ASPEN DIGITAL
aspen institute

Cybersecurity in a Post-Mythos World: Upgrading Defenses and Improving Governance

By Jeff Greene, Matt Altomare, Rob Joyce, Rob T. Lee, Joe Levy, Sezaneh Seymour

AI has brought cybersecurity to an inflection point that requires more than new tools: we need better, more nimble governance policies. Models like Anthropic's Mythos can now find and exploit software flaws faster than humans can fix them. But rather than create a new class of threats, they accelerate existing ones. Put differently, the core reasons that networks are insecure and the basics of securing them have not changed, but the skills an adversary needs to identify and exploit vulnerabilities are now dramatically lower.

Organizations must now significantly improve their approach to cybersecurity to meet this new reality.

Robust cybersecurity requires people, governance, and financial resources, yet many executives and boards still view it as merely a technical problem. This disconnect is even more acute when speed matters, because the biggest barrier to rapid deployment of security upgrades is almost always organizational: people and policies. Instead, cybersecurity should begin with risk management and governance, with tone and guidance set at the top of the organization

This paper focuses first on governance, offering questions executives and boards can use to assess the strengths, weaknesses, and gaps of their cybersecurity and risk management programs. It then outlines practical technical steps management can take now to strengthen their defenses in a rapidly evolving threat landscape. You should approach this moment with discipline and urgency, and take advantage of the short window we have before these AI capabilities become widely available.

Questions Executives & Board Members Should Ask

The questions below are intended to help board members and CEOs get a strategic and governance-level understanding of their organization's security profile and whether it can move rapidly enough to address the growing risks associated with AI. The questions are not intended to insert leaders into day-to-day technical decision-making but should give them a clearer line of sight into an organization's preparedness, priorities, and resilience.

1. **If all our critical systems went down tomorrow morning, how long before we're operational again, and have we ever actually tested that?** A good answer includes a

timeline that your team can back up with strong procedures and recent testing. If the team cannot support a claim of rapid recovery or the answer is “weeks” or "we're not sure," that's a foundational risk that executives and the board should know about and support fixing. **Highest Priority.**

2. **What is our plan if we get hacked?** A good answer begins with one or more incident response firms on retainer, or a cyber insurance policy with direct access to an incident response provider, and a managed detection and response (MDR) service. It also includes an outreach plan to named government partners, a back up “out of band” communications plan in case an attacker is on your network, and draft internal and external messaging. If the answer is silence or "we'd figure it out," that's a problem, because minutes matter with incident response. **Highest Priority.**

3. **How quickly can we patch a critical vulnerability?** A good answer is measured in hours and includes an explanation of how your team tracks this metric and whether you are improving your response time. If nobody is tracking the response time or the answer seems uncertain, your organization probably does not know whether it's improving or falling behind. **Highest Priority.**

4. **How are we using AI for our own defense?** The answer you want to hear is specific: we either implement AI agents ourselves or partner with vendors who can. Elements of defense should include a full review of all code before it ships, incident response triage, and vulnerability management and audit work. If the answer is "we're studying it," or "we have concerns about the risk of using AI," that is itself a risk decision. As leadership, you need to make this decision with the understanding that your adversary is using AI, and you need tools that match the pace of the threat. **Highest Priority.**

5. **If you could fix one security problem to most reduce your AI risk, what would it be, and what do you need to make it happen?** Security and IT leaders often know exactly what needs to happen but lack the authority, budget, or political capital to implement the solution. This question gives your team a chance to identify the most troubling gaps and demonstrates that leadership prioritizes security.

6. **Financially, can we withstand a major cyber event without compromising strategic priorities?** Cyber incidents are capital-intensive events that can materially strain liquidity and erode business value. Organizations need to be robust enough across insurance coverage, retained risk, and ready access to capital. You need to know if your organization can quickly mobilize funds in a way that sustains operations, protects growth initiatives, and preserves investor and customer confidence.

7. **How many people have administrative access to our systems, and who reviews that access?** A good answer includes a specific number that is small relative to the size of your organization. It also includes an established process for reviewing these

permissions. If the answer is "it will take a while to check," or the number is surprisingly large, you have identified a structural problem as well as the possibility that these high-risk privileges have been accumulating unchecked for years.

8. **Are we using multi-factor authentication on everything, and if not, why not?** A good answer is a confident "yes." However, "no" is not necessarily a bad answer – some legacy systems may not support phish-resistant MFA. In this case, an acceptable answer includes knowledge of which systems lack MFA, a description of compensating controls such as zero-trust architectures, and a plan for replacing them. A bad answer is an uncertain "yes" or a "yes" with reference to text messages as your MFA. A "no" followed by "we haven't gotten to it yet "or "we don't have the budget" signals that your organization has not prioritized this fundamental security tool - you need to change this.

9. **Do we have an accurate picture of which third parties we depend upon and what data they have access to?** A good answer includes a tiered inventory of third-party dependencies ranked by business impact. Your team should be able to back this up with modeling that evaluates the impact of key third parties going down or losing sensitive data in a breach. Conversely, a simple list of your primary vendors without a deeper understanding of how their compromise could cascade across your organization or impact its reputation reveals a failure to consider your true exposure.

10. **Are we running any important systems or software that the vendor doesn't support or update?** A good answer is a confident "no." However, a "yes" can be acceptable if your team knows which systems may be at risk, has put compensating controls in place, and has a plan for replacing them. A "yes" followed by "we don't have the budget" gives you an opportunity to provide your team the support it needs to close this vulnerability.

11. **What is our organization's policy on employee use of AI, and how do we know if people follow it?** A good answer is a realistic assessment of the gaps between your policy and how employees use AI, and names the tools that block unsanctioned use. You must know what's sanctioned, what's tolerated, what's prohibited, and how you check, because employees across the business are almost certainly using AI tools whether or not you've formally approved them. If the answer is "we don't have a defined policy" or "we have a policy" without the sense of the actual usage or how to check, your organization has a governance gap that can result in data leakage, intellectual property exposure, and new attack surfaces that nobody is defending.

Technical Steps to Take Now

The list below includes both operational and policy actions you should take now. You should not treat them as a one-size-fits-all checklist; some may not be relevant to your organization.

Choose the measures that best fit your business, technology environment, and risk profile, with the aim of reducing exposure, limiting damage, and improving resilience. Experienced security teams are likely taking many of the actions below and may already be considering how to incorporate AI into their security stack. Smaller or less mature organizations may need to focus on fundamental hygiene and structural hardening. We have also flagged the highest priority actions – those you should assure your organization has taken or will take.

1. **Harden identity and deploy phish-resistant MFA.** Every account, remote access tool, and SaaS application – without exception – should use either modern hardware security keys, passkeys, or an authenticator app with number-matching. If a system can't support MFA, put compensating controls in place and make a plan to retire it. Separate and minimize administrative accounts, minimize privileged access, and aggressively monitor service accounts. Extend this rigor to AI agents, treating them as high-privilege non-human identities with access to critical code and data. Before deploying AI workflows, define strict blast-radius limits and human override mechanisms. Audit the entire agent framework – including prompts, tool definitions, and permission scopes – and integrate these assets into your identity governance and behavioral monitoring systems. **Highest Priority.**

2. **Turn on automatic updates on every asset that supports them.** Workstations, laptops, phones, browsers, routers, firewalls. For anything that can't auto-update, assign a named person to check for and deploy patches on a regular cadence. **Highest Priority.**

3. **Aggressively reduce external exposure, with a focus on internet-facing and edge systems.** Prioritize rapid patching, isolation, and hardening for externally exposed systems, especially edge devices, web applications, identity platforms, and remote access infrastructure. Decommission internet-facing systems that are not essential or which are no longer supported by vendors. Retire anything you cannot defend and modernize anything that you must keep. Create an emergency patching track for critical CVEs on internet-facing systems, and create a triage model for prioritizing high-volume vulnerability findings. **Highest Priority.**

4. **Authorize automated responses and accept the potential for false positives.** Pick five or six security response actions where the cost of a false positive is lower than the cost of waiting even 15 minutes for a human. Define and get executive sign-off for automated containment actions, including isolating a compromised host, blocking a known-malicious IP, killing an anomalous process, or revoking a credential that just did something impossible. Deploying an AI security tool without this functionality strips it of its effectiveness. **Highest Priority.**

5. **Implement zero-trust and microsegmentation.** Inventory and then separate crown-jewel systems, production administration, backups, and sensitive business processes

from general user environments. Restrict the ability to move sideways ("east-west") between internal systems and reduce implicit trust so systems don't automatically trust each other just because they are on the same network.

6. **Secure your backups and prioritize resilience.** Make secure backup and recovery a top-tier priority by protecting backups with separate credentials, isolation, and immutability where available. Validate and exercise your ability to restore the systems the business cannot survive without.

7. **Expand and improve your logging beyond the endpoint, and build in detection.** Centralize all of your logs, whether from identity systems, cloud control planes, edge devices, DNS, admin actions, virtualization, and key business applications. Put AI agents at the front line triaging alerts, enriching them with intelligence, and investigating anomalies continuously - logs without detection are just expensive stored data. Make sure you retain them long enough to investigate slow-moving intrusions.

8. **Turn AI inward; point agents at your own code, pipelines, and configurations now.** Begin by having a coding agent perform security reviews of your highest-risk code before it ships, and build toward AI-driven review as a standard gate in your continuous integration and continuous deployment (CI/CD) pipeline. Ultimately, there should be no code, whether human-written or AI-generated, reaching production without machine-speed security analysis. Do the same for configuration reviews, infrastructure-as-code, and third-party dependencies.

9. **Move from periodic penetration testing to continuous attack simulation.** Annual or quarterly penetration tests are snapshots that are stale before the report is written. You need to discover what an AI-enabled attacker would find before they find it. If you can't run this in-house, engage a firm that offers continuous red-teaming as a service with AI-augmented tooling.

10. **Update and exercise your incident response plans.** Review your incident response plans and update them to reflect the accelerating pace of threats. Exercise them regularly with the goal of streamlining decision making and reducing response time. Pre-authorize specific containment actions for defined scenarios and use automation where it speeds up low-regret actions.

11. **Get your organization into a coordinated vulnerability program.** Plug your organization into CISA's disclosure and known-exploited-vulnerability channels or the equivalent channels run by your national CERT/CSIRT or cybersecurity agency. Ensure you have a disclosure intake and response process, and make sure your major vendors publish clear vulnerability disclosure policies and rapidly tell you when findings affect you.

12. **Mandate, don't suggest, AI agents use across your security team.** Coding agents and AI-assisted workflows are mature enough to accelerate nearly every security function today: vulnerability triage, alert investigation, threat intelligence analysis, audit evidence collection, detection engineering, and incident response. Make agent use a standard expectation for every role on the security team, from the SOC analyst to the CISO, with appropriate guardrails and training.

13. **Fast-track procurement and governance for defensive technology.** Most organizations are slow to evaluate and onboard new security tools. If the time between vulnerability disclosure and active exploitation is hours, an procurement cycle is itself a vulnerability. Stand up a cross-functional fast-track process bringing security, legal, engineering, and procurement together to evaluate and deploy priority defensive technologies in weeks, not quarters.

Ultimately, the race to meet AI-driven threats is about more than technology, it requires a shift in organizational leadership that reflects this new environment. The technical steps above will provide necessary protection, but they will not be fully effective without a nimble governance framework. Executives and boards must move beyond passive oversight and treat security as a core business discipline. By removing the friction that slows down everything from patching to procurement, by mandating AI integration across the team, and by setting a security-first tone across the organization, leadership can meet this challenge. The window is closing; the time for decisive action is now.