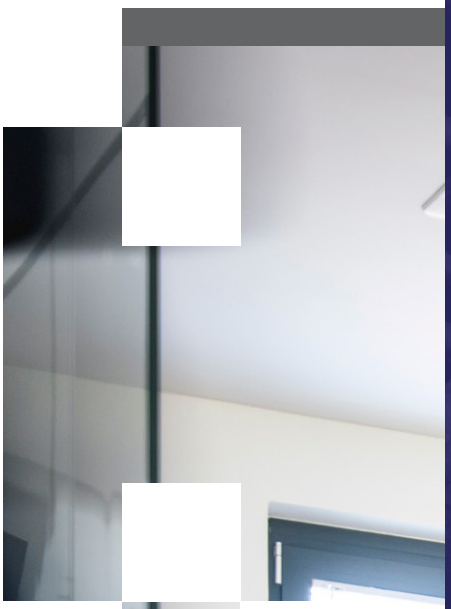

Future is Now.

How Project Mythos Rewrites the Rules of Cybersecurity

By Richard Itri, Chief Innovation Officer, ECI

April 2026

Foreword: The Cracks Were Always There

In March of this year, I published Full *Throttle, Full Control*—a decision framework for how alternative investment managers should evaluate enterprise AI platforms on security, compliance, and resiliency. That paper opened with a simple premise: the regulatory framework for governing AI in financial services already exists, and firms that wait for AI-specific rules will find themselves explaining why they weren't following the ones already on the books.

Six weeks later, Anthropic proved the point—
in the most dramatic way imaginable.

On April 7, 2026, Anthropic announced Claude Mythos Preview and Project Glasswing: a frontier AI model that had autonomously discovered thousands of zero-day vulnerabilities across every major operating system and every major web browser, some of which had survived 27 years of human review and decades of conventional wisdom that said "if it were that bad, someone would have found it by now." Someone finally did. It just wasn't a person.

I've spent 28 years in financial services technology. I've watched this industry survive the dot-com crash, Sarbanes-Oxley, the 2008 financial crisis, the SolarWinds supply chain compromise, and the ransomware epidemic. Each time, the industry adapted. This time is different—not because Mythos is a more powerful tool, though it is, but because it operates at a speed and scale that makes the gap between "good enough" security and actual security a matter of hours, not months.

This paper is about what Mythos revealed. The five structural weaknesses outlined here are not the result of negligence. They are the natural accumulation of practical compromises that every firm and every managed service provider has made—compromises that were entirely tolerable in a world where attackers moved at human speed. That world ended on April 7th. And the problem is only going to accelerate.

This is the New Norm

Mythos is not a one-time event. It is the beginning of a permanent shift. Every major AI lab is on the same capability trajectory. Researchers estimate Mythos-class capabilities will reach open-source models within 12 to 18 months. The threat landscape will never again be slower than it is today. Every framework, process, and partnership described in this paper must be designed not for a single moment of crisis, but for a permanently accelerated operating environment.

The Mythos Moment

Claude Mythos Preview is a general-purpose frontier AI model that was not explicitly trained for cybersecurity. Its capabilities emerged as a downstream consequence of improvements in code reasoning and agentic execution—the same improvements that make it better at writing software also make it better at breaking it. The model discovered and exploited zero-day vulnerabilities in every major operating system and web browser, found a 27-year-old bug in OpenBSD, and in one case chained four separate vulnerabilities into a browser exploit that escaped both renderer and OS sandboxes. It didn't just find bugs. It wrote the attack.

Three characteristics set Mythos apart from previous security tools. It chains vulnerabilities into end-to-end exploit paths rather than flagging individual bugs. It operates autonomously, producing working exploits within hours without human guidance. And it will not stay exclusive—researchers at AISLE tested Mythos's showcase vulnerabilities on small, open-weight models costing as little as eleven cents per million tokens and found that eight out of eight detected the flagship FreeBSD exploit.

By the Numbers

The following data points, drawn from Anthropic's published testing and independent research, quantify the scale of the capability shift:

- **Exploit success rate:** Anthropic's prior model, Opus 4.6, had a near-0% success rate at autonomous exploit development. Mythos succeeded 181 times out of several hundred attempts on Firefox's JavaScript engine alone, plus 29 additional instances of register control.

- **Zero-day volume:** Thousands of zero-day vulnerabilities discovered across major open-source and commercial codebases in weeks of testing—many critical, many 10 to 20+ years old.

- **Cost of replication:** AISLE researchers demonstrated that smaller, openly available models detected the same flagship vulnerabilities when given targeted code segments—at costs as low as $0.11 per million tokens.

- **Systemic response:** Treasury Secretary Bessent and Fed Chair Powell convened an emergency meeting with the CEOs of Citigroup, Morgan Stanley, Bank of America, Wells Fargo, and Goldman Sachs to address the cyber risks. JPMorgan Chase joined as a Glasswing launch partner.

- **Regulatory acceleration:** The EU AI Act's next enforcement phase takes effect August 2, 2026, with penalties up to 3% of global revenue. SEC cybersecurity disclosure rules and DORA's ICT third-party risk framework are already in force.

Anthropic launched Project Glasswing—a consortium including Amazon, Apple, Cisco, CrowdStrike, Google, JPMorgan Chase, Microsoft, and Palo Alto Networks—to harden critical systems before these capabilities proliferate. They committed $100 million in usage credits. The Glasswing partners will publish their findings in early July 2026, triggering what analysts expect to be a patch tsunami across operating systems, browsers, and core infrastructure software.

For alternative investment managers, the implication is direct: the Glasswing partners are hardening their platforms now. The firms that depend on those platforms will benefit over time. But the gap between when defenders harden and when equivalent capabilities reach adversaries is the exposure window—and it is narrowing.

1. Hygiene and the Exception Graveyard

"The basics were never optional. Now they're existential."

The overwhelming majority of successful breaches against financial services firms exploited basic hygiene gaps. An unpatched server. A service account with a password that hasn't been rotated in years. A conditional access policy that exempts a senior executive because MFA was creating friction. A firewall rule opened for a vendor integration that was completed three years ago and never closed. Mythos doesn't need to chain four zero-days together if your environment has an open RDP port and stale admin credentials. It will walk through the open door first.

The result is an exception graveyard — a growing catalog of policy deviations that accumulate naturally over time, often without a clear owner or a defined expiration. Each exception is a node in a potential exploit chain. These exceptions don't accumulate because of negligence. They accumulate because both firms and their service providers are optimizing for the same thing: keeping the business running with minimal disruption. Each individual decision makes sense in context. It's the aggregate — the full catalog of small, reasonable compromises compounding over months and years — that creates the exposure.

Partnership Action Framework: Hygiene

The following shared accountability model converts the hygiene diagnosis into assignable, measurable actions with explicit ownership:

The Firm Owns	The MSP Owns	Target Timeline
Authorize full exception audit; designate internal approver for exception lifecycle decisions	Conduct comprehensive audit of all active exceptions across firewall, CA policies, MFA, service accounts; deliver categorized inventory with risk scores	Complete by May 15
Adopt 90-day maximum exception lifecycle policy; approve exception review cadence (monthly)	Implement automated exception expiration with 14-day advance notification; build exception dashboard with aging metrics	Policy adopted by May 1; tooling by June 1
Identify and communicate business-critical systems where patching downtime is constrained	Classify all managed assets into Tier 1 (patch within 72 hrs), Tier 2 (7 days), Tier 3 (14 days) based on exposure and criticality	Tiering complete by May 15
Require MFA for all personnel without exception; communicate policy to senior leadership	Remediate all MFA exemptions; implement phishing-resistant MFA (FIDO2) for privileged accounts	Zero MFA exceptions by June 1
Approve service account credential rotation policy (90-day max)	Execute credential rotation; migrate to managed identities where platform supports; eliminate shared credentials	First rotation complete by June 15

2. Configuration Drift

"Your environment is not what you deployed. It's what it became."

Configuration drift is the distance between what was designed and what is running. A penetration test on day one shows a clean environment. By day 180, firewall rules have accumulated, conditional access policies have been loosened for traveling executives and never re-tightened, security group memberships have expanded as project teams formed and disbanded, and endpoint protection configurations have diverged as individual machines missed updates.

Consider a firm that deployed a conditional access policy requiring MFA for all users eighteen months ago. Today it may have fourteen exemptions, three undocumented, two applying to repurposed service accounts, and one exempting an entire IP range because a satellite office had connectivity issues during onboarding. None of this appears in the monthly security summary. AI-powered adversaries don't need to be lucky—they need to be thorough. A model that can enumerate your environment's actual configuration will find every deviation, every orphaned rule, every policy gap that drift created.

The question worth asking

If you asked your MSP right now to produce a diff between your environment's current state and its designed baseline—every firewall rule, every conditional access policy, every security group membership—could they do it? This isn't a gotcha question. It's a diagnostic one. If the answer is yes, you have a strong foundation. If the answer is no, that's the single most important gap to close before July.

Partnership Action Framework: Configuration Drift

Drift is a shared failure. Neither party can solve it alone. The following actions establish a continuous drift detection and remediation operating model:

The Firm Owns	The MSP Owns	Target Timeline
Approve establishment of a documented configuration baseline as the authoritative reference	Build and deliver environment baseline: firewall rules, CA policies, security groups, endpoint configs, DNS, GPOs; store as versioned artifact	Baseline delivered by May 30
Participate in quarterly drift review; approve or reject identified deviations	Deploy automated drift detection (e.g., Azure Policy, SentinelOne Ranger, config-as-code diffing); generate monthly drift reports against baseline	Detection tooling live by June 15
Designate single internal owner for configuration change approval	Implement change logging with mandatory baseline reference; flag any change without an associated ticket or approval	Process in place by May 15
Fund annual configuration re-baselining engagement	Conduct full re-baseline annually; integrate drift metrics into quarterly security review scorecards	First re-baseline scheduled for Q3 2026

3. Application and Data Sprawl

"Sprawl isn't just more apps. It's more seams."

The typical alternative investment firm operates a technology stack that would have been unrecognizable a decade ago: Bloomberg, portfolio management, OMS/EMS, risk analytics, compliance monitoring, CRM, investor portal, fund accounting, collaboration platforms, cloud storage—and now, AI tools layered on top. The real exposure is not in the applications themselves but in the connective tissue between them. API keys stored in configuration files. OAuth tokens with no expiration. Service accounts bridging systems with permissions inherited from initial setup and never scoped down.

Shadow AI has accelerated this sprawl dramatically. As I outlined in Full *Throttle, Full Control*, the adoption of AI tools by investment professionals has naturally outpaced the governance structures designed to contain them. None of this is surprising—the tools are compelling and the business pressure to adopt is real. But each ungoverned AI interaction adds another seam to an already complex attack surface.

AIMB has a structural characteristic that amplifies sprawl risk: vendor concentration. When hundreds of firms use the same portfolio management platform, the same fund administration system, and the same OMS/EMS provider, a vulnerability in any of those systems is not one firm's problem. It is a systemic event—exactly the dynamic the SolarWinds compromise demonstrated.

Partnership Action Framework: Application & Data Sprawl

Sprawl is a discovery problem first and a governance problem second. You cannot secure what you cannot see:

The Firm Owns	The MSP Owns	Target Timeline
Disclose all known SaaS, AI tools, and third-party integrations; mandate disclosure of new tool adoption	Conduct application and data flow discovery: enumerate all SaaS integrations, API connections, OAuth grants, and service accounts across the managed environment	Discovery complete by May 30
Adopt shadow AI acceptable use policy; communicate to all staff	Deploy shadow AI detection: monitor DNS, proxy logs, and endpoint telemetry for unauthorized AI tool usage (ChatGPT, Gemini, Copilot, etc.); report monthly	Detection active by June 1
Approve API/OAuth governance framework: no non-expiring tokens, mandatory scoping, documented purpose for every integration	Audit all existing API keys and OAuth tokens; revoke unused grants; enforce 90-day token rotation; implement centralized API key vault	Audit complete by June 15; vault by July 1
Designate data classification owner; approve classification taxonomy	Implement DLP policies aligned to classification; detect and alert on sensitive data movement across integration seams	DLP policies active by June 30

4. Vendor Supply Chain Risk

"You are only as secure as the least secure vendor in your stack."

A 50-person hedge fund does not write software. What it does have is a stack of 15 to 25 vendor dependencies—and essentially zero visibility into whether those vendors are patching at the speed that Mythos now demands. Consider the typical AIMB technology stack. The fund admin platform likely runs on software that has not had a Glasswing-level review. The OMS vendor may be a mid-sized company patching on a monthly cycle. The investor portal may have been built by a fintech startup without a dedicated security team. And in most cases, none of these vendors are contractually obligated to disclose how fast they patch.

The SEC's cybersecurity disclosure rules and DORA's ICT third-party risk framework both demand that firms demonstrate oversight of their critical vendors' security posture. The EU AI Act's next enforcement phase takes effect August 2, 2026, with penalties up to three percent of global revenue. The regulatory direction is clear: firms are expected to demonstrate vendor oversight. The window for treating vendor security as someone else's problem is closing.

Partnership Action Framework: Vendor Supply Chain

Neither the firm nor the MSP can assess every vendor alone. This framework distributes the work by leveraging the MSP's technical depth and the firm's contractual relationships:

The Firm Owns	The MSP Owns	Target Timeline
Classify vendors into Tier 1 (critical: PMS, OMS, fund admin, cloud) and Tier 2 (supporting) based on data sensitivity and business dependency	Develop vendor security assessment questionnaire aligned to SEC disclosure rules, DORA ICT requirements, and AI-era patching expectations	Classification and questionnaire ready by May 15
Issue vendor security questionnaire to all Tier 1 vendors; escalate non-responders	Conduct technical assessment of Tier 1 vendor attack surface: external scanning, certificate hygiene, DNS posture, public vulnerability history	Tier 1 assessments complete by June 15
Negotiate contractual patch SLA requirements into renewals: Tier 1 vendors must commit to 72-hour critical patch deployment	Monitor Tier 1 vendor patching cadence; flag vendors that miss SLAs; provide quarterly vendor risk scorecard	Contractual language drafted by June 1; scorecards by Q3
Allocate budget for vendor risk management tooling (e.g., SecurityScorecard, BitSight)	Deploy continuous vendor risk monitoring; integrate alerts into SOC workflow; escalate material changes in vendor posture	Tooling evaluation by June 15; deployment by July 1

5. Operational Tempo

"The patching process that worked last year won't survive the next six months."

The most operationally consequential implication of Mythos is the compression of exploitation timelines. The window between vulnerability disclosure and weaponized exploitation has shrunk from weeks to hours. Threat actors are already using AI to reverse-engineer patches within 72 hours of publication. When the Glasswing partners publish their findings in July, firms that haven't expanded beyond a monthly patching cadence will face a significant challenge absorbing that wave.

For alternative investment managers, patching is also a business continuity problem. Reboots mean downtime. Downtime means trading systems offline and the front office calling the CIO. Many firms have understandably built a culture around scheduled maintenance windows. But when exploitation timelines compress from weeks to hours, the cadence needs to evolve with the threat. The goal isn't chaos—it's building the operational muscle to move quickly when the situation demands it.

The solution is not to eliminate change management. It is to build emergency change paths that are pre-approved, pre-tested, and ready to execute without convening a change advisory board. When a critical vulnerability is disclosed at 2:00 PM on a Tuesday, the response should not be a meeting. It should be a pre-defined runbook that executes immediately, with automated rollback in case of failure.

Partnership Action Framework: Operational Tempo

Speed without structure creates chaos. The following actions build the emergency change paths that enable machine-speed response while preserving accountability:

The Firm Owns	The MSP Owns	Target Timeline
Pre-approve emergency change protocol: authorize MSP to deploy CVSS 9.0+ patches within 72 hours with notification (not approval)	Build pre-tested emergency patching runbooks with automated rollback for each major platform (Windows, network devices, cloud services)	Protocol signed by May 15; runbooks by June 1
Expand maintenance windows: authorize weekly (not monthly) patching windows for the next 90 days to absorb Glasswing-driven CVE volume	Implement automated patch testing pipeline: stage, test, validate, deploy with zero-touch where possible; report outcomes within 24 hours	Pipeline operational by June 15
Designate 24/7 emergency contact for out-of-cycle critical patches; authorize phone-based approval for CVSS 10.0	Establish emergency patching SLA: CVSS 10.0 / active exploit = deploy within 4 hours; CVSS 9.0+ = 72 hours; CVSS 7.0+ = 7 days	SLA adopted by May 15
Fund tabletop exercise simulating July Glasswing disclosure wave	Design and facilitate tabletop exercise: 15 critical CVEs disclosed simultaneously across OS, browser, firewall, and hypervisor; test response time, communication, and rollback	Exercise scheduled for June, pre-July

The New Paradigm: From Vendor Relationship to Security Partnership

The traditional MSP model was built for a world of periodic threats and scheduled maintenance. The client purchased a service; the MSP delivered it. Accountability flowed through SLAs measured in uptime and ticket resolution. That model cannot survive a world where AI-accelerated adversaries discover and weaponize vulnerabilities faster than any monthly patching cycle can address, and where the threat landscape compounds in complexity every quarter as the next generation of models arrives.

What's needed is not a better vendor relationship. It's a genuine security partnership—one where both parties share ownership of outcomes, not just activities. The following framework reflects the operating model we believe the post-Mythos environment demands.

The Firm Owns	The MSP/MSSP Owns	Shared Accountability
Risk appetite, policy decisions, and budget allocation for remediation	Continuous posture monitoring, drift detection, and configuration enforcement against agreed baselines	Joint quarterly security reviews with shared scorecards measuring outcomes, not just operational SLAs
Timely approval of emergency patches, reboots, and change requests	Pre-built emergency change runbooks with automated rollback, ready to execute within hours of critical disclosure	Exception lifecycle management with mandatory expiration, shared review cadence, and mutual sign-off
Educating staff on security hygiene and acceptable use policies	Application and data flow inventory maintained as a living asset, with shadow AI detection	Shared vendor governance—MSP assesses and monitors; the firm enforces contractual SLAs
Honest communication about business constraints and risk tolerance	Transparent reporting on drift, exceptions, patching coverage—not just green dashboards	Joint roadmap planning aligning security investment with the evolving threat landscape, reviewed semi-annually

Maturity Self-Assessment:
Where Does Your Firm Stand?

The five structural weaknesses described in this paper do not affect every firm equally. The following maturity model provides a framework for self-assessment across each domain. Most AIMB firms will find themselves in the Reactive or Structured column for the majority of domains—this is expected and is not an indictment. The purpose is diagnostic: identify where you are, decide where you need to be by July, and prioritize the partnership actions accordingly.

Domain	Reactive	Structured	Adaptive
Hygiene & Exceptions	Exceptions tracked informally or not at all; no expiration policy; MFA gaps exist for some users or service accounts	Exception register maintained with annual review; MFA enforced for interactive users; service account inventory exists but rotation is inconsistent	Automated exception lifecycle with mandatory expiration; zero MFA exceptions; managed identities replacing static credentials; continuous compliance monitoring
Configuration Drift	No documented baseline; drift discovered only during incidents or annual audits; changes not consistently tracked	Baseline documented but updated infrequently; manual drift reviews quarterly; change logging in place but not enforced for all systems	Versioned baseline as code; automated drift detection with alerting; monthly drift reports; deviations require documented justification and expiration
App & Data Sprawl	No centralized inventory of SaaS or integrations; shadow AI usage unmonitored; API keys managed ad hoc	Application inventory exists; shadow AI policy communicated; API key audit conducted annually; DLP policies in place for email only	Continuous discovery of SaaS and AI tools; centralized API key vault with rotation; DLP across all egress paths; data classification enforced
Vendor Supply Chain	Vendor security assessed only at onboarding; no contractual patch SLAs; no ongoing monitoring of vendor posture	Tier 1 vendors identified; annual security questionnaire issued; vendor risk discussed in security reviews but not scored	Continuous vendor risk monitoring (SecurityScorecard/BitSight); contractual patch SLAs enforced; quarterly vendor scorecards; Tier 1 vendor incident response plans tested
Operational Tempo	Monthly patching cycle; all patches require CAB approval; no emergency change path; patch compliance measured quarterly	Bi-weekly patching for critical systems; emergency change process exists but rarely tested; patch compliance measured monthly	Risk-tiered patching SLAs (4hr / 72hr / 7d); pre-approved emergency runbooks with automated rollback; continuous patch compliance monitoring; tabletop exercises conducted semi-annually

How to use this assessment: Score your firm honestly in each row. Any domain in the Reactive column is a P0 priority. Firms should target Structured across all five domains by July 2026. Adaptive is the long-term operating model—the firms that reach it first will have a durable advantage in a permanently accelerated threat landscape.

July Readiness Checklist

The Glasswing consortium will publish its findings in early July 2026. The patch cycle that follows will stress-test every assumption about operational tempo, change management, and vendor responsiveness. The following checklist provides a concrete, prioritized set of actions—each with explicit ownership and success criteria—to complete before the Glasswing disclosure.

Priority key: **P0** Complete by May 15 **P1** Complete by June 15 **P2** Complete by June 30

☐	Action	Owner	Priority	Success Criteria
☐	Complete full exception audit and deliver categorized inventory with risk scores and expiration dates	MSP	P0	100% of active exceptions documented with assigned owners and expiration dates
☐	Sign emergency patch protocol: MSP authorized to deploy CVSS 9.0+ within 72 hours with notification	Joint	P0	Signed protocol on file; designated 24/7 emergency contacts confirmed
☐	Achieve zero MFA exceptions across all interactive accounts including VPN, RDP, and admin consoles	MSP	P0	Zero exceptions verified in Entra ID reporting; phishing-resistant MFA on all Tier 0/1 accounts
☐	Classify all managed assets into Tier 1/2/3 patching priority based on exposure and business criticality	MSP	P0	Tiering document delivered and acknowledged by firm
☐	Establish documented configuration baseline for firewall, CA policies, security groups, and endpoint configs	MSP	P1	Versioned baseline artifact delivered; automated drift detection deployed
☐	Complete Tier 1 vendor security assessments and issue patch SLA requirements	MSP	P1	Tier 1 vendor scorecards delivered; contractual patch SLA language drafted for next renewal
☐	Deploy shadow AI detection and conduct full application/integration discovery	MSP	P1	Discovery report delivered; shadow AI monitoring active; unauthorized tools flagged
☐	Audit all API keys, OAuth tokens, and service principals; revoke unused; enforce 90-day rotation	MSP	P1	Audit report delivered; zero non-expiring tokens in Tier 1 systems
☐	Build and test emergency patching runbooks with automated rollback for OS, network, and cloud platforms	MSP	P1	Runbooks documented and tested in staging; rollback validated

☐	Action	Owner	Priority	Success Criteria
☐	Conduct tabletop exercise simulating Glasswing-scale multi-platform CVE disclosure	Joint	P2	Tabletop completed; after-action report with improvement items delivered
☐	Deploy continuous vendor risk monitoring (SecurityScorecard, BitSight, or equivalent)	Joint	P2	Platform live; Tier 1 vendors monitored; alerts integrated into SOC workflow
☐	Update incident response plan with AI-driven attack scenarios and compressed response timelines	Joint	P2	Updated IRP reviewed and approved; distributed to all stakeholders

This Only Gets Harder

The AI capability curve in cybersecurity is compounding. OpenAI classified its GPT-5.3-Codex as the first model to reach high-capability status for cybersecurity tasks. Google, Meta, and leading Chinese labs are pursuing similar capabilities. Open-weight models are closing the gap faster than anyone predicted. The challenges described in this paper will not stabilize—they will intensify in complexity, speed, and sophistication.

The partnership model above cannot be a one-time implementation. It must be a living operating rhythm that evolves as the threat landscape evolves. The firms and MSPs that build this muscle now will compound their advantage over time. Those that treat it as a project with a completion date will find themselves re-learning the same lessons at an accelerating cost.

The Partnership Test

Ask your MSP three questions: Can you show me the actual state of my environment versus its designed baseline—today, not last quarter? Can you execute an emergency critical patch within 24 hours without a change advisory board meeting? Can you tell me the patch SLA and last security assessment date for every Tier 1 vendor in my environment? If the answers are yes, you have a partner. If not, you have a starting point for the conversation that needs to happen before July.

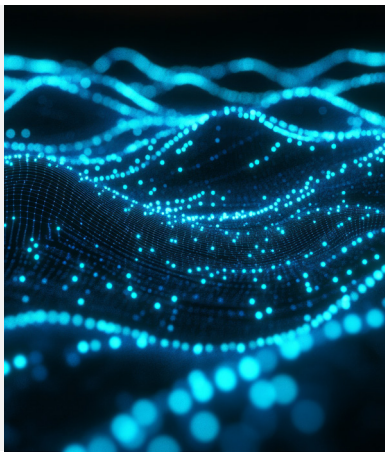

The Clock is Running

The Glasswing report drops in July. The patch cycle that follows will test every assumption that firms and MSPs have made about operational tempo, change management, and risk tolerance. Between now and then, the industry has a window—measured in weeks, not months—to strengthen these foundations. Audit your exceptions. Measure your drift. Map your attack surface. Assess your vendors. Accelerate your patching. And sit down with your MSP—not to issue demands, but to build the shared operating model that the post-Mythos world requires.

The cracks were always there. Mythos just proved that something is now fast enough, smart enough, and thorough enough to find every single one of them. The good news: every weakness in this paper is fixable. The firms and MSPs that build true partnerships now will turn this moment into a durable competitive advantage. Not just for one quarter. For the permanently accelerated world we all now operate in.

About the Author

Rich Itri is the Chief Innovation Officer at ECI, serving approximately 750 alternative investment management and boutique financial services firms globally. With 28+ years as a CTO and CIO at hedge funds and investment banks, Rich leads ECI's ELLA platform strategy. This is the second paper in a series; the first, Full Throttle, Full Control, was published in March 2026.

ECI

US: +1 800 752 1382

Singapore: +65 6622 2345

UK: +44 207 071 6802

Hong Kong: +852 3189 0101

eci.com

This paper provides an independent assessment of enterprise AI platforms for alternative investment managers and boutique financial services firms. It is not a product pitch. It is a decision framework built on real-world deployment experience across 750+ firms managing complex regulatory and security requirements.