LMNTRIX

BE THE HUNTER | NOT THE PREY

CYBER RISK QUANTIFICATION FOR SECURITY OPERATIONS

Using FAIR Modeling to Prioritize, Investigate and Respond to Cyber Threats in Real Time

WHITEPAPER

LMNTRIX USA
19800 MacArthur Blvd,
Suite 850
Irvine, CA 92612
sales@lmntrix.com
888-388-1879

LMNTRIX UK
Kemp House, 152 – 160
City Road, London, EC1V
2NX
sales@lmntrix.com
+44.808.164.9442

LMNTRIX INDIA
VR Bengaluru, Level 5, ITPL Main
Rd, Devasandra Industrial Estate,
Bengaluru, Karnataka 560048, India
sales@lmntrix.com
+91-22-49712788

LMNTRIX AUSTRALIA
Level 25, 100 Mount street,
North Sydney 2060
sales@lmntrix.com
+61.288.805.198

LMNTRIX SINGAPORE
60 Kaki Bukit Place, #05-19,
Eunos TechPark
sales@lmntrix.com
+65-3129-2639

lmntrix.com

Contents

EXECUTIVE SUMMARY

Modern Security Operations Centers (SOCs) generate unprecedented volumes of telemetry and security alerts, yet many organizations continue to prioritize operational efficiency over measurable business risk. Conventional performance metrics such as Mean Time to Detect (MTTD), Mean Time to Respond (MTTR), alert volumes and case closure rates provide valuable operational insight but fail to quantify the financial consequences of cyber incidents or identify which threats pose the greatest organizational exposure.

Executive leaders increasingly require security teams to demonstrate how cyber events translate into business disruption, regulatory exposure and financial loss. As cyber threats become more sophisticated and resource constraints intensify, security operations must evolve from managing alerts to actively reducing enterprise risk through evidence-based decision-making.

The **Factor Analysis of Information Risk (FAIR)** model provides a recognized framework for quantifying cyber risk in financial terms by estimating both the probable frequency and probable magnitude of future loss events. Rather than relying solely on technical severity or static risk scores, FAIR enables security teams to prioritize investigations according to expected business impact, allowing scarce analytical resources to focus on the incidents that present the highest potential financial exposure.

This white paper proposes extending FAIR beyond periodic governance and enterprise risk assessments by embedding quantitative risk modeling directly into day-to-day SOC workflows. Continuous risk quantification enables every investigation, intelligence update and telemetry event to dynamically recalculate organizational exposure, producing prioritization decisions that reflect changing business conditions in near real time.

The paper further explores how artificial intelligence and automation can continuously evaluate FAIR variables, correlate telemetry, assess control effectiveness and recommend response actions while maintaining human oversight for strategic decision-making. The result is a Security Operations Center that functions not simply as an alert-processing capability, but as an intelligent business risk reduction engine that continuously aligns cyber defense activities with measurable organizational outcomes.

WHY SECURITY OPERATIONS MEASURE THE WRONG THINGS

Security Operations Centers (SOCs) have evolved significantly over the past decade, becoming highly effective at collecting telemetry, detecting threats, and orchestrating incident response. However, despite advances in automation, artificial intelligence (AI), and extended detection and response (XDR), most SOCs continue to measure success using operational metrics rather than business outcomes. Metrics such as alert volume, Mean Time to Detect (MTTD), Mean Time to Respond (MTTR), false positive rates, analyst utilization, and incidents closed provide insight into operational efficiency but reveal little about the organization's actual cyber risk exposure (NIST, 2024).

These metrics answer important tactical questions, including *What happened?*, *How quickly was the incident identified?*, and *How efficiently did analysts respond?* Yet they fail to answer the questions most relevant to executive leadership and boards of directors: *Which incident creates the greatest financial exposure? Which compromised asset represents the highest business risk? How much potential loss has been avoided through effective response?* Without quantitative answers to these questions, SOCs risk optimizing operational performance while overlooking the threats that could have the greatest business impact (The Open Group, 2023).

Traditional severity ratings further illustrate this challenge. Security Information and Event Management (SIEM) and XDR platforms frequently prioritize incidents based on technical indicators such as Common Vulnerability Scoring System (CVSS) ratings, detection confidence, or the number of alerts generated. While technically useful, these approaches often fail to account for business context (Google Cloud, 2025). A ransomware intrusion targeting a domain controller supporting critical business operations may generate relatively few alerts, while widespread commodity malware affecting low-value workstations can generate hundreds. Conventional SOC metrics frequently emphasize the latter because of alert volume rather than the former because of business impact.

This disconnect reflects a broader issue within cybersecurity operations: technical severity is not synonymous with organizational risk. Cyber risk is determined by the interaction between threat capability, organizational vulnerabilities, asset criticality, control effectiveness, and the potential financial consequences of a successful attack (Verizon, 2025). Organizations that measure only operational activity lack a consistent mechanism for understanding whether their security investments are reducing enterprise risk or merely improving process efficiency.

Industry research reinforces this distinction. The *Verizon Data Breach Investigations Report* consistently demonstrates that a relatively small number of successful attacks account for a disproportionate share of financial losses, while IBM's *Cost of a Data Breach Report* continues to show that business disruption, regulatory penalties, reputational damage, and recovery costs often exceed the immediate technical impact of an incident (IBM, 2025; Verizon, 2025). These findings suggest that the

effectiveness of a SOC should ultimately be measured by its ability to reduce expected business loss rather than simply process security alerts more efficiently.

Modern security operations therefore require a shift from operational metrics toward risk-centric decision-making. Rather than asking which alert should be investigated first based solely on technical severity, security teams should prioritize incidents according to their probable financial impact on the organization. This transition requires a quantitative framework capable of translating technical events into measurable business risk—a capability that traditional SOC metrics were never designed to provide. The Factor Analysis of Information Risk (FAIR) model addresses this gap by providing a structured methodology for estimating the probable frequency and probable magnitude of cyber loss events, enabling security operations to prioritize investigations based on expected financial exposure rather than alert volume or technical severity alone (The Open Group, 2023; FAIR Institute, 2025).

UNDERSTANDING FAIR

The cybersecurity industry has traditionally relied on qualitative methods to assess risk. Terms such as *critical*, *high*, *medium*, and *low* have become standard across vulnerability management platforms, Security Information and Event Management (SIEM) systems, and Extended Detection and Response (XDR) platforms. While these classifications provide useful operational guidance, they remain subjective and frequently fail to communicate cyber risk in terms that resonate with executive leadership. A "critical" vulnerability may represent vastly different levels of financial exposure depending on the asset affected, the effectiveness of existing security controls, and the organization's business context. As a result, security teams often prioritize incidents based on technical severity rather than measurable business impact (NIST, 2024).

The FAIR model was developed to address this challenge by providing a standardized framework for quantifying cyber risk in financial terms. Rather than assigning subjective risk ratings, FAIR estimates the probable frequency and probable magnitude of future loss events (The Open Group, 2023). These estimates enable organizations to express cyber risk as an expected financial exposure, allowing security professionals, executives, and boards of directors to evaluate cybersecurity investments using the same language applied to other forms of enterprise risk management.

Unlike traditional risk matrices, FAIR decomposes cyber risk into measurable variables that collectively determine the likelihood and potential financial impact of a loss event. By evaluating each variable independently, organizations can build transparent, repeatable, and evidence-based estimates of cyber risk while avoiding the ambiguity associated with subjective scoring systems (FAIR Institute, 2025).

At the highest level, FAIR defines risk as the **probable frequency and probable magnitude of future loss**. To calculate these estimates, the framework analyzes six fundamental components that describe both the probability of a successful attack and its expected business consequences.

Loss Event Frequency (LEF) estimates how often an organization can expect a loss event to occur over a given period. Rather than counting the number of attacks, LEF measures the expected occurrence of incidents that actually result in business loss. This distinction is important because organizations experience millions of malicious events each year, yet only a small percentage result in material financial impact.

Threat Event Frequency (TEF) measures how often threat actors are expected to act against a particular asset or organization. TEF considers factors such as attacker motivation, capability, opportunity, and exposure to determine how frequently malicious activity is likely to occur.

Vulnerability within FAIR differs from the traditional definition of a software vulnerability. Instead of referring solely to a technical weakness, FAIR defines vulnerability as the probability that a threat action will successfully overcome existing security controls when attempted. This broader definition recognizes that organizational resilience depends not only on technology but also on people, processes, detection capabilities, and response effectiveness.

Contact Frequency (CF) estimates how often threat actors come into contact with a potential target. Internet-facing systems, cloud applications, publicly exposed APIs, and widely used identities naturally experience significantly higher contact frequencies than isolated internal systems.

Probability of Action (PoA) measures the likelihood that a threat actor will exploit an opportunity once contact has been established. Different adversaries exhibit different behaviors and objectives. Opportunistic cybercriminals, financially motivated ransomware groups, insider threats, and nation-state actors all possess different probabilities of acting based on their objectives, resources, and expected return on investment.

Finally, **Loss Magnitude (LM)** estimates the financial impact should a successful loss event occur. FAIR encourages organizations to consider both primary losses, such as incident response costs, business interruption, forensic investigations, legal expenses, and technology recovery, as well as secondary losses including regulatory penalties, contractual liabilities, reputational damage, customer attrition, and long-term revenue impacts. By evaluating these factors together, organizations gain a comprehensive understanding of the total economic consequences of cyber incidents rather than focusing solely on technical remediation costs (Jones, 2020).

These components interact to produce estimates such as **Single Loss Expectancy (SLE)** and **Annualized Loss Expectancy (ALE)**. Although these concepts have existed within information security for many years, FAIR strengthens their accuracy by replacing static assumptions with structured probability analysis based on available evidence and expert judgment. Rather than generating a single deterministic value, FAIR typically produces a range of probable financial outcomes that better reflects the uncertainty inherent in cybersecurity risk.

One of FAIR's greatest strengths is that it does not attempt to predict the future with absolute certainty. Cybersecurity remains inherently uncertain because threat actor behavior, emerging vulnerabilities, and organizational environments continuously evolve. Instead, FAIR applies probability distributions, historical observations, expert analysis, and scenario modeling to estimate the range of

likely financial outcomes. This probabilistic approach provides decision-makers with a far more realistic understanding of risk than traditional qualitative scoring methods.

The practical value of FAIR extends beyond enterprise risk management and board reporting. When integrated into security operations, FAIR enables analysts to evaluate every investigation through the lens of expected business impact rather than technical severity alone. Instead of asking which alert generated the highest confidence score or affected the largest number of endpoints, analysts can prioritize the incident expected to produce the greatest financial loss if left unresolved. This subtle but important shift transforms incident response from a reactive operational function into a business-driven decision process.

Importantly, FAIR should not be viewed as a replacement for existing cybersecurity frameworks such as the **NIST Cybersecurity Framework**, **MITRE ATT&CK**, or **MITRE D3FEND**. These frameworks describe how organizations identify, protect, detect, respond to, and recover from cyber threats. FAIR complements these operational frameworks by providing the quantitative layer that explains **which incidents matter most from a business perspective**. Together, they enable security teams to combine technical evidence with financial analysis, ensuring that limited resources are consistently directed toward reducing the organization's greatest sources of cyber risk.

The following section explores how this quantitative methodology can move beyond periodic enterprise risk assessments and become an integral component of daily Security Operations Center workflows, continuously recalculating organizational exposure as new telemetry, threat intelligence, and investigative findings emerge.

Illustration: From Alerts to Financial Exposure

Traditional SOC Approach

Focus: What happened? How serious is the attack? Which asset was affected?
Limitation: *Does not quantify business risk or financial impact.*

FAIR-Enabled SOC Approach

Critical Alert Generated	Alert or threat is detected.
Quantify Risk (FAIR)	FAIR model calculates likelihood and impact based on current conditions.
Estimated Financial Exposure	Produces Single Loss Expectancy (SLE) and Annualized Loss Expectancy (ALE).
Prioritised Investigation	Incidents are prioritised by expected financial exposure, not just technical severity.
Targeted Response	Response actions focus on reducing the highest financial exposure.
Exposure Reduction Measured	Recalculate risk to measure reduction in expected financial loss.

Focus: What is the financial exposure right now? Which incident threatens the organisation most? What response will reduce the most expected loss?

Outcome: Resources are directed where they reduce the most business risk.

BRINGING FAIR INSIDE THE SOC

Traditional implementations of FAIR have focused primarily on strategic governance activities, enterprise risk management, and periodic cyber risk assessments. In many organizations, FAIR analyses are performed quarterly or annually to support board reporting, cyber insurance decisions, regulatory compliance, or investment planning. While these assessments provide valuable insight into enterprise risk posture, they represent only a snapshot of organizational exposure at a specific point in time. Unfortunately, cyber risk evolves continuously, often changing within minutes as new vulnerabilities emerge, threat actors shift tactics, security controls fail, or critical business assets become exposed.

Modern SOCs operate in an environment of constant change. Every security alert, endpoint detection, intelligence update, identity compromise, or vulnerability disclosure has the potential to alter the organization's overall risk profile. Despite this dynamic operating environment, most SOCs continue to prioritize investigations using static severity scores, predefined playbooks, or analyst judgment rather than continuously evaluating expected business impact. As a result, incident priorities established at the beginning of an investigation frequently remain unchanged even as new evidence significantly alters the level of organizational exposure.

This white paper proposes a fundamentally different approach: embedding FAIR directly into day-to-day SOC operations as a continuous decision-support capability rather than treating it as a standalone governance exercise. Instead of calculating cyber risk periodically, FAIR variables should be continuously updated as telemetry, threat intelligence, asset context, identity information, and investigative findings become available. Every new piece of evidence should contribute to a real-time recalculation of expected financial exposure, enabling incident priorities to evolve alongside the threat landscape.

Under this model, risk quantification becomes an active component of every investigation rather than an activity performed after incidents have concluded. As analysts uncover new information, automated workflows continuously reassess the probability of a successful attack, estimate potential business loss, evaluate the effectiveness of existing controls, and determine whether investigative resources should be redirected toward incidents presenting greater organizational risk. The SOC therefore shifts from managing alerts based primarily on technical severity to managing investigations based on measurable financial exposure.

The following examples demonstrate how continuous FAIR modeling transforms routine security investigations into dynamic risk assessment processes. Whether responding to a credential theft campaign or an emerging ransomware threat, each telemetry update, intelligence feed, and analyst observation contributes to an updated estimate of business risk. This continuous recalculation enables security teams to prioritize the incidents that present the greatest expected financial loss, ensuring that limited analyst resources consistently focus on reducing the organization's most significant cyber risks rather than simply processing the highest volume of alerts.

Section 3: Bringing FAIR Inside the SOC

FAIR is traditionally performed on a periodic basis (e.g., quarterly or annually). This section proposes a fundamental shift: embedding FAIR into daily SOC operations to continuously quantify and reprioritize risk as new events, telemetry and intelligence emerge.

Example 1: Credential Theft

Credential Theft Detected Phishing alert indicates possible credential compromise.	Initial identity risk increases.
Identity Risk Increases Compromised credentials provide potential access to additional systems.	Risk score increases.
Business-Critical Account Identified Access to privileged account with high business impact is detected.	Business impact validated.
FAIR Engine Recalculates Expected Financial Exposure Updated variables increase Single Loss Expectancy (SLE) and Annualized Loss Expectancy (ALE).	Expected financial exposure rises.
Incident Priority Automatically Increases FAIR risk score exceeds defined threshold.	Priority escalated automatically.
Analyst Investigates High-Risk Incident Focused investigation on high-value identity compromise.	Investigation focused.
Containment Actions Executed Credentials reset, sessions terminated, access restricted.	Risk mitigated.
Expected Financial Exposure Reduced FAIR model recalculates lower exposure post-containment.	Exposure reduced.

Continuous FAIR modeling dynamically reprioritizes investigations based on changing business risk.

Example 2: Ransomware Threat

New Ransomware IOC Detected Threat intel feed or EDR alert identifies new ransomware indicators.	New threat introduced.
Threat Intelligence Correlated IOC linked to active ransomware campaign targeting similar organizations.	Threat relevance confirmed.
Critical Business Servers Identified Exposed systems mapped to high-value business services.	Critical assets at risk.
Exposure and Control Effectiveness Assessed Current controls evaluated for ability to prevent ransomware impact.	Control gaps identified.
FAIR Risk Score Recalculated Updated variables increase SLE and ALE for ransomware scenario.	Financial exposure recalculated.
Incident Priority Automatically Escalates Risk score exceeds threshold for automatic escalation.	Priority escalated.
Recommended Response: Isolation + Containment Host isolation, blocking, backups verified, and monitoring increased.	Containment in progress.
Projected Financial Loss Reduced FAIR model projects lower loss based on containment effectiveness.	Expected loss reduced.

Every telemetry update continuously recalculates organizational exposure, ensuring analysts focus on the incidents with the highest expected business impact.

Continuous Risk Quantification: Every event, detection, or intelligence update feeds the FAIR engine, which recalculates risk in real time. SOC teams prioritize what matters most—reducing the greatest expected financial loss.

BUILDING DYNAMIC RISK SCORES

Traditional Security Operations Centers (SOCs) assign risk scores at the point of detection. A security alert is generated, assigned a severity rating, triaged by an analyst, and ultimately closed once remediation activities are complete. While this approach has served organizations for many years, it assumes that risk remains relatively static throughout the lifecycle of an incident. In reality, cyber risk is highly dynamic. Every new piece of intelligence, every investigative finding, and every change to the threat landscape has the potential to significantly alter an organization's exposure. As a result, risk should not be viewed as a fixed value assigned when an alert is created, but as a continuously evolving estimate that reflects the organization's current operating environment.

From Static Scores to Continuous Risk Assessment

Building dynamic risk scores requires moving beyond traditional severity classifications toward continuous quantitative risk assessment. Rather than assigning labels such as *Critical*, *High*, or *Medium*, a modern SOC should continuously estimate the expected financial exposure associated with every active investigation. As new information becomes available, that estimate should be recalculated automatically, allowing incident priorities to evolve in parallel with changing business conditions. This enables security teams to focus resources where they will reduce the greatest amount of enterprise risk instead of simply responding to the loudest or most technically severe alerts.

At the center of this approach is FAIR™ model, which provides the quantitative framework needed to estimate probable loss frequency and probable loss magnitude. However, FAIR calculations are only as effective as the quality and breadth of the information feeding them. Modern SOCs possess an enormous volume of contextual data that can continuously refine these calculations. Rather than relying on periodic assessments, organizations can leverage operational telemetry to create dynamic risk models that become increasingly accurate as investigations progress.

The Data That Drives Dynamic Risk

One of the most influential inputs is **threat intelligence**. External intelligence feeds, industry information-sharing organizations, commercial intelligence providers, and internal threat research continuously identify emerging adversary tactics, active ransomware campaigns, exploited vulnerabilities, and indicators of compromise (IOCs). As new intelligence becomes available, FAIR variables such as Threat Event Frequency (TEF) and Probability of Action (PoA) can be updated automatically to reflect changes in adversary behavior. For example, the discovery that a vulnerability is being actively exploited by multiple ransomware groups immediately increases the likelihood of compromise and should elevate the financial exposure associated with affected assets (MITRE, 2025).

Asset value represents another critical variable in dynamic risk quantification. Not all systems contribute equally to organizational operations. A compromised development workstation and a compromised enterprise identity provider may generate similar technical alerts while representing vastly different business consequences. Dynamic risk models incorporate business context by evaluating asset criticality, revenue dependency, operational importance, regulatory sensitivity, and data classification. This ensures that investigations prioritize systems whose compromise would create the greatest financial and operational disruption.

Similarly, **identity criticality** has become one of the defining factors of modern cyber risk. Identity has replaced the traditional network perimeter as the primary target for attackers, making privileged accounts, service accounts, cloud administrators, and executive identities particularly valuable. A credential theft alert involving a standard user account may warrant routine investigation, whereas identical activity affecting a Global Administrator account could dramatically increase expected financial exposure. Incorporating identity context into FAIR calculations enables SOCs to recognize these differences immediately and adjust investigation priorities accordingly.

Dynamic risk models should also incorporate **attack path analysis**. Modern attack frameworks demonstrate that adversaries rarely compromise their ultimate target directly. Instead, they progress through multiple stages of reconnaissance, privilege escalation, credential theft, lateral movement, persistence, and data exfiltration. By continuously mapping observed attacker behavior against known attack paths, organizations can estimate how close an adversary is to achieving strategic objectives. As attackers move deeper into critical business environments, expected financial exposure increases, prompting automatic reprioritization of investigative efforts (MITRE ATT&CK®, 2025).

Business context extends beyond technology alone. **Business unit criticality** should also influence quantitative risk calculations. Systems supporting manufacturing operations, healthcare services, financial transactions, customer-facing platforms, or executive decision-making often carry significantly greater operational and regulatory consequences than internal administrative systems. Integrating business ownership into FAIR calculations ensures that security priorities remain aligned with enterprise objectives rather than purely technical considerations.

Evaluating Defensive Posture in Real Time

Another important variable is **exploit maturity**. Vulnerabilities evolve rapidly following public disclosure. Initially, exploitation may require sophisticated technical expertise. Within days or weeks, publicly available exploit code, automated scanning tools, and ransomware playbooks frequently reduce the technical barriers for attackers. Dynamic risk scoring accounts for this progression by continuously increasing the estimated likelihood of successful exploitation as exploit maturity advances from proof-of-concept to widespread weaponization.

Equally important is the assessment of **control effectiveness**. Traditional vulnerability management frequently assumes that every vulnerable asset represents identical risk. In reality,

Organizations deploy multiple overlapping security controls including endpoint protection, multi-factor authentication, privileged access management, network segmentation, identity protection, email security, and continuous monitoring. FAIR recognizes that effective controls reduce both the probability and magnitude of successful loss events. As controls fail, degrade, or improve, expected financial exposure should be recalculated automatically to reflect the organization's true defensive posture.

Modern SOCs should also evaluate the **confidence of their telemetry**. Security decisions are only as reliable as the data informing them. Missing endpoint visibility, incomplete cloud telemetry, disabled security agents, logging failures, or blind spots within network monitoring all reduce confidence in risk estimates. Rather than assuming incomplete visibility represents lower risk, dynamic FAIR models should recognize uncertainty itself as a contributor to organizational exposure. Reduced telemetry confidence may increase estimated uncertainty and justify additional investigation before risk can be accurately assessed.

From Data to Financial Risk

Historical incident data provides another valuable source of continuous learning. Every investigation produces evidence regarding attacker behavior, control effectiveness, response times, financial impact, and recovery costs. Rather than archiving this information after incidents close, organizations should continuously incorporate historical observations into future FAIR calculations. Over time, this creates progressively more accurate estimates of probable loss frequency and magnitude while improving confidence in executive reporting and investment decisions.

To manage these continuously changing variables, organizations increasingly employ probabilistic modeling techniques such as **Bayesian updating** and **Monte Carlo simulation**. Bayesian updating allows the FAIR model to revise risk estimates whenever new evidence becomes available. Rather than replacing previous assessments entirely, each new observation incrementally improves the organization's understanding of current risk. Monte Carlo simulation complements this approach by evaluating thousands of possible outcomes across varying assumptions, generating a range of probable financial losses instead of a single deterministic value. Together, these techniques provide decision-makers with a realistic understanding of uncertainty while supporting more informed prioritization and investment decisions (Carnegie Mellon Software Engineering Institute, 2024).

A Foundation for the Next Generation SOC

Ultimately, dynamic risk scoring transforms cybersecurity from a reactive operational discipline into a continuously adaptive business function. Every detection, intelligence update, analyst observation, and control change contributes to a living estimate of enterprise cyber risk. Rather than asking whether an incident is technically severe, the SOC can continuously answer a far more valuable question: **How much financial exposure exists at this moment, and which action will reduce it most effectively?** This capability provides the foundation for the next generation of AI-assisted security operations, where investigations are prioritized not by static alert severity, but by continuously evolving business risk.

SOC WORKFLOW TRANSFORMATION

The integration of FAIR into SOC workflows fundamentally changes how cyber incidents are prioritized, investigated, and remediated. Traditional SOCs are designed to optimize operational efficiency, focusing on processing alerts as quickly as possible through standardized triage and response procedures. While this approach improves operational performance, it often fails to distinguish between incidents based on their actual business impact. A FAIR-enabled SOC shifts this paradigm by embedding continuous risk quantification throughout the incident lifecycle. Every alert is evaluated within the context of business criticality, threat intelligence, asset value, control effectiveness, and estimated financial exposure. As new evidence emerges during an investigation, risk calculations are automatically updated, enabling analysts to continuously reprioritize their efforts toward the incidents that pose the greatest threat to the organization. The result is a transformation from an alert-driven workflow into a business risk reduction process, where success is measured not by the number of incidents processed, but by the measurable reduction in enterprise cyber risk and expected financial loss (The Open Group, 2023; NIST, 2024).

Section 5: SOC Workflow Transformation

Embedding FAIR into SOC operations fundamentally transforms how security teams prioritize, investigate, and respond to incidents. The traditional workflow is event-driven and operationally focused, optimized for speed and efficiency. The future workflow is risk-driven and outcome-focused, optimized for reducing financial exposure. By quantifying business impact in real time, SOCs evolve from alert-processing centers into intelligent risk reduction engines that continuously align security actions with the organization's most critical business outcomes.

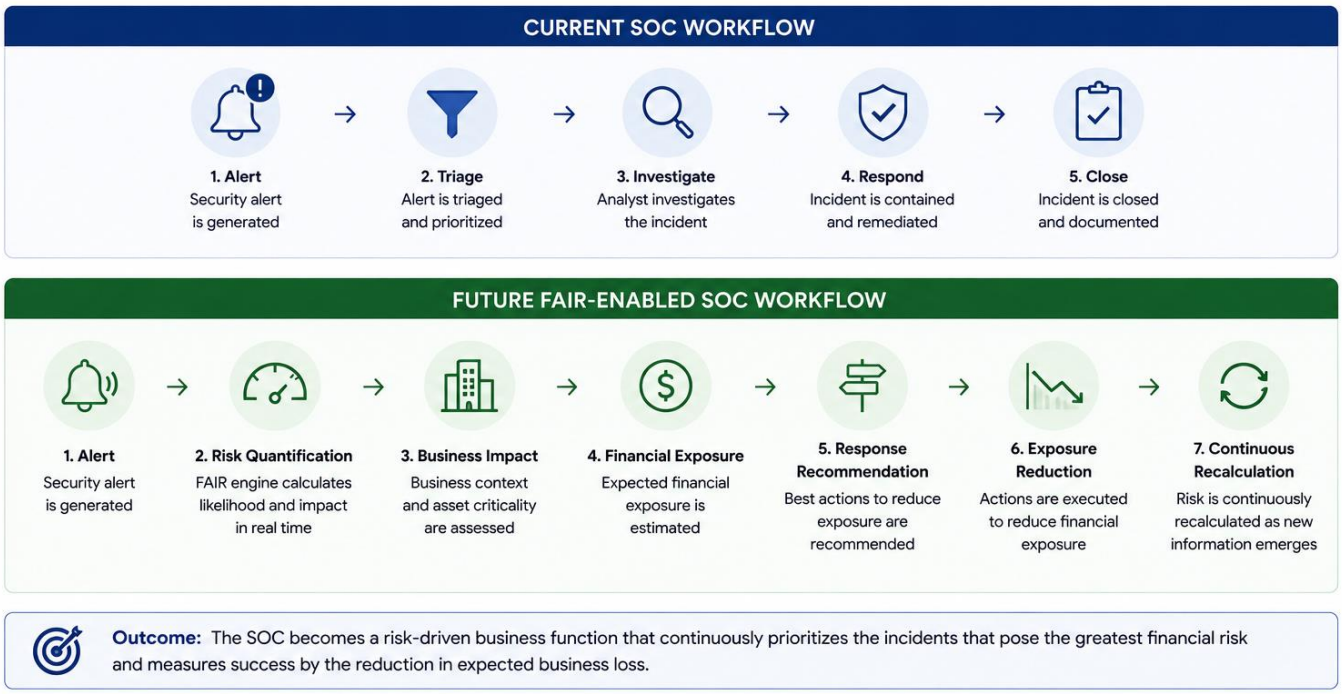

Outcome: The SOC becomes a risk-driven business function that continuously prioritizes the incidents that pose the greatest financial risk and measures success by the reduction in expected business loss.

Integrating AI into Quantified Security Operations

The increasing complexity, speed, and scale of modern cyberattacks have exposed the limitations of traditional Security Operations Center (SOC) workflows. Analysts are expected to process thousands of alerts daily, correlate telemetry from dozens of disparate security tools, evaluate evolving attacker behavior, and determine which incidents warrant immediate attention. Even with mature automation capabilities, the volume of data often exceeds the capacity of human analysts to make consistently informed decisions in real time. As organizations move toward continuous cyber risk quantification using the **Factor Analysis of Information Risk (FAIR™)** model, Artificial Intelligence (AI) is emerging not as a replacement for human expertise, but as the analytical engine capable of continuously calculating, updating, and contextualizing organizational risk at machine speed (NIST, 2024).

Unlike traditional Security Orchestration, Automation and Response (SOAR) platforms that primarily automate predefined workflows, AI-powered security operations can continuously ingest telemetry, correlate attack paths, evaluate changing business context, and estimate the financial implications of every investigation. Rather than simply automating repetitive tasks, AI enables the continuous updating of FAIR variables—including Threat Event Frequency (TEF), Vulnerability, Loss Event

Frequency (LEF), and Loss Magnitude (LM)—as new evidence becomes available. Every new endpoint alert, identity anomaly, threat intelligence feed, vulnerability disclosure, or analyst observation contributes to a constantly evolving estimate of organizational exposure. The result is a Security Operations Center that continuously reprioritizes investigations based on expected business risk instead of static alert severity.

AI as the Quantitative Risk Analyst

One of the greatest strengths of AI lies in its ability to analyze relationships that would be impossible for human analysts to evaluate manually. Modern cyber incidents rarely consist of isolated alerts; they involve hundreds or thousands of interconnected events occurring across identities, endpoints, cloud environments, networks, and business applications. AI can continuously correlate these data sources, identify attack progression, recognize patterns consistent with known adversary behaviors, and determine how each new observation affects the organization's financial exposure.

This capability fundamentally changes the role of the SOC analyst. Rather than spending valuable time collecting evidence from multiple consoles and manually assessing incident priority, analysts receive continuously updated recommendations based on quantified business risk. AI becomes the quantitative analyst, while human investigators provide contextual judgment, validate assumptions, and authorize response actions. This human-in-the-loop approach ensures that automation accelerates decision-making without removing accountability for high-impact operational decisions.

From Automation to Intelligent Decision Support

The next generation of AI-enabled SOCs extends beyond automation by providing continuous decision support throughout the investigation lifecycle. As telemetry is ingested, AI can automatically evaluate changes in asset criticality, monitor identity privilege levels, identify emerging attack paths, assess the effectiveness of existing security controls, and estimate the reduction in financial exposure associated with different response options.

For example, isolating a compromised endpoint may reduce expected financial loss by only a marginal amount if an attacker has already obtained privileged credentials elsewhere in the environment. Conversely, rapidly disabling a compromised privileged identity could significantly reduce the probability of ransomware deployment or data exfiltration. AI continuously models these scenarios, allowing security teams to prioritize the actions that produce the greatest reduction in enterprise risk rather than simply the fastest operational response.

LMNTRIX: Enabling Continuous FAIR-Based Security Operations

LMNTRIX is particularly well positioned to capitalize on this evolution in security operations because its platform already combines autonomous AI with continuous human oversight through its complementary AI capabilities, **Artemis** and **LISA**. Together, these technologies provide many of the foundational capabilities required to operationalize continuous FAIR-based risk quantification inside the SOC.

Artemis functions as the autonomous investigative engine, continuously ingesting telemetry from endpoints, identities, cloud environments, threat intelligence, and security controls to identify malicious activity, correlate related events, and execute investigations at machine speed. Rather than treating alerts as isolated events, Artemis builds a contextual understanding of attacker behavior, mapping activity across the MITRE ATT&CK® framework while continuously evaluating how each new finding influences organizational exposure. This rich contextual awareness creates an ideal foundation for dynamically updating FAIR variables such as Threat Event Frequency, Vulnerability, Loss Event Frequency, and Loss Magnitude as investigations evolve.

Complementing Artemis is **LISA**, LMNTRIX's natural language AI interface that bridges complex cybersecurity operations with human decision-making. Through conversational interactions in platforms such as Microsoft Teams and Slack, LISA enables analysts, incident responders, and business stakeholders to query investigations, understand why incidents have been prioritized, explore estimated business impact, and review recommended response actions using language that extends beyond technical indicators. As FAIR introduces financial risk modeling into day-to-day SOC operations, LISA has the potential to translate quantitative risk calculations into business-focused insights that are readily understood by security leaders and executive stakeholders alike.

Together, Artemis and LISA support a model in which AI performs continuous quantitative analysis while experienced analysts retain strategic oversight. Artemis rapidly processes vast volumes of telemetry, evaluates attacker behavior, identifies attack paths, and recommends containment actions. LISA provides transparency by explaining the rationale behind prioritization decisions, exposing the assumptions behind risk calculations, and enabling analysts to challenge or refine AI-generated recommendations before action is taken. This combination reflects an increasingly important principle within AI-enabled cybersecurity: AI should augment human expertise rather than replace it.

Human-Led AI for the Future SOC

This approach closely aligns with LMNTRIX's philosophy of **Human-Led AI**, where artificial intelligence accelerates investigation, quantifies business risk, and recommends optimized response actions while human analysts remain responsible for governance, validation, and strategic decision-making. Rather than automating security for its own sake, AI becomes an intelligent decision-support capability that continuously measures how every investigative action affects enterprise cyber risk.

As organizations increasingly seek to align cybersecurity investments with measurable business outcomes, platforms capable of combining continuous telemetry analysis, FAIR-based risk quantification, and explainable AI will become increasingly valuable. By integrating autonomous investigation through Artemis with human-centric decision support through LISA, LMNTRIX is well positioned to help organizations transition from alert-driven security operations toward continuously quantified, financially informed cyber defense. In doing so, the SOC evolves from a reactive operational function into a strategic business capability focused on reducing enterprise risk, improving resilience, and maximizing the value of every security decision.

CONCLUSION

Security Operations Centers are approaching a fundamental inflection point. As cyber threats continue to grow in complexity, speed, and financial impact, traditional operational metrics such as alert volumes, Mean Time to Detect (MTTD), and Mean Time to Respond (MTTR) are no longer sufficient measures of success. While these metrics provide valuable insight into operational performance, they offer limited visibility into the one question that matters most to business leaders: **How much cyber risk has been reduced?**

By integrating FAIR into day-to-day security operations, organizations can transform technical investigations into measurable business decisions. Continuous risk quantification enables every alert, intelligence update, and investigative finding to contribute to a real-time assessment of financial exposure, allowing security teams to prioritize the incidents that present the greatest organizational risk. Rather than operating as alert-processing functions, modern SOCs can become intelligent business risk reduction engines that continuously align security operations with enterprise objectives.

Artificial intelligence further accelerates this transformation by providing the computational capability required to continuously evaluate FAIR variables, correlate telemetry, model attack paths, and recommend response actions at machine speed. Combined with human expertise and oversight, AI enables faster, more consistent, and financially informed decision-making without sacrificing governance or accountability.

Organizations that embrace this evolution will be better positioned to maximize the value of limited security resources, improve executive visibility into cyber risk, and demonstrate measurable reductions in expected financial loss. The future of security operations will not be defined by the number of alerts processed, but by the organization's ability to continuously quantify, prioritize, and reduce enterprise cyber risk in support of broader business resilience.

REFERENCES

FAIR Institute. (2025). *FAIR Institute Knowledge Base*. https://www.fairinstitute.org

Google Cloud. (2025). *Threat Horizons Report*. https://cloud.google.com/security/resources/threat-horizons

IBM. (2025). *Cost of a Data Breach Report 2025*.

National Institute of Standards and Technology. (2024). *Cybersecurity Framework (CSF) 2.0*. https://www.nist.gov/cyberframework

The Open Group. (2023). *The Open FAIR™ Standard*. https://www.opengroup.org/openfair

Verizon. (2025). *2025 Data Breach Investigations Report*. https://www.verizon.com/business/resources/reports/dbir/

Jones, J. (2020). *Book: Measuring and Managing Information Risk*. Wiley.

MITRE. (2025). *MITRE ATT&CK® Framework*. https://attack.mitre.org

Carnegie Mellon Software Engineering Institute. (2024). *Cyber Risk and Resilience Resources*. https://www.sei.cmu.edu

National Institute of Standards and Technology. (2024). *AI Risk Management Framework (AI RMF 1.0)*. https://www.nist.gov/itl/ai-risk-management-framework

OWASP Foundation. (2025). *OWASP Top 10 for Large Language Model Applications*. https://owasp.org/www-project-top-10-for-large-language-model-applications/