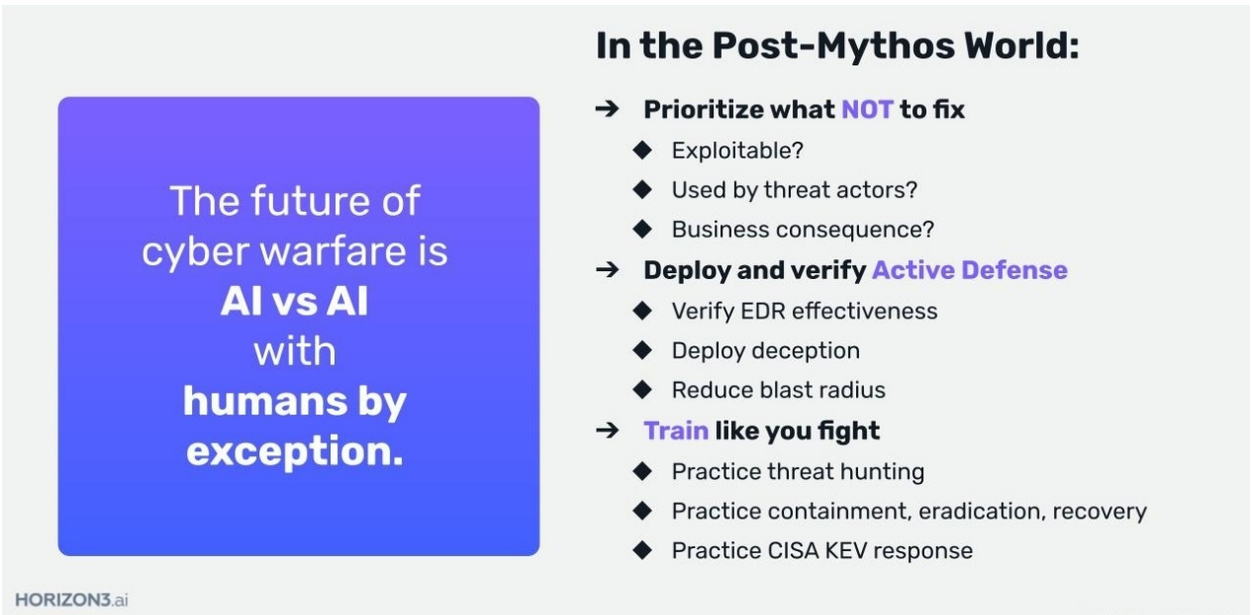

Figure 1: Operating Framework for the Post-Mythos Era

I have spent my career at the intersection of enterprise technology, national security, and the building of autonomous security systems. From my time as CTO at Splunk and CIO roles at GE Capital, through my service as the first CTO of Joint Special Operations Command, to co-founding Horizon3.ai, one lesson has been consistent:

Don't tell me I'm secure, show me, then show me again tomorrow, then again next week, because the environment is always changing and the adversary has a vote.

This mindset has never been more important than in the age of Gen AI, accelerated by Mythos and similarly capable open-source models

Traditional vulnerability management, periodic audits, and static defenses were already insufficient. In the Post-Mythos world, they are actively dangerous because they create a false sense of progress while adversaries operate at machine speed.

Mythos and its successors do not simply find more vulnerabilities—they compress the time from discovery to exploitation dramatically. What once required teams of skilled researchers now happens in minutes or hours, at scale, against anyone. The volume of potential issues

explodes. The cost to attackers plummets. The blast radius of a single successful chain grows. This is not hype. It is the new operational reality for every bank CIO, every CISO, and every Management Body that oversees critical financial infrastructure and customer data.

The European Central Bank's letter of 7 July 2026 (SSM-2026-0301) from Claudia Buch makes this reality unavoidable for the 110 significant institutions under its direct supervision. The letter requires a comprehensive Action Plan by 31 October 2026 that addresses both immediate priorities and longer-term structural measures. It correctly identifies the shift as structural rather than temporary, and places responsibility squarely with bank Management Bodies.

This paper offers a practical operating system that CIOs can use to meet that requirement with clarity and credibility.

The Three Pillars: Minimum Viable Operating System for Post-Mythos Resilience

The three pillars in Figure 1 are not suggestions. They are the minimum viable operating system for cyber resilience going forward. I expand on each from the perspective of someone who has lived the pain of misallocated security resources and the urgency of real-world defense.

1. Prioritize What NOT to Fix

This is the hardest and most important shift. As a CIO I used to dread Patch Tuesday—not because of the patches themselves, but because of the impossible prioritization that followed. Everything is labeled critical. Every vendor pushes urgency. Security teams burn out chasing noise while real, exploitable attack paths remain open.

In the Post-Mythos world this problem becomes existential. AI can generate candidate exploits and validate them against code at a pace humans cannot match. The backlog will grow faster than any team can triage using traditional CVSS scores or vendor severity alone.

The three filters in Figure 1 must be applied ruthlessly and continuously:

- ***Exploitable?*** Not every vulnerability is reachable in your environment. Not every misconfiguration can be chained into impact. Continuous, safe testing that actually attempts exploitation (in production or faithful replicas) separates signal from noise. Prove exploitability; do not merely list CVEs.

- ***Used by threat actors?*** Intelligence feeds matter, but context matters more. Is this vulnerability appearing in campaigns targeting the financial sector or similar environments? Is it in the CISA Known Exploited Vulnerabilities (KEV) catalog for a reason? AI lowers the bar for less sophisticated actors, so "commodity" exploits become dangerous at scale.

- ***Business consequence?*** This is where security meets the board. Can this path lead to data exfiltration that triggers regulatory fines under DORA or GDPR? Operational disruption that halts payments or customer access? Intellectual property or competitive loss? In my DoD experience we thought in terms of mission impact. Banks must do the same—map to regulatory capital, customer trust, and franchise risk.

The organizations that survive will treat prioritization as a continuous, data-driven discipline informed by real attack paths, not annual risk registers. Deciding what not to fix frees resources for what actually moves the needle on risk. This pillar directly answers the ECB's short-term call to accelerate vulnerability and patch management at scale, especially for internet-facing, cloud, third-party, and open-source assets.

2. Deploy and Verify Active Defense

Passive defense is a losing strategy against adaptive, autonomous adversaries. You must assume breach and build mechanisms that detect, slow, mislead, and contain in real time. But assumption is not enough—you must ***verify*** that these mechanisms work against the same techniques attackers will use.

- ***Verify EDR Effectiveness:*** Endpoint Detection and Response tools are table stakes, yet many organizations have never tested whether their specific configuration and tuning actually catch the behaviors that matter in their environment. AI attackers will use living-off-the-land techniques, fileless execution, and rapid lateral movement that signature-based or poorly tuned behavioral rules miss. Continuous testing that exercises the full kill chain against your EDR deployment is the only way to know. If an AI-driven attacker can establish persistence or exfiltrate data without triggering high-fidelity alerts, you have a problem that needs fixing now.
- ***Deploy Deception***: Deception technologies—honeypots, honeytokens, decoy credentials, fake data—force attackers to reveal themselves and waste cycles on low-value targets. In an AI-versus-AI world, deception becomes even more powerful because autonomous agents can be lured into observable loops or forced to make detectable mistakes. Research shows that frontier AI models fall for well-placed deception significantly more often than expert human operators and frequently proceed even after recognizing the trap. Deception must be dynamic and integrated with your detection stack; static decoys are quickly fingerprinted. Use it primarily as a high-yield, early-warning detection layer rather than pure misdirection.
- ***Reduce Blast Radius:*** This is both architectural and operational. Micro-segmentation, least-privilege access, just-in-time credentials, and strong identity controls limit how far an initial foothold can travel. But these controls drift. AI can discover and chain subtle misconfigurations faster than governance processes can correct them. The only reliable way to maintain a reduced blast radius is continuous end-to-end attack-path validation. When testing shows a path from a compromised

workstation to crown-jewel data or payment systems in under ten minutes, that is the signal to harden the specific controls that failed—not to add another layer of monitoring.

Active defense is not about replacing your existing tools; it is about closing the loop between deployment and proven effectiveness. "Deploy and verify" is deliberate: deployment without verification is theatre. This pillar addresses the ECB's requirements for stronger monitoring, detection and AI-enabled defensive capabilities, as well as third-party and supply-chain risk management.

3. Train Like You Fight

This principle comes directly from military and special operations culture, where realistic, high-tempo training builds the muscle memory that separates success from failure under stress. Cyber is no different—except that the "enemy" now moves at machine speed and never sleeps.

Tabletop exercises and annual compliance drills are necessary but insufficient. They lack the realism of actual toolchains, live production-like environments, and the cognitive load of real alerts mixed with false positives. In the Post-Mythos era your blue team must regularly face AI-augmented red-team activity that mirrors what nation-states and sophisticated criminals are deploying.

Three critical practice areas:

- ***Practice threat hunting:*** Proactive hunting for indicators of AI-driven activity—unusual authentication patterns, anomalous data access, living-off-the-land sequences that evade EDR. AI can assist with correlation at scale, but human analysts with domain expertise must stay sharp.

- ***Practice containment, eradication, and recovery:*** When an AI attacker moves in minutes, your response must be measured in minutes, not hours. This requires pre-approved playbooks, automated containment options (network isolation, credential revocation, process termination), and muscle memory from realistic drills. Eradication is harder than containment—removing persistence, validating clean images, and restoring trust.

- ***Practice CISA KEV (and equivalent) response:*** Treat high-signal vulnerability catalogue additions as priority events requiring rapid assessment and action. Regular drills ensure that when a new entry drops (and they will drop faster with AI assistance), your team does not start from zero.

Training like you fight also means purple teaming—red and blue working together in real time. With AI on both sides, these exercises become laboratories for understanding how autonomous agents behave and how to counter them. The organizations that institutionalize this will develop a decisive advantage.

Snehal S. Antani, CEO @ Horizon3.ai, http://www.linkedin.com/in/snehalantani

Mapping these pillars to the ECB July 7th letter

ECB Short-Term Priority	Pillar	Evidence of Progress
Accelerate vulnerability & patch management at scale (especially exposed assets)	1	Continuous exploitability-based prioritisation; reduced time-to-remediate for high-impact findings
Enhance monitoring, detection & AI-enabled defensive capabilities	2	Continuous EDR verification; dynamic deception as detection layer; measured coverage of AI-relevant techniques
Verify third-party risk management is fit for purpose	2	Supply-chain attack-path visibility; updated assurance and contractual mechanisms

ECB Longer-Term / Structural Priority	Pillar	Evidence of Progress
Reinforce defence-in-depth & cyber hygiene; modernise infrastructure	2	Continuous blast-radius validation; legacy reduction roadmap
Improve response & recovery, crisis management, information sharing	3	Frequency and effectiveness of purple-team, containment and KEV-style drills; mean time to contain

Key metrics to capture and report to key stakeholders

1. ***Mean Time to Remediation (MTTR)*** – The time to close critical exploitable risks
2. ***Mean Time to Detection (MTTD)*** – The time from initial access to SOC detection
3. ***Critical Exploitable Weaknesses Over Time*** – Improvement & regression trends

The question is no longer whether you will be tested by AI-driven attacks. The question is whether you will have practiced, prioritized, and verified your defenses before those attacks arrive. The banks that answer yes will not merely survive the coming wave—they will define the standard for resilience in the AI era and will submit Action Plans that satisfy both the letter and the spirit of the ECB's requirements.

Snehal Antani is Co-Founder and CEO of Horizon3.ai. He previously served as the first Chief Technology Officer for Joint Special Operations Command (JSOC), CTO and SVP at Splunk, and held CIO roles at GE Capital. He is a software engineer by training (Purdue, RPI) and holds multiple patents in distributed systems and security.