



TOWARDS NEW CYBERSECURITY PERFORMANCE MANAGEMENT

Whitepaper

Measuring success in the
next era of cybersecurity

September 2025

Foreword



Dr. Hesham Altaleb

Chairman, Future of Cybersecurity
Knowledge Community
Saudi Information Technology
Company - SITE

Cybersecurity is entering a new chapter—defined by autonomy, evolving expectations, and greater strategic relevance. As technology transforms how organizations operate, there is a growing need to rethink how cybersecurity performance is measured and aligned with long-term value creation.

We invite business leaders, policymakers, and cybersecurity professionals to engage with the ideas in this report—and to join us in advancing a new generation of cybersecurity performance that empowers innovation, resilience, and long-term value creation.

This white paper reflects the collective effort of the Future of Cybersecurity Knowledge Community, supported by the Global Cybersecurity Forum (GCF). It brings together multidisciplinary expertise to explore how cybersecurity performance measurement must adapt in a world of accelerating innovation and complexity.

Knowledge Community: The Future of Cybersecurity

The Future of Cybersecurity Knowledge Community is committed to exploring the potential opportunities and threats presented by an ever-evolving Cyberspace. By bringing together a diverse array of expertise from various stakeholder groups, it seeks to develop mechanisms that maximize the benefits and address the risks of this new and challenging dimension.

The community welcomes leading technology companies, global cybersecurity organizations, cybersecurity research centers, reputable think tanks, academic institutions, and other stakeholders with a vested interest in exploring and acting upon the future of cybersecurity.

Authors & Contributors

- **Dr. Manar Alohaly**, Saudi Information Technology Company (SITE)
- **Heelah Alraqibah**, Saudi Information Technology Company (SITE)
- **Shoaib Yousuf**, Boston Consulting Group (BCG)
- **Radu Balanescu**, Boston Consulting Group (BCG)
- **Alberto Pardo**, Boston Consulting Group (BCG)
- **Dr. Stefan Deutscher**, Boston Consulting Group (BCG)
- **Duna Alghamdi**, Boston Consulting Group (BCG)
- **Chaimae Haska**, Boston Consulting Group (BCG)
- **Dr. Amira Khattab**, Deloitte
- **Dr. Almerindo Graziano**, CyberRanges
- **Dr. Faisal Sibai**, Accenture
- **Dr. Dan Bogdanov**, Cybernetica
- **Abdulrahman A. Almanea**, Sirar
- **Dr. Antonio Jara**, Libelium
- **Dikmen Edgu**, Axon Partners Group
- **Ed Sleiman**, Microsoft
- **Abdulrahman Almusfir**, Schneider Electric
- **Eilaf Algebail**, Schneider Electric
- **Aaron Ng**, CrowdStrike
- **Sulaiman Almohsen**, National Company of Telecommunications and Information Security (NTIS)
- **Neil Ginns**, International Business Machines Corporation (IBM)
- **Abdulrahman Alosaimi**, International Business Machines Corporation (IBM)
- **Riku Valpas**, Fortinet
- **Richard Schoebel**, ID Quantique
- **Abdullah Almohammed**, Cisco



Contents

Executive Summary	04
1. Cybersecurity Performance Management Stands at a Strategic Turning Point	05
1.1 The rise of agentic AI and autonomy	05
1.2 Cybersecurity is now expected to demonstrate business enablement	06
2. Three Emerging Strategic Priorities	08
2.1 Adapt cybersecurity KPIs to the era of agentic AI and autonomy	08
2.2 Introduce metrics to promote alignment between business and cyber teams	10
2.3 Institutionalize continuous KPI evolution	10
3. Way Forward – A Blueprint to Adapt the Next Era of Cybersecurity	11
Conclusion	14
Appendices	15
Appendix A – NIST CST 2.0	15
Appendix B – Detailed KPI disruption analysis	16
Appendix C – Metrics to support business and cybersecurity synchronicity	21
Bibliography	22

Disclaimer

This document has been published by the Global Cybersecurity Forum (GCF) in collaboration with Knowledge Partners as part of their efforts to promote thought leadership in cybersecurity. While GCF and the knowledge partners have made every effort to ensure the accuracy and reliability of the information provided, neither party assumes any responsibility for errors, omissions, or inconsistencies in the content, nor for any consequences arising from its use or interpretation. The content is provided for general information purposes and may be subject to change without prior notice at the discretion of GCF. This publication is protected by copyright law. No part of this report may be reproduced, distributed, or transmitted in any form or by any means—whether electronic or mechanical—without prior written permission from both GCF and the Knowledge Partners. All requests for such permissions should be directed to KC@GCFForum.org.

Executive Summary

Cybersecurity stands at a strategic turning point—with two transformative shifts redefining the landscape. First, the rapid rise of agentic AI and autonomy is challenging long-standing assumptions about control, oversight, and threat detection. Second, cybersecurity is no longer seen solely as a defensive function—it is increasingly expected to enable business goals, innovation, and organizational agility.

These shifts have exposed fundamental limitations in current performance measurement frameworks. Traditional Key Performance Indicators (KPIs) have been built for an era of manual control and risk mitigation. But they now fail to reflect the speed, scale, and autonomy of modern digital environments. They also often overlook cybersecurity's growing business role in enabling innovation, product integrity, regulatory trust, operational continuity, and market access.

This white paper—developed by the Future of Cybersecurity Knowledge Community—offers a practical blueprint to help organizations modernize cybersecurity performance management, enabling senior leaders to adapt to next-generation AI-driven architectures and evolving business objectives.

Our research reveals three emerging strategic priorities:

- **Adapt cybersecurity KPIs to the era of agentic AI and autonomy:** Cybersecurity KPIs face widespread disruption. Most existing and widely adopted metrics will either require recalibration, become obsolete, or

miss blind spots that call for new KPIs. A strategic analysis of KPIs reveals that Tech & Data and People & Culture metrics will be most disrupted in autonomous-led environments, while governance-linked KPIs remain largely future-proof. New metrics—such as Explainability Score and Autonomy Risk Index—are needed for the effective governance of AI.

- **Introduce future-proof performance metrics to promote alignment between business and cyber teams:** Cybersecurity metrics must now demonstrate business enablement. 85% of CEOs say cybersecurity is critical for business. Metrics such as % of projects with early cybersecurity engagement, and business user satisfaction with cyber support, are critical for tracking cybersecurity's contribution to innovation and trust.
- **Constantly refresh cyber performance metrics & frameworks:** Going forward, cyber performance metrics and frameworks need a constant upgrade. The report introduces a 5-step blueprint to help organizations continuously reassess, evolve, and realign cybersecurity KPIs with AI-driven change and business priorities.

By rethinking how success is measured, organizations can ensure that cybersecurity becomes not just a shield, but a driver of resilience and growth in the next era of autonomy.

1. Cybersecurity Performance Management Stands at a Strategic Turning Point


The performance expectations placed on cybersecurity are evolving rapidly, driven by two powerful and simultaneous shifts. On one side, the emergence of agentic AI is reshaping how organizations operate and make decisions, introducing both opportunity and risk. On the other, cybersecurity is no longer seen as just a protective function, but as a critical enabler of business growth, resilience, and trust.

1.1 The rise of agentic AI and autonomy

AI agents are advancing at an unprecedented pace—with the average length of tasks they can reliably perform doubling approximately every seven months.

Early adopters are already realizing transformative gains: AI-first firms report up to 34× higher revenue per employee, while reducing the cost of high-skill tasks from more than \$100,000 to just \$2,000–\$3,000, and slashing time-to-knowledge from 3 hours to just 5 minutes.

Today, AI Agents can reach '1h' of automation - doubling every 7 months



Source: Measuring AI Ability to Complete Long Tasks arXiv:2503.14493 [cs.AI]; illustrative diagram

1. State of The Art 2. Time taken by human experts is strongly predictive of model success on a given task: current models have almost 100% success rate on tasks taking humans less than 4 minutes, but succeed <10% of the time on tasks taking more than around 4 hours

Figure 1: Rising length of AI autonomy

However, these gains come with significantly added complexity. Agentic systems require access to sensitive data and are increasingly linked to exposure risks, including GenAI-related breaches. These factors contribute to the average breach cost now exceeding USD 4.4 million. These systems can also behave unpredictably, pursue unintended goals, and generate outcomes that are difficult to trace or control. As a result, two-thirds of executives now cite cybersecurity and data privacy as their top GenAI-related concerns.

In this context, cybersecurity measurement frameworks play a critical

role. Unlike ad hoc investments in tools or processes, robust measurement frameworks provide a systematic way to track emerging risks, evaluate the effectiveness of controls, and benchmark resilience across peers and time. By translating complex, machine-led dynamics into quantifiable metrics, these frameworks help leaders prioritize resources, demonstrate accountability, and adjust oversight models as AI autonomy grows. In doing so, they ensure that organizations are not only defending against new threats but also continually improving their ability to manage and govern autonomous systems effectively.

1.2 Cybersecurity is now expected to demonstrate business enablement

Expectations for cybersecurity are shifting from protection to business enablement. Digital infrastructure is becoming central to growth, which means that cybersecurity now directly impacts customer trust, market access, operational continuity, and brand value.


Among CEOs, 85% say cybersecurity is critical for business growth. Additionally, according to the IBM Institute for Business Value, two-thirds of business leaders now view cybersecurity primarily as a revenue enabler, and organizations with more mature security approaches achieve on average 43% higher revenue growth than their less mature peers. However, realizing this potential requires more than intent; it demands measurement frameworks that capture how well cybersecurity supports broader business priorities.

A key barrier is the lack of alignment across cyber and business functions, further referred to as cyber synchronicity. Research highlights a persistent gap between cybersecurity and business priorities—a challenge which affects even leading organizations. Most CISOs reported limited alignment between cyber and business functions, with misalignment most evident in three areas: the halting of major digital initiatives due to security concerns, insufficient understanding of cyber risks and investments among business leaders, and the absence of a business-oriented mindset within cyber teams (see Figure 2). Closing this gap is critical to embedding cybersecurity at the core of enterprise strategy and decision-making.

Synchronicity | Alignment across business and cyber initiatives can act as another dimension to measure the performance of cyber teams

Self-assessment of cyber synchronicity

Q: At your company, to what extent do you agree that:



Source: BCG & GLG CISO Survey (April 2025, Total N = 300)

Figure 2: CISOs self-assessment of business & cyber synchronicity

This white paper tackles the twin challenges of aligning business and cybersecurity, while confronting the disruptive force of AI autonomy. It provides a practical blueprint for modernizing cybersecurity performance measurement, and offers a step-by-step

approach to building a KPI framework that is fit for AI-driven environments, as well as being aligned with strategic business outcomes. The goal is to equip cybersecurity leaders to respond to change with clarity, rigor, and relevance.



2. Three emerging strategic priorities

Our analysis identified three emerging priorities that should guide this transformation:

- First, existing cybersecurity KPIs must be recalibrated to remain relevant in autonomous and AI-driven environments
- Second, new business-aligned KPIs are needed to connect cybersecurity performance directly to enterprise outcomes
- Finally, organizations must institutionalize continuous KPI evolution, ensuring that measurement frameworks remain adaptable

2.1 Adapt cybersecurity KPIs to the era of Agentic AI and Autonomy

Cybersecurity KPIs must evolve to adapt to autonomous environments and new business demands. The shift toward agentic AI and autonomy is rendering many traditional metrics obsolete or unreliable—especially those built for human-centered, manual workflows. Organizations must reassess existing metrics to prepare for the next era of cybersecurity.

To guide this reassessment, we classified KPIs along two dimensions: their degree of relevance in AI-driven environments, and their required degree of disruption or adaptability. This 2x2 lens makes it possible to cluster KPIs into four categories:

1. Recalibrated KPIs require modification to reflect hybrid or AI-assisted workflows. For example: Mean Time to Detect (MTTD) measures how quickly threats are identified. It remains a relevant metric, but in autonomous environments, thresholds must shrink from minutes to seconds.

2. Obsolete KPIs will lose utility as organizations transition toward autonomy. For example: Detection Rules are traditionally used to track manual rule creation in SIEM tools, but this

metric becomes less relevant as agentic AI handles threat detection dynamically, without relying on a static rule set.

3. New KPIs needed. For example: Explainability Score gauges how transparent and auditable AI decisions are. This metric is essential for building trust and governance in AI-powered systems.

4. Future-proof KPIs remain effective. For example: Time to Full Recovery measures the duration required to fully restore operations; an essential metric, whether the recovery process is AI-assisted or manual.

To guide the evolution of cybersecurity performance measurement, we conducted a structured analysis of over 100 KPIs, informed by expert consultations. The KPIs were assessed for relevance in autonomy-driven environments and mapped across the six functions of NIST Cybersecurity Framework (CSF) 2.0 (Govern, Identify, Protect, Detect, Respond, Recover) and the four foundational asset classes: People & Culture, Processes & Operations, Technology & Data, and Policies. From this analysis, KPIs were grouped into the above four categories.

Three key insights emerged from this analysis:

1. Tech & data-driven KPIs are most exposed to disruption

While some metrics remain future-proof, almost every NIST CSF function here will require adaptation. Most tech and data KPIs will need significant recalibration (e.g. MTTD), while new metrics should be put in place (e.g. AI correlation accuracy). This underscores the challenge of measuring system performance and integrity as architecture evolves towards greater autonomy.

2. Human-centered KPIs will need a major adaptation

A significant concentration of obsolete

KPIs (e.g. number of compliance assessment reports) is evident across the “People & Culture” dimension—particularly in the Protect and Respond functions. This signals a need to rethink how we measure human factors such as awareness, behavior, and decision-making in an AI-mediated environment.

3. Governance and policy KPIs remain relatively stable

KPIs tied to policies and governance appear largely future proof. These areas may benefit from refinement but face less disruptive change—suggesting that governance mechanisms will serve as stabilizing anchors amid rapid transformation elsewhere.




Figure 3: Strategic KPI disruption heatmap - consolidated insights

The executive insights and KPI distribution are visualized in Figure 3 and detailed in Appendix B. It is worth noting that these insights not only clarify how specific metrics should evolve, but

also directly underpin the three emerging strategic priorities: the need to recalibrate existing KPIs, introduce new business-aligned measures, and institutionalize continuous KPI evolution.

2.2 Introduce metrics to promote alignment between business & cyber teams

Cybersecurity must now demonstrate business enablement. As organizations accelerate AI & autonomy, cybersecurity is expected not only to manage risk but to actively enable growth, innovation, and trust.

To support this shift, the community identified business-aligned KPIs that link cybersecurity performance directly to enterprise outcomes. These include:

- **Cost of cyber incidents as % of revenue** quantifies the financial impact of such incidents relative to total revenue, incorporating both direct losses (e.g., downtime, disrupted operations) and indirect costs (e.g., response efforts, crisis communications, reputational repair)

- **% of projects with early cybersecurity engagement** measures the proportion of projects where cybersecurity is engaged during the initial requirements stage, indicating proactive risk alignment and early integration into business planning
- **Business user satisfaction with cyber support measures** the satisfaction score (e.g., via surveys or feedback ratings) of business users with cybersecurity team support

A full list of business-aligned KPIs is available in Appendix C.

2.3 Institutionalize continuous KPI evolution

The speed of technological change requires organizations to move beyond static dashboards and annual reviews. But evolution is not just about frequency—it is about timing, phasing, and anticipating when metrics will lose or gain relevance.

Over time, the balance will keep shifting, and what is “new” today may require recalibration tomorrow, and become obsolete later. To manage this lifecycle, organizations should build performance frameworks that explicitly include:

- **Regular review cycles (quarterly or biannual, rather than annual, to keep pace with AI autonomy)**

- **Forward-looking horizon scanning (identifying emerging risk domains where new KPIs will be needed)**
- **KPI lifecycle management (clear guidance on when to retire, recalibrate, or introduce metrics across near-, mid-, and long-term horizons)**

By treating KPIs as dynamic assets with life cycles, organizations can ensure that measurement frameworks remain relevant, future-oriented, and able to keep pace with AI’s accelerating impact on cybersecurity.



3. The Way Forward – a Blueprint to Adapt to the Next Era of Cybersecurity

This section builds on our findings by outlining a practical blueprint for CISOs and senior leaders who seek to align performance measurement with both the disruptive potential of AI, and cybersecurity’s evolving role as a business enabler.

It provides a structured approach to reviewing current metrics, identifying critical gaps, and designing new or adapted KPIs that capture autonomy,

risk, and strategic value. This process is iterative—it should be continuously refined as technologies advance and business priorities shift.

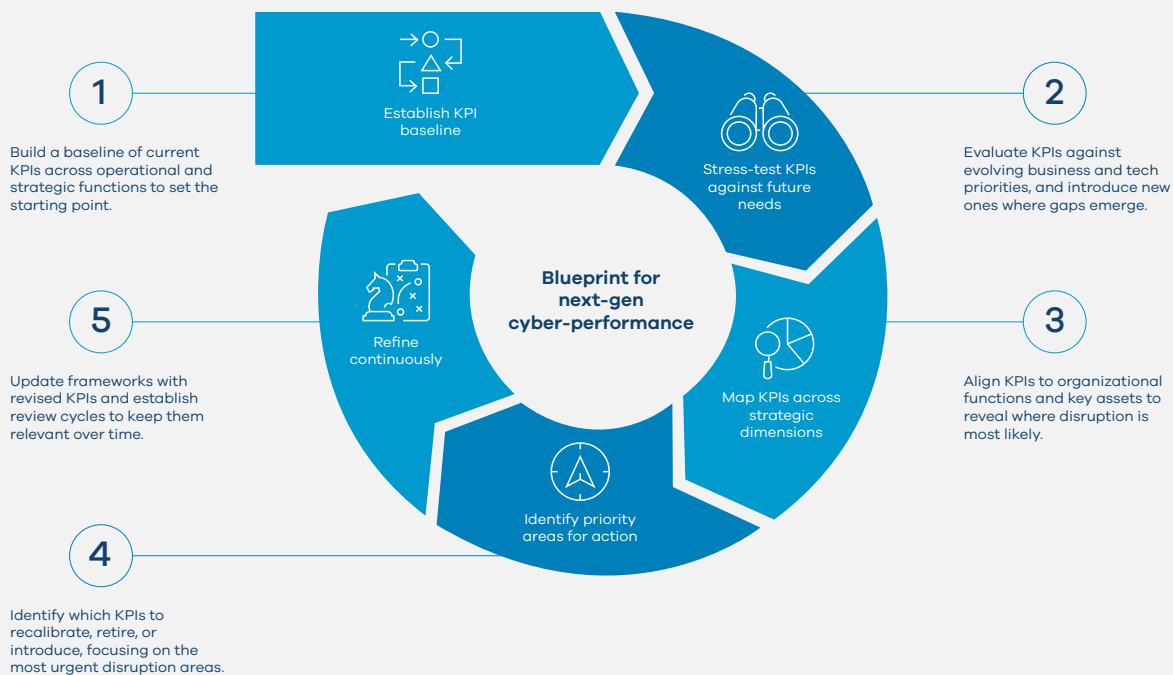


Figure 4: Blueprint for next-gen cyber performance

Step 1: Establish a baseline list of KPIs

A clear and meaningful baseline is essential to any performance management framework. Establishing this requires a structured evaluation of existing KPIs, not only in terms of what they measure but also why they exist.

The objective is to understand their purpose and relevance: which KPIs assume human rather than agentic or autonomous effort, which are reactive versus proactive, and which support business enablement versus those maintained for compliance.

To structure our baseline, we used the NIST CSF 2.0 as a guide, using its core

functions— Govern, Identify, Protect, Detect, Respond, and Recover—as categories to systematically consider and develop relevant KPIs. We complemented this functional approach by reflecting on key organizational essential assets such as People & Culture, Processes & Operations, Technology & Data, and Policies, allowing us to comprehensively identify KPIs across critical areas.

While the NIST functions and these organizational assets proved useful for structuring this step, there is no one-size-fits-all approach. Organizations should choose the framework, or combination of frameworks that best fits their context, maturity, and objectives.

Step 2: Categorize KPIs by future relevance

Key questions

To guide this step, organizations should address questions including:

- Which aspects of cybersecurity performance are our KPIs currently measuring?
- How many of our current KPIs assume human effort as opposed to AI-automation?
- Do our existing cyber-security teams' KPIs capture strategic business outcomes and alignment with business goals and innovation?
- Where are the gaps in readiness for agentic AI and autonomous systems?
- Where are the gaps that prevent better alignment with business goals and innovation?

Once the baseline inventory is established, the next step is to categorize each KPI based on its projected long-term value. This process requires a nuanced understanding of how KPIs relate not only to technological trends, such as AI adoption, but also to evolving business priorities and strategic objectives.

Metrics that monitor human review cycles, manual triage, or static threshold detection may lose utility as these processes become automated or delegated to AI agents. Conversely, new KPIs will emerge, focusing on alignment with business growth and innovation goals.

KPIs can be broadly categorized into four types: new KPIs that need to be developed, obsolete KPIs that may no longer add value, future-proof KPIs that remain relevant over time, and those that require recalibration. This categorization serves as a helpful lens for analysis but is not intended as a prescriptive framework. Organizations are encouraged to tailor their approach based on their specific context and needs.

Step 3: Map KPIs across strategic dimensions

Key questions

To guide this step, questions to be addressed include:

- Which KPIs will become irrelevant as decision-making becomes machine-led?
- Are any metrics too tightly tied to manual processes or outdated assumptions?
- Where are new KPIs needed to reflect AI governance, explainability, or customer trust?
- Can any existing KPIs be retained with recalibrated targets to reflect hybrid human-AI workflows?
- Are existing business-cyber alignment metrics effective to support organizational priorities and innovation?

To ensure contextual relevance, KPIs identified in earlier steps should be mapped across both functional and strategic dimensions. In this analysis, the NIST CSF 2.0 was used to align KPIs with the six core cybersecurity functions — Govern, Identify, Protect, Detect, Respond, and Recover — providing clarity on areas where disruption is most likely to occur.

Additionally, KPIs were mapped to key organizational asset categories: Policies, Processes & Operations, Technology & Data, and People & Culture.

This dual-layer approach highlights which metrics capture strategic value. It also reveals measurement gaps, and illustrates how cybersecurity performance links to broader business objectives.

While NIST CSF 2.0 offered a useful structure, again, there is no one-size-fits-all approach. Each organization should tailor its mapping method based on its strategic goals, cybersecurity maturity, and regulatory environment.

Step 4: Identifying and assessing areas of disruption

Key questions

To guide this step, organizations should address questions including:

- Which cybersecurity assets does each KPI measure or influence within the organization?
- How do KPIs align with strategic cybersecurity functions or processes?
- Are there gaps or blind spots in the current measurement approach that this mapping reveals?
- Where will AI-driven autonomy have the greatest impact on existing performance metrics?
- Are new performance metrics needed to better align cybersecurity teams with business and innovation functions?

After each KPI is mapped to the relevant cybersecurity functions and key organizational assets selected in earlier steps, these relationships can be visualized in a single matrix or heatmap. This two-dimensional view highlights where KPIs intersect across both axes and makes it clear where the greatest disruption is expected to take place.

High-disruption areas can be color-coded (e.g., orange or red) to enable quick, visual prioritization. An illustrative example of this mapping is provided in Appendix C, giving readers a concrete view of how disruption manifests across functions and asset categories.

This structured mapping enables the prioritization of KPIs for adaptation, replacement, or enhanced visibility, supporting cybersecurity's evolving role as a strategic enabler of both business and AI-driven innovation.

Step 5: Repeat throughout the organization cycle

Key questions

Issues to be considered include:

- Where is the highest concentration of disruption across the mapped KPIs? Which asset classes are most affected by this disruption?
- Which strategic dimensions should be prioritized for adaptation based on the intensity of disruption indicated?
- Which cybersecurity functions show the greatest impact? What are the implications for each cybersecurity department?
- How can KPIs be leveraged to measure and support both AI-transformation and broader business enablement objectives?

As cybersecurity landscapes are constantly evolving, KPI frameworks cannot remain static. Embedding regular review and refinement into the performance management process is essential to ensure ongoing relevance and effectiveness.

Organizations should treat the previous steps not as a one-time exercise, but as part of a continuous cycle. This means routinely reassessing KPIs, remapping them as cybersecurity environment and business priorities shift, and updating frameworks to reflect new technologies, emerging risks, and evolving strategic goals.

By institutionalizing this iterative approach, organizations can ensure that their KPI framework remains agile and adaptive, enabling them to anticipate future disruptions rather than simply react to them.



Conclusion

The rapid emergence of agentic AI and autonomous systems has ushered in a new era of cybersecurity—one that demands a fundamental shift in how performance is measured and managed. As cyber capabilities become greater and more complex, the frameworks used to evaluate their impact must evolve in parallel.

Traditional metrics designed for manual oversight are quickly losing their relevance. Because of this, organizations that recalibrate existing KPIs will be better positioned to maintain real-time visibility into machine-led environments.

At the same time, cybersecurity can no longer be measured in isolation from the enterprise. Business-aligned KPIs — such as cyber cost as a share of, or satisfaction with, cyber support — will reframe performance as a driver of value creation and business enablement, not simply risk mitigation.

Crucially, these KPIs cannot remain static. They will need to be introduced, recalibrated, or retired as technologies and threats evolve, ensuring that measurement frameworks remain both relevant and resilient. By managing KPIs as dynamic assets, organizations can turn cybersecurity into a strategic function that sustains innovation, strengthens trust, and enables growth.

Those who modernize their cybersecurity measurement will not only adapt to the age of autonomy — they will set the pace. In an environment where digital trust is a competitive differentiator, leadership in cybersecurity performance will define who thrives and who does not.

Appendices

Appendix A: NIST Cybersecurity Framework 2.0

The NIST CST 2.0 provides a globally recognized structure for managing digital risk, organized around six core functions — Govern, Identify, Protect, Detect, Respond, and Recover. Each function is further divided into

categories and subcategories (identifiers) that specify key outcomes and activities. This offers organizations a comprehensive and detailed reference model for assessing, prioritizing, and improving cybersecurity performance.

NIST CSF 2.0: Core Functions, Categories & Identifiers

Function	Category	Category Identifier
Govern (GV)	Organizational Context	GV.OC
	Risk Management Strategy	GV.RM
	Roles, Responsibilities, and Authorities	GV.RR
	Policy	GV.PO
	Oversight	GV.OV
	Cybersecurity Supply Chain Risk Management	GV.SC
Identify (ID)	Asset Management	ID.AM
	Risk Assessment	ID.RA
	Improvement	ID.IM
Protect (PR)	Identity Management, Authentication, and Access Control	PR.AA
	Awareness and Training	PR.AT
	Data Security	PR.DS
	Platform Security	PR.PS
	Technology Infrastructure Resilience	PR.IR
Detect (DE)	Continuous Monitoring	DE.CM
	Adverse Event Analysis	DE.AE
Respond (RS)	Incident Management	RS.MA
	Incident Analysis	RS.AN
	Incident Response Reporting and Communication	RS.CO
	Incident Mitigation	RS.MI
Recover (RC)	Incident Recovery Plan Execution	RC.RP
	Incident Recovery Communication	RC.CO

Appendix B: Detailed KPI disruption analysis

To guide the evolution of cybersecurity performance measurement, the Future of Cybersecurity Knowledge Community conducted a structured analysis of over 100 KPIs, informed by expert consultations. The KPIs were evaluated against defined criteria — including whether they were measurable, strategically relevant, and could be scored meaningfully — and were vetted through expert review to ensure their robustness.

The refined set was then assessed for relevance in AI-driven environments and mapped across the six functions of the NIST Cybersecurity Framework 2.0 (Govern, Identify, Protect, Detect, Respond, Recover) and the four foundational asset classes (People & Culture, Processes & Operations, Technology & Data, and Policies).

	GOVERN	IDENTIFY	PROTECT	DETECT	RESPOND	RECOVER
POLICIES	% compliance to national or sectoral cybersecurity regulations	Third-party risk assessment coverage			Service level agreement reporting	Integration of post-incident learnings
	% compliance to CS requirements (OpCos)					
	Explainability score					
	Project compliance with internal cybersecurity policies					
TECH & DATA		Risk level distribution	% of assets that have the minimum CS technologies deployed	Automation risk index	Investigation Automation Level	Time to full recovery
			% coverage of effective security controls implementation (IT)	Adversarial detection rate	Average incident response cost	
				Adaptation rate to novel attacks		
				Threat prioritization accuracy		
				Detection rules		
				Emerging threat coverage rate		
				Exposure intelligence		
				Attacker attribution		
				Volumes of alerts generated		
				Mean time to detect (MTTD)		
				False positives vs. false negatives		
				AI correlation accuracy		
PROCESSES & OPERATIONS	% of compliant third-parties	% of projects covered by cyber risk assessments		Adaptive threat detection use cases	Mean time to respond (MTTR)	% reduction in incident recovery time YoY
	CS impact on project timelines	% progress toward cybersecurity maturity			Root cause distribution	# overdue remediations of critical cybersecurity compliance findings
	Industry incident benchmarking				% cybersecurity compliance findings with remediation plan	% success rate of operational testing outcome
						Integration of post-incident learnings
						Cyber incident cost as % of revenue
PEOPLE & CULTURE	% of compliant third-parties			% projects with early cybersecurity engagement	Emerging threat coverage rate	Business user satisfaction with cyber support
	# of compliance assessment reports			Cybersecurity awareness score		# of compliance assessment reports
	# of critical cybersecurity third-party incidents			% cybersecurity compliance findings with remediation plan		Coverage of awareness campaigns
				Coverage of awareness campaigns		Alert distribution by department

■ New KPIs
 ■ Future proof KPIs
 ■ Recalibrated KPIs
 ■ Obsolete KPIs

C-1: New KPIs

- **AI correlation accuracy:** Measures the percentage of security alerts that are accurately correlated by AI into validated, true-positive incidents. This measures the effectiveness of AI in consolidating multiple low-level alerts into meaningful cases, reducing noise and improving SOC efficiency by lowering false groupings and alert fatigue.
- **Explainability score:** Measures the score (e.g., on a 0–100 scale) that reflects how clearly AI-generated actions and decisions can be explained, audited, and understood by human operators. It measures the transparency of AI reasoning, including whether rationale and evidence can be traced and communicated to stakeholders for accountability and trust.
- **Autonomy risk index:** Measures the composite index that evaluates the reliability and risk exposure of AI systems by tracking the proportion of autonomous actions taken without human input, the frequency of manual overrides, and the consistency with defined risk thresholds. It provides an aggregate measure of how safely and predictably AI agents operate within approved boundaries.
- **Adaptation rate to novel attacks:** Measures the average time (in hours or days) required for AI systems to detect, learn from, and incorporate defenses against previously unknown or stealthy attack patterns. This measures the responsiveness of agentic AI in closing detection gaps and generating new protective logic through autonomous learning.
- **Adversarial detection rate:** Measures the percentage of novel or evasive threats (outside known signatures or training data) that are successfully identified by the system. This includes detection of zero-day malware, AI-generated phishing, or anomalous behaviors such as unusual access patterns, providing a benchmark for the system's ability to counter emerging adversarial tactics.

C-2: Recalibrated KPIs

- **Mean time to detect (MTTD):** Measures the average time (in seconds or minutes) required to identify and validate a cybersecurity incident after it occurs. In autonomous environments, thresholds must shrink significantly compared to traditional baselines, reflecting the speed of AI-enabled detection.
- **Mean time to respond (MTTR):** Measures the average time (in seconds or minutes) taken to contain, mitigate, or neutralize a confirmed cybersecurity incident once detected. This KPI reflects the efficiency of response workflows and must adapt to faster, machine-led containment cycles in AI-driven operations.
- **False positives vs. false negatives:** Measures the ratio or percentage comparison between incorrectly flagged incidents (false positives) and undetected actual threats (false negatives). This KPI reflects detection accuracy and incorporates AI classification behavior and confidence scoring to balance sensitivity to real threats with minimizing noise.
- **Threat prioritization accuracy:** Measures the AI system's ability not only to detect threats but to accurately rank them based on contextual factors such as business impact, criticality, and operational relevance. This ensures that the most consequential threats are addressed first, aligning detection with organizational risk priorities.
- **Investigation automation level:** Measures the extent to which post-incident investigations are automated, from manual to fully autonomous, based on how data is gathered, correlated, and analyzed to generate actionable insights.
- **Average incident response cost:** Measures the time and effort required to manage a security incident end-to-end, with benchmarks recalibrated as

autonomous systems drive faster, more efficient resolution—from hours to minutes.

- **Risk level distribution:** Measures the distribution of risks across categories by incorporating asset criticality, business context, and potential impact. This KPI moves beyond static scoring to provide a dynamic, business-aware view of organizational risk exposure.
- **% reduction in incident recovery time YoY:** Tracks how much faster systems recover from incidents year over year. As autonomous systems optimize recovery processes, this metric needs recalibration to reflect reduced human intervention and faster recovery baselines.
- **% success rate of operational testing outcomes:** Indicates the percentage of operational cybersecurity tests (e.g., red teaming; breach and attack simulation) that meet predefined success criteria. With AI-assisted testing and adaptive defense mechanisms, success thresholds should adapt accordingly.
- **% coverage of effective security controls implementation (IT):** Measures the proportion of organizational assets protected by actively functioning security controls. With AI-driven automation accelerating deployment, this KPI shifts from basic implementation tracking to assessing real-time effectiveness, not just whether controls are in place, but whether they actively prevent or mitigate threats.
- **Service level agreement reporting:** Captures the percentage of security incidents resolved within defined SLA windows. In AI-enhanced environments, this metric must shift to align with faster resolution expectations and autonomous interventions.
- **Emerging threat coverage rate:** Measures the percentage of novel or evolving threat vectors that are successfully identified by detection systems, including AI. This KPI

evaluates AI's real-time adaptability and foresight in addressing emerging threats that fall outside traditional detection logic.

- **Adaptive threat detection use cases:** Measures the number and relevance of detection use cases aligned to specific threat scenarios, business risks, and regulatory requirements. This KPI reflects the system's ability to autonomously generate and update detection logic, ensuring continuous adaptation and sustained coverage.

C-3: Obsolete KPIs

- **Detection rules:** Measures engineering effort—tracking how many detection rules analysts manually create in SIEM tools to identify known threats (e.g., triggering an alert after five failed logins). However, as agentic AI increasingly handles event correlation across data sources without relying on static rules, the relevance of this metric declines.
- **Volume of alerts generated:** Traditionally used to gauge system activity, this KPI becomes misleading in AI-led environments. Agentic AI consolidates multiple alerts into single incidents, making alert volume less meaningful. Over time, alert counts reflect noise, not insight. Future metrics must focus on incident-level precision, since many alerts may be filtered or merged by AI and never surface as distinct security events.
- **# overdue remediations of critical cybersecurity compliance findings:** Tracks the number of unresolved high-priority compliance issues past their due date. Compliance workflows are expected to remain largely manual and governed by policy, making automation-driven remediation metrics less relevant in an AI-led operations model.
- **Cybersecurity awareness score:** Measures employee preparedness against threats like phishing, typically through quizzes and simulations. This KPI is expected to

become obsolete as agentic AI increasingly reduces human error by handling threats in real time, diminishing the need for employee-driven defense.

- **% of cybersecurity compliance findings with remediation plan:** Tracks whether identified compliance gaps have associated corrective actions and timelines, ensuring closure and accountability.
- **# of compliance assessment reports:** Tracks the number of formal reports documenting adherence to cybersecurity policies or regulations. As AI enables continuous, real-time compliance monitoring, the need for periodic manual reporting diminishes, reducing the relevance of this metric.
- **% of employees reached by awareness campaigns:** Proportion of employees who participated in or were reached by awareness activities (e.g., training modules, phishing simulations).

C-4: Future-proof KPIs

- **Time to full recovery:** Duration required to restore operations after an incident, whether recovery is AI-assisted or manual.
- **Average time to implement post-incident learnings:** Average duration (in days/weeks) from incident closure to implementing related control or procedural updates.
- **Third-party risk assessment coverage:** Measures the proportion of external vendors, partners, or service providers that have undergone formal cybersecurity risk assessments. This KPI captures the extent of oversight applied to third-party entities, with AI supporting data gathering and prioritization.
- **Cyber incident cost as % of revenue:** Quantifies the financial impact of cyber incidents relative to total revenue, incorporating both direct losses (e.g., downtime, disrupted operations) and indirect costs (e.g., response efforts, crisis

communications, reputational repair).

- **Industry Benchmark Incident Rate:** Compares the organization's monthly incident volume with peers in the same industry, enabling external performance benchmarking and contextual risk assessment.
- **Root cause distribution:** Analyzes the most common root causes behind incidents—such as security misconfigurations or inadequate user training—offering guidance for targeted remediation and systemic improvement.
- **Exposure intelligence:** Evaluates the level and spread of exposure across assets, departments, and business units, supporting proactive risk reduction and prioritization efforts.
- **Attacker attribution:** Identifies the threat actor or group responsible for an attack, based on indicators such as techniques, infrastructure, and digital forensics.
- **% compliance to national or sectoral cybersecurity regulations:** Measures adherence to applicable national or sector-specific cybersecurity requirements. Essential for maintaining regulatory alignment, reducing risk exposure, and avoiding legal or financial penalties.
- **Project compliance with internal cybersecurity policies:** Assesses the extent to which business projects align with established cybersecurity requirements. This KPI ensures that cybersecurity is embedded early in planning and execution, reducing risks and strengthening organizational resilience.
 - **% of projects compliant with internal policies:** Measures the proportion of projects that fully adhere without deviation to cybersecurity requirements. Demonstrates effective integration of cybersecurity into business practices
 - **% of projects non-compliant with internal policies:** Captures

the percentage of projects failing to meet baseline cybersecurity standards. Highlights risk exposure and signals areas requiring immediate corrective action

- **% of projects with approved exceptions:** Tracks the share of projects that received formal exemptions from specific cybersecurity requirements. Ensures transparency in governance and provides oversight on managed risk trade-offs
- **# of critical cybersecurity third-party incidents:** Tracks breaches or issues originating from third parties that significantly impact the organization. Drives focus on supplier vetting and monitoring.
- **% of compliant third parties:** Indicates how many vendors meet the organization's minimum cybersecurity requirements. Critical for ensuring that security extends beyond internal boundaries.
- **% progress toward cybersecurity maturity:** Measures advancement across recognized maturity models (e.g., NIST CSF tiers). Demonstrates strategic progression from reactive to proactive postures.
- **% compliance to CS requirements (OpCos):** Reflects how well operational companies within a group adhere to central cybersecurity standards. Ensures consistency across distributed environments.
- **% of assets with minimum CS technologies deployed:** Measures the percentage of organizational assets that have the minimum required cybersecurity technologies installed and active.
- **Business user satisfaction with cyber support:** Measures satisfaction score (e.g., via surveys or feedback ratings) of business users with cybersecurity team support.
- **% of projects covered by cybersecurity risk assessments:** Measures percentage of projects that underwent formal cybersecurity risk assessment during planning or execution phases. Includes risk workshops, control evaluations, or third-party assessments with documented outcomes.
- **% of projects delivered on time with cybersecurity integration:** Tracks the percentage of projects completed within planned timelines while meeting any required cybersecurity reviews and controls, reflecting effective integration of cyber without causing delivery delays.
- **% projects with early cybersecurity engagement:** Measures the proportion of projects where cybersecurity is engaged during the initial requirements stage, indicating proactive risk alignment and early integration into business planning.

Appendix C: Metrics to support business and cybersecurity synchronicity

- **Business user satisfaction with cyber support:** Measures satisfaction score (e.g., via surveys or feedback ratings) of business users with cybersecurity team support.
- **% of projects covered by cybersecurity risk assessments:** Measures percentage of projects that underwent formal cybersecurity risk assessments during planning or execution phases. Includes risk workshops, control evaluations, or third-party assessments with documented outcomes.
- **% of projects delivered on time with cybersecurity integration:** Tracks the percentage of projects completed within planned timelines while meeting any required cybersecurity reviews and controls, reflecting effective integration of cyber without causing delivery delays.
- **% of projects with early cybersecurity engagement:** Measures the proportion of projects where cybersecurity is engaged during the initial requirements stage, indicating proactive risk alignment and early integration into business planning.
- **% progress toward cybersecurity maturity:** Measures advancement across recognized maturity models (e.g., NIST CSF tiers). Demonstrates strategic progression from reactive to proactive postures.
- **% compliance with CS requirements (OpCos):** Reflects how well operational companies within a group adhere to central cybersecurity standards. Ensures consistency across distributed environments.
- **Cyber incidents cost as % of revenue:** Quantifies the financial impact of cyber incidents relative to total revenue, incorporating both direct losses (e.g., downtime, disrupted operations) and indirect costs (e.g., response efforts, crisis communications, reputational repair).
- **% reduction in incident recovery time YoY:** Tracks how much faster systems recover from incidents year after year. As autonomous systems optimize recovery processes, this metric needs recalibration to reflect reduced human intervention and faster recovery baselines.
- **Average incident response cost:** Measures the time and effort required to manage a security incident end-to-end, with benchmarks recalibrated as autonomous systems drive faster, more efficient resolution—from hours to minutes.

Bibliography

1. Boston Consulting Group (2025). A Synchronized Approach to Digital Risk.
2. Boston Consulting Group (2025). Agentic AI Ecosystems in the Tech for AI.
3. Boston Consulting Group & Gerson Lehrman Group (2025). CISO Survey.
4. Boston Consulting Group & Gerson Lehrman Group (2024). IT Buyer Pulse Check 8.0.
5. Boston Consulting Group (2024). Measuring AI Ability to Complete Long Tasks.
6. Gartner (2024). Gartner CEO & Senior Business Executive Survey (fieldwork)
7. International Business Machines (IBM) (2022). Prosper in the Cyber Economy.
8. International Business Machines (IBM) (2025). Cost of a Data Breach Report.
9. International Scientific Panel (2025). International AI Safety Report.
10. IEEE SA (2024). The IEEE Global Initiative 2.0 on Ethics of Autonomous and Intelligent Systems.
11. International Finance Forum (IFF) (2024). Global AI Competitiveness Index Report – Part 1: Enterprise Perspective.
12. Journal of Information Security (2025). Evolving Cybersecurity Policy: Addressing Modern Threats and Enhancing Resilience in a Digital Age.
13. National Institute of Standards and Technology (NIST) (2024). Cybersecurity Framework 2.0.
14. Stanford University (2022). One Hundred Year Study on AI: “Gathering Strength, Gathering Storms”.
15. United Nations / OECD / EU (2024). Interim Global Scientific Report on Advanced AI Safety.
16. US Department of Defense (2025). Strengthening Multimedia Integrity in the Generative AI Era.
17. World Economic Forum (2025). Artificial Intelligence and Cybersecurity: Balancing Risks and Rewards (WEF white paper).
18. World Economic Forum (2023). Chief Risk Officers Outlook.
19. World Economic Forum (2025). Global Cybersecurity Outlook 2025.
20. World Economic Forum (2025). The Global Risks Report 2025.

