Internal Audit Collective

Adaptive Risk: A Framework built for today's agentic threat landscape

Assurance designed to close the gaps created by adversaries moving at machine speed.

David Malcom
Cybersecurity Practice Leader / Field CISO
InBalance

Date Oct 29, 2026
www.InternalAuditCollective.com

Internal Audit Collective

Engagement Requirements

Engage Actively:
Participate actively by responding to all polling questions.

Technical Preparedness:
Ensure your device is set up with a stable internet connection, functioning audio, and video capabilities if required.

Stay Engaged:
To receive the credited hour, you must remain on the call for the entire duration of the webinar.

Provide Feedback:
Complete the post-webinar surveys to share your insights and improve future sessions.

Join the Internal Audit Community Collective Today

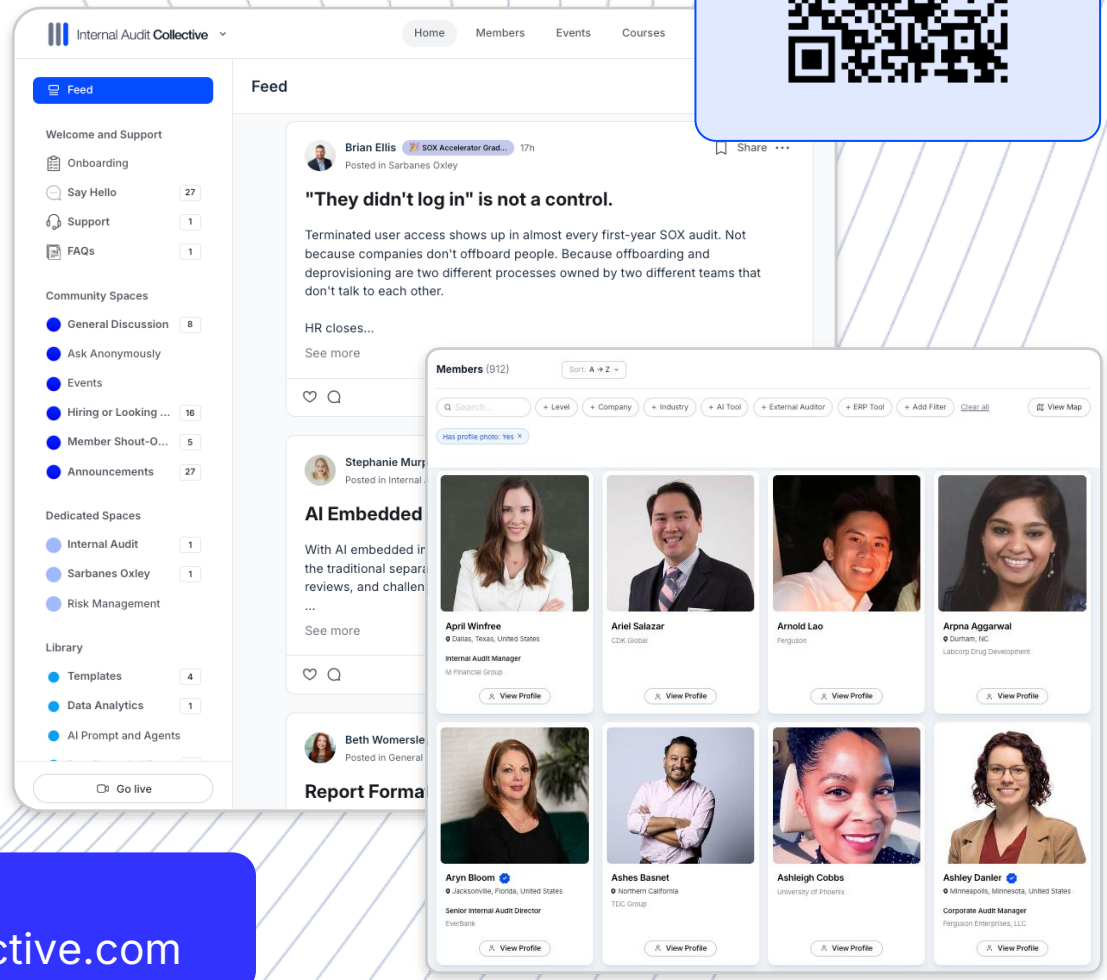
Join the Collective

Use Code WEB1029 for 20% off

Join an online, "managed" community for SOX and Internal Audit professionals

The Collective is a membership site designed for SOX and Internal Auditors to network, share, learn, and grow.

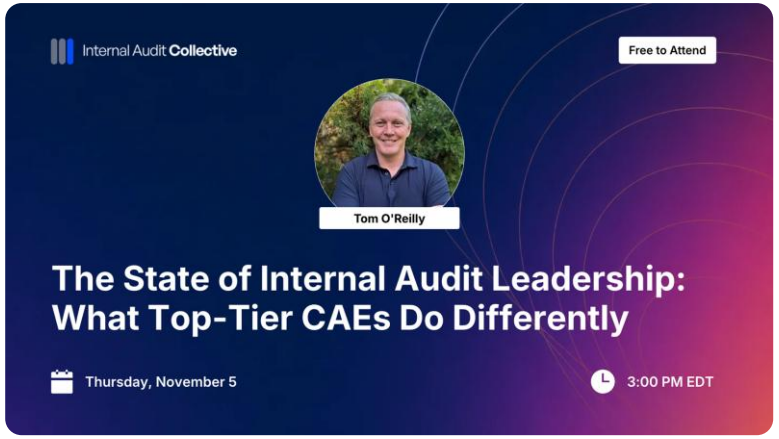
Get access to an online resource to ask questions, get feedback, and exchange ideas with your peers.

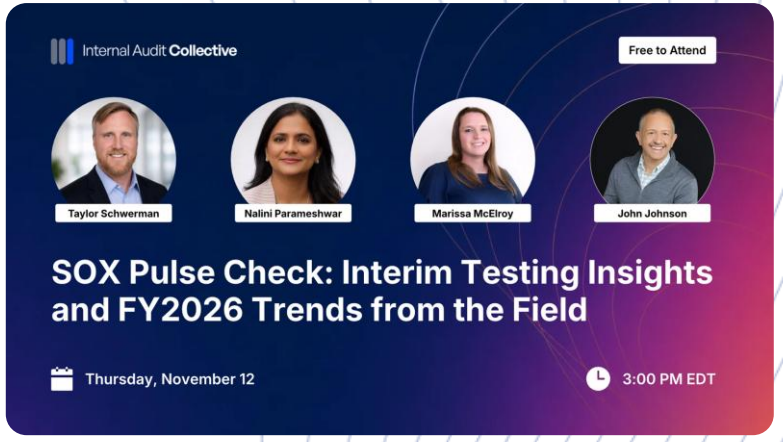


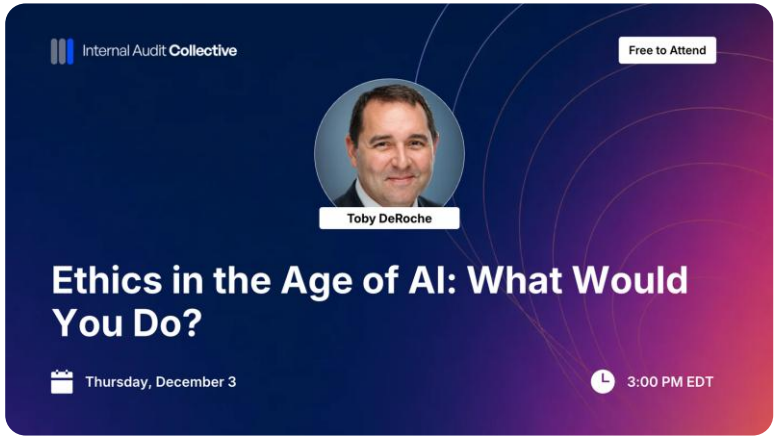
Join today at:
www.InternalAuditCollective.com

Internal Audit Collective

Join Us For These Upcoming Webinars









Internal Audit **Collective**



President and CEO

malcom risk advisors.

InBalance

David Malcom is a seasoned executive with over 26 years of experience in cybersecurity, internal audit, technology consulting, and data analytics. In 2024, with a mission to make cybersecurity services available and to small and medium sized businesses (SMBs), David founded **Malcom Risk Advisors.** In October 2025, Malcom Risk Advisors formed a working partnership with **In Balance IT Solutions**, a Midwest-based technology services provider. As a **Managing Partner and Cybersecurity Practice Leader**, David is responsible for the company's go-to-market cybersecurity strategy and offerings, and as **Field CISO**, oversees all client service delivery.

David's prior executive leadership experience includes:

- **Chief Audit Executive** at **Groupon**
- **Chief Information Security Officer** at **Hyatt Hotels Corporation**
- **Global IT Audit and Data Analytics Leader** at **Accenture**
- **Managing Director, SVP Digital Risk and Data Analytics** at **The Northern Trust Company**.

A creative thought leader, David was named the **2023 Auditor of the Year** by the Institute of Internal Auditors (Chicago Chapter) in recognition of his use of data-driven strategies, automation, and emerging technologies to lead future-focused internal audit teams. He frequently speaks at conferences, podcasts, and webinars, where he discusses today's cyber threat landscape, digital transformation, emerging technologies, and security metrics that truly define 'what good looks like'.

David is a Certified Information Systems Auditor (**CISA**) and a Certified Internal Auditor (**CIA**). He holds a **Bachelor's Degree in Accounting** from Illinois Wesleyan University and recently obtained his **Masters Degree in Biblical Preaching** at Moody Bible Institute. He serves as a **Resident Pastor** at Renew Church in Westmont, IL.

Internal Audit Collective

Today's Agenda

Topics to Discuss

01

The Adaptive Risk Imperative

Why static control frameworks leave assurance gaps and how Adaptive Risk closes them.

02

Pillar I: Threat Driven Metrics

Leveraging continuously monitored Key Risk Indicators (KRIs) to detect material changes in your organization's cyber risk posture.

03

Pillar II: Dynamic Audit Priorities

Breached KRIs activate dynamic audit activities in response to escalating threats.

04

Pillar III: Control Resilience Testing

Audit testing extends beyond legacy control frameworks to validate the ability of your controls to hold up against agentic adversaries.

05

Adaptive Risk in Action

Examples of metrics that identify emerging cyber risks, options for obtaining additional coverage, and procedures to test hardened controls.

SECTION I

The Adaptive Risk Imperative

*Control frameworks **weren't designed to withstand adversaries moving at machine-speed**. **Adaptive Risk** provides Internal Audit with the tools necessary to identify and **address the threats that lay undetected.***

Internal Audit Collective

Adversaries Are Operating at Machine Speed

The window between initial access and business impact has collapsed from days to minutes

29 min

Average eCrime breakout time in 2025, 65% faster than 2024

CrowdStrike 2026

27 sec

Fastest observed breakout time – time from initial breach to data access

CrowdStrike 2026

48 hrs

Typical time-to-exploit after vulnerability disclosure, down from about a week

Fortinet 2026

42%

Of exploited vulnerabilities where the attack occurred before public disclosure

CrowdStrike 2026

CYBERCRIME HAS INDUSTRIALIZED

122B

exploitation attempts have increased 25% in 2025

4.62B

stolen user credentials sold on the dark web, up 79%

7,831

ransomware victims, up from about 1,600

89%

increase in AI-enabled adversary operations

Sources: CrowdStrike 2026 Global Threat Report; Fortinet FortiGuard Labs 2026 Global Threat Landscape Report

Internal Audit Collective

The Mythos Inflection Point

AI capabilities are used maliciously to automate large portions of the cyber kill chain; what once required coordinated human efforts now run as autonomous workflows across your environment.

Thousands of high-severity vulnerabilities found in every major operating system and browser within weeks, **one missed by experts for 27 years**

83%+ of vulnerabilities produced a working exploit on the first attempt

4.7 months time by which the top AI models' hacking ability doubles, down from 8 months in late 2025

2027 when a Mythos-class capability is expected to be readily available to all classes of adversaries

ARE SECURITY LEADERS READY?

4.0 / 10



Average self-rated preparedness for a Mythos-level threat today (52 CISOs)

6.7 / 10



Expected preparedness 12 months from now; 85% rate themselves 5 or below today

Sources: AIUC-1 Consortium, After Mythos: Defending at Machine Speed (2026), citing UK AI Security Institute

Controls Remain Calibrated for Human-Speed Attackers

Reliance on manual detection and response introduces ‘fragility’ into the organization. In other words, it’s the latency created between the moment a breach occurs and when full recovery is accomplished

Typical control parameter	Threat reality in 2026	Resulting assurance gap
90-day critical patch SLA	Exploits are weaponized within 24 to 48 hours, with patches being reverse-engineered into exploits	Months of exposure for a control deemed ‘compliant’
Quarterly access reviews	29-minute average breakout using valid user credentials	Compromised access persists between review cycles
Annual vendor questionnaire	Mid-contract, vendors enable new AI features that touch your data	Within months, vendor risk rating is stale and irrelevant
Analyst triage measured in hours	Agentic kill chains complete in under two minutes	Response begins long after the damage has already been done
Signature-based detection	49% of suspicious activity abuse legitimate admin tools	Attacks appear in audit logs as normal operations
Annual penetration test	AI-driven discovery continuously surfaces new zero-day exploits	Results are out of date before the report is issued

Sources: AIUC-1 Consortium 2026; CrowdStrike 2026 Global Threat Report; Fortinet 2026; In Balance Human/AI SecOps Operating Model

Internal Audit Collective

AI Now Operates Inside Your Perimeter, At-Scale

Shadow AI, autonomous agents, and non-human identities are the fastest growing attack surface

Shadow AI

43%

of breached organizations were impacted in part by shadow AI use, up from 20%

$5.39M

average cost of a data breach caused by shadow AI

AI Agents and NHIs

65%

of organizations run AI agents in production

80%

Lack the ability to consistently explain why an NHI took a privileged action

AI-Driven Attacks

56%

increase in AI-driven attacks year over year

92%

of AI-related breaches occurred due to insufficient AI access controls

$4.99M average cost of a data breach, up 12%. **16 min** for red teamers to breach most enterprise AI systems tested.

Sources: IBM Cost of a Data Breach Report 2026; AIUC-1 Consortium 2026; Delinea 2026 Identity Security Report; Zscaler ThreatLabz 2026

Internal Audit Collective

What an Adaptive Defense Posture Looks Like

Six capability pillars, governed by a judgment layer that defines where machines act and when humans decide

Identity Security

Temporary access and alerts for signals of misused access by staff and agents

Data Security & Shadow AI

Locate sensitive data, label it, block leaks, and control what gets pasted into AI tools

Cloud Security

Identify and fix cloud misconfigurations and stop attacks on live systems

Judgment Layer

Governance that defines the actions that run autonomously, the ones requiring human approval, and the ways in which these boundaries are tested (Human + AI Operating Model)

Outcome: defensive capabilities that match adversary speed

Application and API Security

Security embedded into software development, AI code review, and identifying every API

Agentic SOC

AI-driven alert triage, investigation, and containment

Agentic Security (Secure AI)

Complete agent inventory, guardrails, and live monitoring of agents' activity

Sources: In Balance Adaptive Defense framework; SACR, Unified Agentic Defense Platforms (2026)

Reshaping Internal Audit With Forward-Looking Risk Intelligence

From Assurance to Intelligence: the Future of Internal Audit - *key takeaways from Protiviti's 2026 white paper.*

Oversight → **Hindsight** → **Insight** → **Foresight**

Where most IA functions operate today: hindsight-oriented, validating past performance through sample-based testing.

Where IA must go: providing risk and decision intelligence in fast-moving environments

Internal Audit's New Mandate

Risk and Decision Intelligence

Internal Audit is expected to provide quality, relevant, risk context - grounded in independent judgement - at or before the moment decisions are made, rather than after the fact.

Core assurance, independence, and objectivity remain essential; earlier engagement is simply added to the list.

1 Control testing is changing

Intelligent automation materially changes how time is spent, with traditional "look back" testing becoming a thing of the past

2 Technology is required

As AI becomes an imperative, increased breadth, depth, timeliness, and quality of audit coverage is the expectation

3 Talent models are shifting

Judgment, skepticism, and credible communication are now core to Internal Audit's identity.

4 Stakeholders expect more

Boards and executives place greater emphasis on forward-looking, timely risk and decision insights.

Source: Protiviti, From Assurance to Intelligence: The Future of Internal Audit (2026)

Internal Audit Collective

Putting Vision Into Practice: Adaptive Risk

Machine-speed defense requires real-time, threat-informed assurance; a fundamental shift in IA's operating model

ADAPTIVE DEFENSE (INFO SEC)

- Continuous telemetry and risk posture signals
- Defenses reprioritized in real-time as risks shift
- Controls hardened to withstand current adversary techniques

ADAPTIVE RISK (INTERNAL AUDIT)

Pillar I

Threat-Driven Metrics

KRIs provide Internal Audit with access to real-time risk signals, enabling IA to act before decisions are made

Pillar II

Dynamic Audit Priorities

Audit plans flex in response to breached KRI thresholds, prioritizing the audit activity aligned to the escalating threat

Pillar III

Control Resilience Testing

Full populations of data are used in audit testing, validating whether controls will hold up against agentic threats.

Internal Audit Collective

Poll 1: How often does Internal Audit perform a cybersecurity risk assessment?

1. Annually
2. Quarterly
3. More Frequently than Quarterly
4. We don’t perform a cyber risk assessment we rely on someone else to do this for us.

Sponsored By:
SYNERGY

LEARN MORE

SECTION II

Pillar I: Threat Driven Metrics

Continuously monitored KRIs detect material changes in your organization's cyber risk posture and reports signals in real-time.

Internal Audit Collective

What Makes a KRI "Threat-Driven"

Five tests that identify true early-warning indicators

1. Leading

Rises before an incident occurs, such as stale privileged credentials or unpatched KEVs

2. Threat-linked

Maps to a known technique adversaries use today (MITRE ATT&CK, MITRE ATLAS)

3. Threshold-bound

Defined tolerance that triggers a pre-defined audit response

4. Automated

Pulled via API from source systems, data refreshed monthly or in real-time

5. Owned

Accountable first-line owner and agreed data definitions

> *Defensive velocity, measured through time to detect, time to contain, and time to revoke compromised credentials, should be treated as a primary business risk indicator.*
>
> FortiGuard Labs, 2026 Global Threat Landscape Report

Internal Audit Collective

Where the Data Lives

Most KRIs already exist in tools your CISO owns; audit needs read-only access and agreed metric definitions

Identity
IdP, PAM, ITDR

Data and AI
DSPM, DLP, AI gateway

Endpoint and Vulnerability
EDR, scanners, EPSS and KEV feeds

Cloud
CSPM, CNAPP

Third Party
TPRM, security ratings

SOC
SIEM and SOAR case data

UNIFIED KRI LAYER

Continuous Controls Monitoring platform or Risk Operations Center

- One asset inventory for every device, identity, and agent
- Correlated, inspectable data tied to owners
- Trends over time measured against thresholds
- Automatic mapping to NIST CSF, ISO, and COBIT

AUDIT OUTPUTS

Threshold breach alerts routed to Internal Audit

Quarterly KRI trends reported to the Audit Committee

Full-population evidence available for continuous auditing

Sources: Panaseer, Buyer's Guide to Continuous Controls Monitoring; Qualys, The Rise of Risk Operations Centers (RSAC 2026)

Internal Audit **Collective**

Four Risk Domains, Updated for the Agentic Era

Sample KRIs for each domain (defined by NIST 800-39 model); metrics added in 2026 are flagged "new"

Identity: Human, NHI, and AI Agent

- Privileged accounts without phishing-resistant MFA
- Mean time from termination to account disablement
- NHIs and AI agents with no accountable owner **NEW**
- AI agents holding standing privileged access **NEW**

Data and AI Governance

- AI tools in use vs. approved AI inventory
- DLP violations in AI prompts
- AI systems in production without a risk assessment
- MCP servers outside the approved registry **NEW**

Endpoint, Ransomware, and Resilience

- Endpoints without EDR coverage
- Backup restore test success rate
- Mean time to remediate KEVs on internet-facing assets **NEW**
- Mean time to contain vs. adversary breakout time **NEW**

Third-Party, Cloud, and Regulatory

- Critical vendors without a current assessment
- High-risk cloud misconfigurations beyond SLA
- Missed regulatory reporting (SEC, DORA, NIS2)
- Vendors enabling AI features without review **NEW**

Sources: Adaptive Risk KRI catalog; Wiz State of AI in the Cloud 2026; Delinea 2026 Identity Security Report

Internal Audit Collective

Examples of KRI Thresholds and Triggers

A metric crossing its threshold moves that area into immediate audit review, ahead of the annual plan

Area	What audit tracks	Triggers immediate review when	Severity
Identity	Admin accounts without phishing-resistant MFA	Any admin account with no MFA, or MFA is provided via text or voice codes	**Critical**
	Time to disable accounts after staff leave	Average time to disable of more than 48 hours	**Critical**
Data and AI	AI tools in use compared with the approved list	More than 3 times the approved number of tools	**Critical**
	AI models in production with no risk assessment	More than 25% of models aren't assessed	**High**
Resilience	Time to patch the flaws attackers are likely to exploit on critical internet-facing systems	Over 72 hours for the most likely targets, or over 7 days for the rest	**Critical**
	Successful backup restore tests	None performed in the last 12 months	**Critical**
Third parties	Critical vendors without a current security assessment	No assessment for a vendor that handles personal data	**Critical**
	AI systems without an AI risk assessment (NIST AI RMF)	Any generative AI system using regulated data	**High**

A sample of the highest-impact triggers; the full list is in the Catalog of Cyber Risk Metrics

Internal Audit Collective

Agentic KRIs: Signals Audit Should Be Watching Today

Added risk indicators to detect AI agents and non-human identities operating outside governed boundaries

KRI	Escalation threshold	Why it matters
Agents in production vs. approved agent inventory	Identification of any unregistered agent with production data access	65% of organizations are already running agents in production
Agents and NHIs using long-lived credentials	More than 25% of agent population use static access privileges	35% of agents use long-lived credentials; only 17% use JIT authorization
Privileged NHI actions without an attributable owner	Any action taken on a Tier-1 systems	80% of organizations cannot consistently explain why an NHI took a privileged action
Prompt injection and guardrail violation detections	Upward trend of 25% QoQ, or upon any successful injection	Malicious prompts injected into GenAI tools observed at 90+ organizations
High-impact agent actions without a human approval gate	An autonomous action taken outside the defined boundary (firewall rule change, closing a security ticket, etc.)	Confirms the judgment layer is operating as designed
Kill switch rehearsal for autonomous systems	Not rehearsed in the last quarter	Untested disablement actions often fail under operational stress

Sources: AIUC-1 Consortium 2026; Delinea 2026 Identity Security Report; CrowdStrike 2026; In Balance Human/AI SecOps Operating Model

Internal Audit Collective

Poll 3: What is the biggest inhibitor preventing your team from being able to monitor and report on Key Risk Indicators?

1. IT Won't Provide You the Data You Need
2. IT is Unable to Provide You with Reliable Data
3. Your Team Lacks the Skills Needed to Analyze Data
4. Your Team Lacks the Capacity Needed to Analyze Data

Sponsored By:
SYNERGY

LEARN MORE

SECTION III

Pillar II: Dynamic Audit Priorities

Breached metric thresholds activate a dynamic audit plan that addresses escalating threats

Internal Audit Collective

From Risk Signal to Action

Automated data collection, calculation, and alerting; auditor judgment makes the calls

1 Collect

Automated feeds pull data from security tools:

- Identity and access
- Vulnerability scanners
- Endpoint and cloud tools
- Data loss prevention
- Vendor risk and GRC

2 Calculate

A KRI engine turns raw data into risk signals:

- Normalizes data across tools
- Calculates each KRI
- Trends results over time
- Measures results against thresholds

3 Alert

Automated alert is triggered by a breached threshold:

- Severity set by threshold tier and crown-jewel exposure
- Posts to the audit dashboard
- Assigns a ticket to an auditor
- Notifies the CAE if critical

4 Act

AI agents: pull full-population evidence, re-perform the KRI calculation, draft the finding

Internal Audit: confirm data with the CISO within 5 business days, perform pre-defined audit activity, report results to the Audit Committee

Feedback loop: results update the risk register and thresholds; internal audit spot-checks how each metric is produced

RECOMMENDED RESPONSE BY SEVERITY

CRITICAL	Audit starts within 30 days; Committee chair notified
HIGH	Added to next quarter's plan and next Committee report
MODERATE	Monitor the trend; targeted inquiry with metric owner

Internal Audit Collective

The Static Annual Plan Becomes Threat-Responsive

The annual plan serves as a baseline - flexing as threat-driven metrics change during the year

	Static annual plan		Adaptive Risk plan
Planning	Annual risk assessment sets next year's plan	→	**Annual baseline plus 15 to 20% time reserved for risk-triggered work**
Scope	Locked at the start of the fiscal year	→	**Re-prioritized when KRI thresholds breach**
Risk refresh	Quarterly reviews identify emerging advisory activities	→	**KRI monitoring (monthly or real-time) with defined action triggers**
Findings	Reflect threats at the time the plan was written	→	**Reflect the current threat landscape**
Audit Committee	IA provides assurance on prior-period activities	→	**Forward-looking identification of emerging risk and IA's response**

Internal Audit Collective

Breaches and the Audit Activities They Trigger

When a KRI crosses its threshold, internal audit acts now instead of waiting for the next planning cycle

Risk domain	Threshold breach that triggers audit	Audit engagement activated
Identity: Human, NHI, and AI Agent	Any privileged account without phishing-resistant MFA; any terminated user still active after 24 hours	Privileged Access and Identity Audit, agentic testing of every joiner, mover and leaver
	NEW Any AI agent or system account with production access that is unregistered or has no named owner	AI Agent Inventory, Privilege and Autonomy Boundary Audit
Data and AI Governance	Under 90% of sensitive data stores scanned; AI use at 3x the approved inventory; AI prompt leaks up 15% in a quarter	Shadow AI and Data Classification Audit, testing the AI acceptable use policy
Endpoint, Ransomware, and Resilience	Restore success below 95%; any critical system without a tested recovery exercise in 12 months	Ransomware Readiness and Recovery Audit, observe a live restore of a tier-1 system
	NEW Known-exploits on internet-facing assets open 72+ hours; over 10% of high and critical vulnerabilities open 30+ days	Emergency Patch and Compensating Controls Audit, testing patch exceptions
Third-Party, Cloud, and Regulatory	Vendor AI features enabled without review; critical vendor findings open 90+ days; public exposed buckets with sensitive data	Vendor AI Due Diligence Audit; same-day cloud configuration review for exposed data
	Any security policy exception past its expiration date	Risk Acceptance Audit, results reported to the Audit Committee

Sample thresholds; set final thresholds with management based on risk appetite. Sources: Catalog of Cyber Risk Metrics; Adaptive Risk KRI catalog

SECTION IV

Pillar III: Control Resilience Testing

Audit testing goes beyond what's required by legacy control frameworks, providing assurance that your controls hold up against agentic adversaries

Internal Audit Collective

Three Levels of Control Testing

Resilience testing asks whether a control holds up against the techniques adversaries are using today

LEVEL 1

Design

Does the control exist and address the risk?

Evidence: Policies, configurations, walkthroughs

LEVEL 2

Operating Effectiveness

Did it operate consistently over the period?

Evidence: Samples, logs, tickets

LEVEL 3

Resilience

Would it stop or detect a current adversarial technique?

Evidence: Adversary emulation, full-population analytics, configuration strength tests

Level 3 reference points

MITRE ATT&CK
Enterprise adversary techniques

MITRE ATLAS
Attacks on AI and ML systems

OWASP Agentic Top 10
Agent hijacking, AI tool misuse, excessive agency

CISA KEV catalog
Vulnerabilities exploited in the wild

Internal Audit Collective

Leveraging AI to Test With Greater Assurance

Agentic audit tools enable Internal Audit to operate on the same cadence as the risks it assesses

SAMPLE-BASED AUDITING

Sample of 25, walkthroughs, and screenshots, produce a point-in-time report issued months after fieldwork

CONTINUOUS ASSURANCE

100% of the population obtained via API, machine-readable evidence from the source, and exceptions surfaced as they occur

Evidence at scale

Replace sample tests with full data sets ingested via API integrations with IAM, SIEM, and ticketing systems

Cross-framework mapping

A finding is mapped to NIST CSF, ISO 27001, SOC 2, NIST AI RMF, AIUC-1, and relevant regulatory requirements

Anomaly detection

Toxic access, privilege creep, configuration drift, slow-drip data exfiltration – all identified in real-time

Continuous assurance

Live posture updates via dashboards and audit-verified evidence packages produced on-demand

65 days faster, $1.93M lower cost: breach outcomes realized by organizations extensively using AI and automation used within their security program (IBM 2026)

Internal Audit Collective

Moving From Baseline to Adaptive Control Testing

NIST CSF 2.0 baselines, the as-is design and testing gaps, and the hardened control IA should test for

Risk domain	NIST CSF 2.0 baseline	As-is design gap	Threat that creates design gap	Testing of hardened control
Identity: Human, NHI, and AI Agent	**PR.AA-03 Authentication**	Any MFA passes; SMS, OTP and push codes stay phishable	Adversary-in-the-middle kits capture one-time codes and session cookies	FIDO2 or passkeys for privileged and remote users; certificate-based workload identities
	PR.AA-05 Least privilege	Standing roles reviewed periodically; AI agents out of scope	Staff and AI agents hold broad standing access that an attacker inherits instantly	JIT elevation, real-time entitlement review, purpose-bound scopes and kill switch for agents
Endpoint, Ransomware, and Resilience	**DE.CM-01 Network monitoring**	Signature and perimeter rules only; identity activity unwatched	49% of suspicious activity abuses legitimate tools that signatures ignore	Behavioral analytics and ITDR mapped to ATT&CK; AI-assisted triage in near real time
	PR.DS-11 Backups	Production admins can reach and delete backups alone	Ransomware crews target backups first and steal data before encrypting	Immutable, isolated backups with multi-person delete approval; quarterly full restores against RTO
	ID.RA-01 Vulnerabilities	Uniform 30-day patch SLA ranked by CVSS, not exploitation	A shipped patch becomes an exploit roadmap within hours	Hours-to-days SLAs for Tier-1 internet-facing assets; KEV and EPSS feeds; compensating controls identified
Third-Party, Cloud, and Regulatory	**GV.SC-07 Supplier risk**	Point-in-time onboarding review; contracts silent on AI features	Vendors roll-out AI capabilities mid-contract and without risk assessment	Continuous third-party monitoring; contract clauses require notice of new AI capabilities

Sources: NIST Cybersecurity Framework 2.0; AIUC-1 Consortium 2026; Fortinet 2026

Internal Audit Collective

Testing Agents Against Machine-Speed Adversaries

Defining new test procedures for AI agents, the identities they use, and the tools they call

Inventory reconciliation

Reconcile agents found in cloud, SaaS, and network scans to the approved registry, including MCP servers and vendor-embedded agents

Least privilege and JIT

Pull agent entitlements; flag standing privilege and long-lived credentials; verify task-scoped, time-bound tokens

Autonomy boundaries

Confirm payments, deletions, and privilege changes require human approval; attempt an out-of-bound action in a sandbox

Injection resistance

Review red team results for direct and indirect prompt injection via documents, email, and tool outputs; confirm remediation

Attributable logging

Trace sampled agent actions to an owner, an authorizing policy, and a logged reasoning chain

Kill switch

Observe a rehearsal of the single-command disabling of an autonomous system; measure the time it takes for a full stop

Sources: Wiz State of AI in the Cloud 2026; OWASP Top 10 for Agentic Applications; Agentic AI as a Cybersecurity Attack Surface (2026)

Internal Audit Collective

SECTION IV

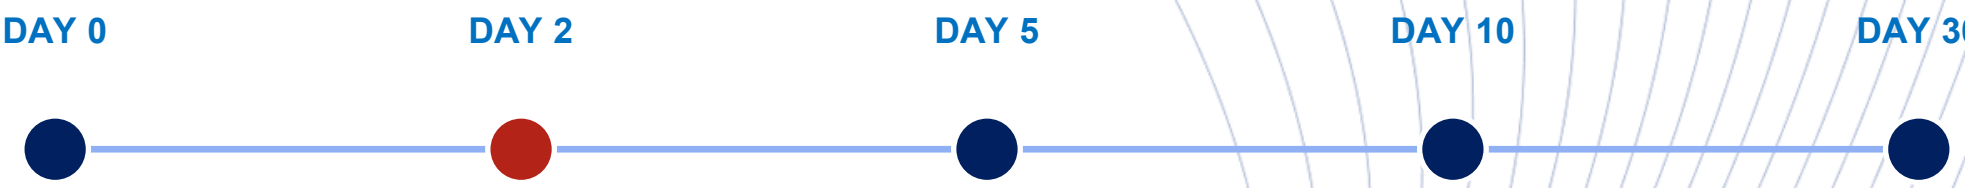
Adaptive Risk in Action

The metrics that identify emerging risks, methods for obtaining additional audit coverage, and extending testing for hardened controls

Internal Audit Collective

Scenario: A Mythos-Level Mass Disclosure Event

With Adaptive Risk, an agentic threat becomes targeted assurance within 30 days (illustrative)



Risk shift

Vendor patches released to address dozens of AI-discovered vulnerabilities; exploitable code appears within 24 hours

KRI breach

31% of Tier-1 internet-facing assets surpass the 72-hour KEV threshold patching SLA

Validate and triage

IA confirms data with CISO; KRI breach rated Critical due to crown-jewel exposure

Targeted audit

Emergency patch and compensating controls audit replaces a deferred Q4 engagement

Report

Audit Committee briefed; IA recommends Tier-1 patching SLA measurement be moved to hours or days

RESILIENT CONTROL TESTING

Verify compensating controls (virtual patching, isolation) can successfully block a known exploit pattern (in test environment)

Rollback readiness and change record confirmed for emergency patches applied on Tier-1 assets

Vendor contracts reviewed to confirm patch SLA and AI-assisted discovery commitments

Sources: Illustrative scenario based on AIUC-1 Consortium 2026 and Fortinet 2026 findings

Internal Audit Collective

Testing a Hardened Control: A Sample Workpaper

Resilience test of an AI agent's privileged access in a Tier-1 environment

CONTROL
AI agents operate with least privilege and time-bound credentials (PR.AA-05)

THREAT TECHNIQUE
Agent hijacking via indirect prompt injection (OWASP Agentic Top 10)

LINKED KRI
Number of AI agents with standing privileged access

Procedure

1. Pull 100% of agents and service principals from IdP, cloud IAM, and AI gateway by API
2. Reconcile the population to the approved agent registry
3. Flag standing privileged entitlements and credentials older than 30 days
4. In a sandbox environment, attempt to perform an out-of-scope action through a test prompt injection; confirm it is blocked and logged
5. Trigger the kill switch for one agent and record the time it takes to perform a full stop

Test of hardened control

Pass criteria

- Zero unregistered agents with production access
- 5% or fewer agents with standing privilege, each with a documented exception
- Out-of-scope action blocked and SOC alerted
- Kill switch completes within the defined target

Internal Audit Collective

How the 2026 Audit Plan Evolved

Original plan (August 2025), breached KRIs and emerging threats, and the revised plan (October 2026)

ORIGINAL PLAN (AUG 2025)	BREACHED KRI / EMERGING THREAT	REVISED PLAN (OCT 2026)	
Q1 SOX ITGC audit	**MFA bypass KRI** AitM phishing kits surge; push-MFA compromises exceed threshold	Q1 SOX ITGC audit	KEPT
Q1 Annual vendor risk review	**Shadow AI KRI** Unregistered GenAI tools exceed 15% of business units	Q1 Phishing-resistant MFA audit	NEW
Q2 Identity and access audit	**Agent privilege KRI** Agents with standing privileged access found in production	Q2 Identity threat detection audit	EXPANDED
Q2 DLP review	**Patch velocity KRI** Mythos-class disclosure wave; 90-day SLA leaves months of exposure	Q2 Shadow AI governance audit	NEW
Q3 Patch and vulnerability management	**Vendor AI KRI** Tier-1 vendors enable AI features without review	Q3 Patch exposure window audit	REPLACED
Q3 Cloud configuration review		Q3 AI agent privilege audit	REPLACED
Q4 Business continuity and DR test		Q4 Ransomware clean recovery	EXPANDED
Q4 EU AI Act readiness		Q4 Vendor AI monitoring audit	REPLACED

Sources: Illustrative Adaptive Risk plan; KRIs drawn from the cross-domain catalog in Section II

Internal Audit Collective

Poll 2: How many cybersecurity audits does your internal audit department perform each year?

1. Zero
2. One or Two
3. At Least Five
4. It’s determined by continuous risk assessments and includes continuous assurance.

Sponsored By:
SYNERGY

LEARN MORE

Internal Audit Collective

The Adaptive Risk Playbook: 30, 60, 90 Days

Start small - one risk domain, one quarter, and one set of metrics

DAYS 1 TO 30

Define threat-driven metrics

- Partner with the CISO and ERM to select 8 to 12 KRIs within one security domain
- Agree on KRI calculation formulas, definitions, owners, and thresholds
- Obtain read-only access to source data

DAYS 31 TO 60

Metric Mapping and Control Hardening

- Map each KRI to a pre-scoped engagement
- Present reserve hours and rules for use to the Audit Committee
- Rewrite tests to address new risks related to identity, shadow AI, and AI agents

DAYS 61 TO 90

Pilot and report

- Perform the first tests on your hardened crown-jewel controls
- Deliver the first KRI trend and impact report
- Track and report mean time from KRI breach to IA action

Keep the Audit Committee informed on the changes in the threat landscape and what Internal Audit is doing about it. CAE and CISO should carry aligned messages.

Internal Audit Collective

Poll 4: How much capacity do you set aside on your annual audit plan to respond to emerging threats?

1. Less than 5% of capacity
2. Less than 10% of capacity
3. More than 11% of capacity
4. 100% of your capacity is allocated to audits on your plan

Sponsored By:
SYNERGY

LEARN MORE

Internal Audit Collective

Q & A

www.InternalAuditCollective.com

Internal Audit **Collective**

AI & Data Analytics Office Hours (For Collective Members Only)



Hosted by Jim Tarantino
Thursday, Oct 29, 12-1pm EDT



Hosted by Lynn Ranf
Tuesday, Nov 03, 12-1pm EDT



Hosted by Dan Hill
Wednesday, Nov 04, 12-1pm EDT



Hosted by Rich Penfil
Tuesday, Nov 10, 12-1pm EDT



Hosted by Dawn Vogel
Wednesday, Nov 11, 12-1pm EDT



Hosted by Abhinav Malhotra & Sophie Chen
Thursday, Nov 12, 12-1pm EDT

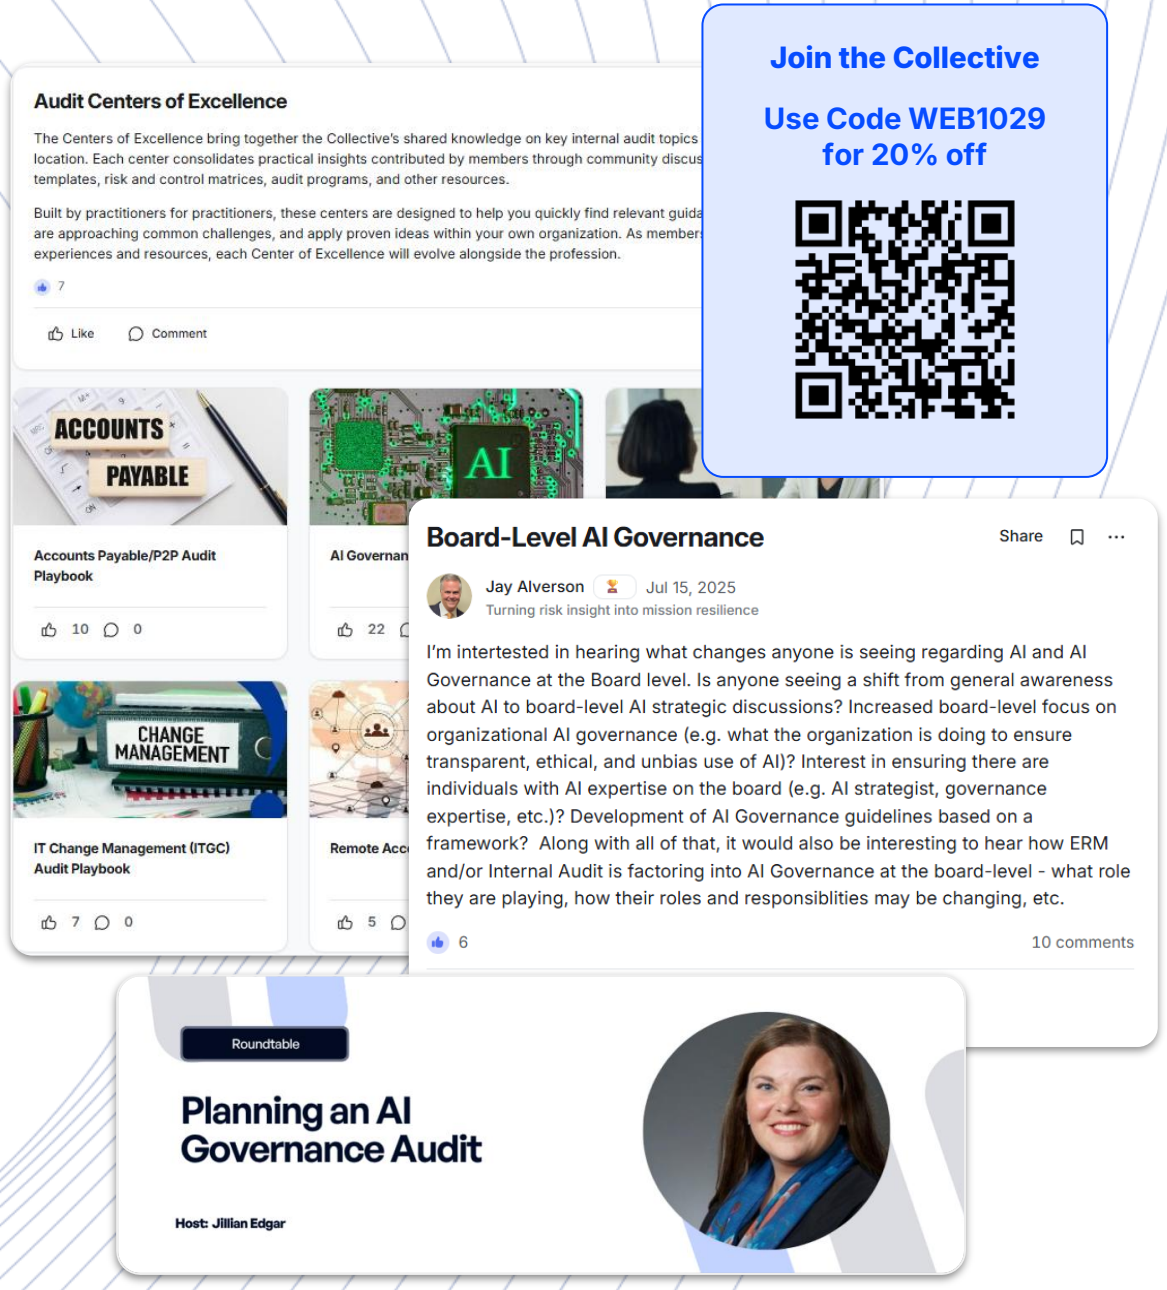
Internal Audit **Collective**

⏰ **Want to join the Collective or register for a course?**
Use coupon code **WEB1008** by Oct 09th for 20% off.

Access to a Library containing AI Prompts, agents, templates, and DA library

Get access to an online resource to ask questions, get feedback, and exchange ideas with your peers.

The Collective is a membership site designed for SOX and Internal Auditors to network, share, learn, and grow.



Explore Our Upcoming Training Programs



Thank You For Participating!