Dataminr

Mending the Broken Cyber Defense Chain

Reconnecting Threat, Posture, and Business Impact to Drive Decisive Action

Dan Cole, Product Marketing Vice President

The architecture described in this book is what the industry actually needs: intelligence tailored to your environment, risk measured in dollars instead of severity scores, and action that flows from understanding rather than panic."

— Jen Miller-Osborn, Co-creator of MITRE ATT&CK

Dataminr

Threat Landscape

Security Posture

Business Impact

One team detects threats. One scans for exposures. Another models the risk. What they can't do — fast enough — is connect all three and act before the adversary gets there first.

57% **of breaches are discovered by someone outside the organization — not by the security team's own defenses.**

[Mandiant, M-Trends 2025 Report]

$2,209,100 **— how much more a breach costs for every 10 days it goes undetected.**

[Dataminr Predictive Threat Exposure Management, proprietary cyber loss data]

0 days **— the effective dwell time between breach and exfiltration in many 2025 attacks. The window for response is collapsing.**

[Dataminr, 2026 Cyber Threat Landscape Report]

Foreword

I'll be direct about something: frameworks don't defend companies. People do. And people can't defend what they can't see coming.

I helped build one of the frameworks this book references — MITRE ATT&CK was an attempt to give defenders a shared language for how adversaries actually operate. I've spent the years since tracking those adversaries in practice: nation-state campaigns, ransomware operations, zero-days that moved faster than the teams trying to contain them. And across every role — analyst, team lead, researcher, executive — the same problem kept showing up.

It wasn't detection. We're actually quite good at that. The problem was what happened after detection — when the signal had to travel from "something is happening in the world" to "here's what it means for us, what it costs us, and what we do about it." That journey broke down in almost every organization I worked with. Not because people weren't talented. Because the systems weren't built to make that journey in one motion.

When I testified before the Senate about Log4Shell, my message was that we needed operational collaboration, not more information. The signals were there. The connections between them were not. That was 2022. The problem hasn't gotten better — it's gotten faster.

That conviction is what drew me to Dataminr. Not a vendor pitch. A belief, grounded in operational experience, that the architecture described in this book is what the industry actually needs: intelligence tailored to your environment, risk measured in dollars instead of severity scores, and action that flows from understanding rather than panic.

Dan and the team have laid out the problem clearly, with evidence, and with a credible path forward. Read it as a practitioner's argument for how defense should actually work — because that's what it is.

— Jen Miller-Osborn

Jen Miller-Osborn co-created the MITRE ATT&CK framework, has testified before the U.S. Senate on cyber defense, and has spent two decades on the frontlines of threat intelligence — from the U.S. Air Force to MITRE to Unit 42 to Dataminr, where she now serves as Field Cyber Intelligence Officer.

Who This Book Is For

Every security program has links in a chain — threat feeds, detection tools, risk frameworks, automation. What most programs lack is an intelligent system that connects them fast enough to make a difference when it matters (i.e. before impact!). When real threats are breaking out and the clock is running — most security teams can't reliably turn what they know into what they do.

The chain between threat intelligence, security posture, and business impact is broken. This book explains why, shows what a mended chain looks like, and introduces the intel-driven threat and exposure management solution suite designed to mend it — from first signal to risk-prioritized action.

If you lead a SOC or incident response (IR) team, you'll see how to close the distance between early warning and confident action — without adding more tools or more headcount to the problem.

If you run a vulnerability management program, you'll see how to stop chasing CVSS scores and start prioritizing by what a vulnerability would actually cost the business — with continuous evidence that your controls are working, not just assumed to be.

If you run a threat intelligence function, you'll see how to make intelligence operational and compounding — instead of rebuilt from scratch every time something breaks.

If you're a CISO or risk leader, you'll see how to connect what your security team knows to core business objectives — in language your board can actually act on.

If you've ever had the right data in the right systems and still couldn't act fast enough, this book is for you.

CHAPTER 1

The Broken Chain

The Paradox of More

Security programs have never been better funded, better staffed, or better equipped. But somehow, better outcomes haven't kept up.

The average enterprise now operates 83 different security solutions from 29 vendors.[3] Teams have threat feeds streaming in from multiple sources. Detection engines flag thousands of events per day. Risk frameworks exist for nearly every compliance mandate and board presentation imaginable. There is no shortage of capability.

And yet: Security teams report an average of 70 minutes to fully investigate a single alert.[4] That doesn't sound like a crisis until you realize what happens next. While analysts are pulling context on one alert, hundreds more alerts are piling up behind it. Industry research shows 62% of SOC alerts are simply disregarded — not because they don't matter, but because nobody has time to figure out whether they do.[5]

The result isn't just inefficiency. It's exposure. In 2025, Dataminr tracked a 225% increase in average monthly threat actor alerts — from 1,490 per month in 2024 to 4,840 per month, with the busiest month in 2024 not even reaching the quietest month.[6] That's the pace of the threat. The pace of the defense hasn't matched it.

We don't have a tools problem. We have a connection problem.

[3] IBM Institute for Business Value, "Capturing the cybersecurity dividend," 2025
[4] Prophet AI, State of AI in Security Operations 2025
[5] Vectra, "2024 State of Threat Detection"
[6] Dataminr, "2026 Cyber Threat Landscape Report"

Three Links That Should Hold Together

Effective cyber defense has always depended on three things working in concert — three links in a chain that only holds when they're connected.

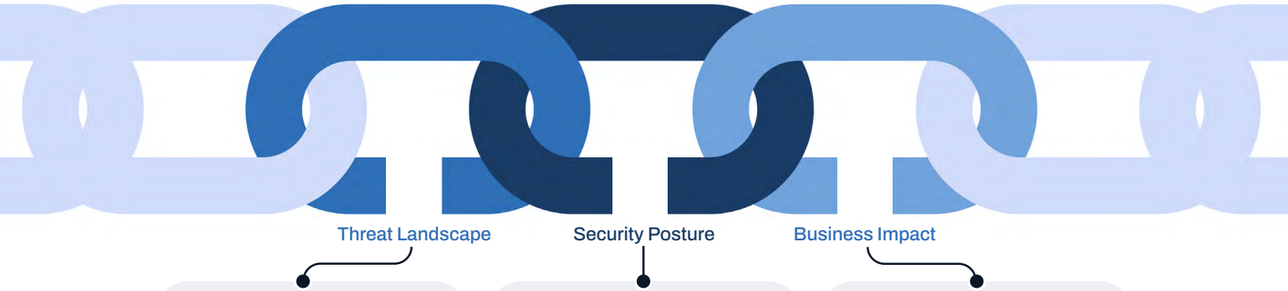

Threat Landscape is the first link: knowing what's coming, early enough to do something about it. Not just indicators of compromise (IoC) feeds or Common Vulnerability and Exposure (CVE) lists, but genuine foresight into which actors are targeting your industry, which exploits are gaining traction, and which campaigns are likely to land next. In 2025, Dataminr tracked more than 5,000 threat actors and generated over 6.3 million external threat alerts across malware, hacktivism, ransomware, and advanced persistent threats.[7] The signals are there. The question is whether security teams can find them fast enough before it's too late.

Security Posture is the second link: understanding your attack surface, where you're actually exposed inside your environment, and how well your defenses mitigate these exposures (if they do at all!). This isn't a quarterly vulnerability scan. It's a continuous answer to the question, "If that threat landed here, today, how would we fare?" We tracked more than 4.8 million vulnerability-related alerts in 2025 alone.[8] Most security teams lack the capacity to assess which of those vulnerabilities actually matter to their specific infrastructure — so they try to patch everything and end up patching the wrong things first.

Business Impact is the third link: measuring the scope and severity of the impact for the organization should a given threat get through. Not in abstract threat scores, but in actual financial impact, downtime, and regulatory repercussions specific to your unique organization. As the Dataminr 2026 Cyber Threat Landscape Report found, technical severity does not always equal business risk. A vulnerability rated "Critical" by CVSS may carry bounded financial impact in your environment, while a "High" may sit on a system that, if compromised, could trigger a nine-figure loss. This is the link that translates a CVSS 9.8 into a board-level conversation — and it's the link most teams don't have at all.

In most organizations, these three links are owned by different teams, powered by different tools, updated at different speeds, and reported in different languages.
The chain is broken.

[7] Dataminr, "2026 Cyber Threat Landscape Report"
[8] Dataminr, "2026 Cyber Threat Landscape Report"

What Happens When the Chain Breaks

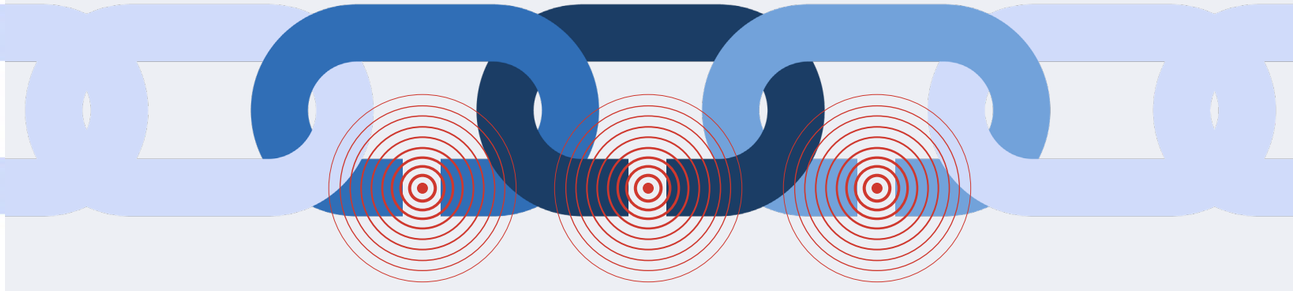

The consequences aren't theoretical. They show up every day in the gaps between what security teams know and what they can do about it.

Intelligence arrives, but nobody can answer whether it applies to their environment without hours of manual investigation. Posture is assessed on a quarterly cycle while threat actors move in hours — the dwell time between initial breach and data exfiltration has, in many cases, effectively been **reduced to zero**.[9] Risk gets discussed in executive meetings using static frameworks that don't reflect what's actually happening this week.

The data tells the story clearly. Security teams missed 57% of cyberattacks compromising their organizations — learning about the initial breach from an outside party, an attacker's own disclosure, or a third-party researcher rather than from their own defenses.[10] Think about that. More than half the time, someone else told them they'd been breached.

Meanwhile, the financial stakes are compounding. In 2025, the cyber loss landscape shifted dramatically — from frequent but contained incidents to fewer events with materially larger impact. Single-incident losses surged, driven by multi-vector attacks combining credential theft, data exfiltration, operational disruption, and regulatory exposure into one event. **Losses at the $100 million and $1 billion-plus level,** virtually absent in 2024, appeared with increasing frequency. This is the era of the "mega-loss" — and disconnected defenses are what make it possible.

The cost of disconnection isn't abstract. Based on proprietary Dataminr cyberloss data, every 10 days a data breach goes undetected increases the total expected financial loss by $2,209,100.[11] That's the price of a broken chain — and it compounds with every gap between signal and action.

Every response in a broken chain starts the same way: rebuilding context from scratch. Searching across disconnected tools for information that should have been connected before the alert ever fired. Trying to answer "how bad is this, for us, right now" with data that was never designed to be assembled under pressure.

The tools aren't the problem. The disconnection between them is.

Security teams missed 57% of cyberattacks compromising their organizations in 2024, only learning about the initial compromise from an external entity.[12]

[9] Dataminr, "2026 Cyber Threat Landscape Report"
[10] Mandiant, "M-Trends 2025 Report"
[11] Dataminr Continuous Control Monitoring with Risk Quantification proprietary cyber loss data
[12] Mandiant, "M-Trends 2025 Report"

The Daily Cost of Disconnection

The industry benchmark is

1 minute to detect | **10** minutes to investigate | **60** minutes to contain

But 95% of organizations don't meet it!![13] Why not?
Here's what often happens:

A Critical Vulnerability Through a Disconnected Organization

STEP 1
THE SIGNAL

Advisory Published

What happens: A vendor or CERT discloses a critical vulnerability in software widely deployed across your industry. A CVSS 9.8. Your threat intel feed fires an alert.

By the numbers: 56% of teams spend more than 20% of their day just deciding which alerts to address first.[15]

Meanwhile: The alert includes IOCs and a severity score — but no mapping to your environment. It sits alongside dozens of other alerts from the same morning.

STEP 2
TRIAGE

An Analyst Picks It Up

What happens: An analyst reads the advisory and asks the first question: "Do we run this software?" Answering requires querying the CMDB or asset inventory in a separate tool. The TI platform doesn't connect to it.

By the numbers: Security teams report an average of 70 minutes to fully investigate a single alert.[16]

Meanwhile: The analyst pivots between three or four tools — TI platform, asset inventory, vulnerability scanner, ticketing system — copying and pasting context between them.

STEP 3
EXPOSURE CHECK

Confirming Whether You're Affected

What happens: The analyst queries the vulnerability management platform (Qualys, Tenable, Rapid7) to check for affected versions. The scan data may be hours or days old. The CMDB may be incomplete. Shadow IT assets won't appear at all.

By the numbers: 82% of cybersecurity professionals report gaps in data visibility across their environment.[17] Even confirmed exposure doesn't answer the next question: how critical are the affected systems to the business?

Meanwhile: The vulnerability is confirmed on multiple systems. But without business context, the analyst has no way to distinguish a CVSS 9.8 on a dev sandbox from a CVSS 9.8 on a payment processing server. Both get the same severity label.

STEP 4
ESCALATION

Getting It Up the Chain

What happens: The analyst escalates via Slack, email, or a ticket. The context they've spent the last stretch building — which advisory, which assets, how exposed — gets compressed into a few sentences. The SOC manager or team lead reviews and decides whether it warrants executive attention.

By the numbers: 55% of teams have missed critical alerts due to ineffective prioritization. Not because the alert didn't fire. Because it couldn't compete for attention against everything else in the queue.[18]

Meanwhile: The escalation sits in a queue. The decision-maker has their own backlog. When they pick it up, they re-read the advisory from scratch because the compressed ticket doesn't carry the full context.

STEP 5
THE QUESTIONS NOBODY CAN ANSWER QUICKLY

Decision Time

What happens: A decision-maker — SOC manager, CISO, or risk lead — reviews the escalation and asks: "What's our actual financial exposure if this gets exploited? Are threat actors actively targeting this in our industry right now?" The team doesn't have fast answers because the threat intel, posture data, and risk models live in different systems with no shared context.

By the numbers: 67% of security professionals believe their senior leadership underestimates cyber threats — partly because the data needed to frame risk in business terms isn't available at decision time.[19]

Meanwhile: Someone goes back to build a business case manually — pulling loss estimates from a spreadsheet, searching for industry targeting data, assembling a narrative for leadership. The clock keeps running.

STEP 6
EVENTUAL ACTION

Patching Gets Approved

What happens: The patch is approved and enters the change management queue. Emergency patches still require testing and a maintenance window in most enterprises.

By the numbers: At 30 days after disclosure, 85% of Known Exploited Vulnerabilities remain unremediated. It takes organizations an average of 55 days to patch just 50% of their critical vulnerabilities.[20]

Meanwhile: Adversaries don't have change advisory boards.

[13] CrowdStrike, 2019 Global Security Attitude Survey
[14] CrowdStrike, 2019 Global Security Attitude Survey
[15] Mandiant, Global Perspectives on Threat Intelligence, 2023
[16] Prophet AI, State of AI in Security Operations 2025
[17] Bedrock Security, "2025 Enterprise Data Security Confidence Index"
[18] Mandiant, Global Perspectives on Threat Intelligence, 2023
[19] Mandiant, Global Perspectives on Threat Intelligence, 2023
[20] Verizon, 2024 Data Breach Investigations Report

What was lost along the way:

Foresight —
The signal arrived, but it wasn't connected to the environment fast enough to drive immediate action.

Focus —
Hours were spent confirming exposure and building business context that should have been available from the start.

Action —
The decision stalled because the people who needed to act didn't have the data to justify acting fast.

Every tool in this workflow did its job.
The problem is that no one connected them.

CHAPTER 2

Why the Chain Is Broken

The Point-Tool Era Created the Gaps

The broken chain wasn't inevitable. It was built — one best-of-breed purchase at a time.

Over the last decade, security programs invested heavily in each link independently. Best-in-class threat feeds. Leading vulnerability scanners. Sophisticated risk models. Each tool solved a real problem. Each purchase was defensible. But each one created a new gap that nobody was responsible for closing.

The result: three critical gaps between the three links, each one a structural failure that no individual tool was designed to fix.

The gap between threat and posture. Threat intelligence is global and generic. Your environment is local and specific. When a new exploit surfaces across the global threat landscape, intelligence feeds may tell you the threat exists — but the information often arrives late and without context. Threat feeds don't tell you whether you're running the affected software, whether it's internet-facing, or whether your compensating controls already mitigate the risk. That translation is manual. In an internal Dataminr customer survey fielded in 2025, 83% of security and CTI analysts said they struggled to find relevant information quickly during investigations.[21] Not because the information didn't exist. Because it lived in a different system, owned by a different team, updated on a different cadence.

The gap between posture and impact. Vulnerability scanners produce findings. Risk models produce scores. But the two rarely talk to each other in real time. When your scanner surfaces 200 critical vulnerabilities, the question isn't which ones have the highest CVSS — it's which ones would cost the business the most if exploited. That answer requires connecting asset criticality, control effectiveness, and financial exposure in a way that most organizations still do manually, if they do it at all. In the same internally-fielded Dataminr customer survey, 80% of respondents said they struggled to prioritize threats by business impact before implementing cyber risk quantification. This is the gap where CVSS 9.8 vulnerabilities on dev sandboxes compete for the same resources as CVSS 7.5 vulnerabilities on payment-processing servers — because the data to distinguish between them doesn't exist in one place.

The gap between impact and action. Even when risk is understood, it often stays stranded in dashboards and reports. The intelligence that could drive a patching decision sits in a threat intel portal. The vulnerability data sits in a scanner. The business context sits in a spreadsheet or a quarterly risk assessment that's already weeks stale. By the time someone assembles these into an actionable recommendation, the window for decisive action has narrowed — or closed entirely. Dashboards don't defend companies when the adversary moves faster than the reporting cycle.

These gaps aren't failures of any one team or tool. They're the natural consequence of an industry that optimized each link in the chain without ever connecting them.

[21] Internal Dataminr customer survey fielded in 2025.

From Feeds to Foresight: The Evolution of Cyber Intelligence

The Industry Is Converging on an Answer: Data Fabrics and Exposure Management

Two frameworks from Gartner have emerged as the clearest articulation of what needs to change.

Unified Cyber Risk Intelligence (UCRI) is an architectural model for fusing internal and external signals into a single intelligence fabric that supports decision-making across the enterprise. The core insight is that intelligence must be unified across sources, contextualized to the organization, and delivered to all security and business functions — not siloed in a TI platform that only analysts see. UCRI calls for three core capabilities: broad collection across varied data types, contextual analysis and fusion to produce a corroborated fact base, and dissemination to all stakeholders so decisions from the SOC to the boardroom are driven by the same risk-aware intelligence.

> **Looking ahead, the future of threat intelligence is unified cyber risk intelligence (UCRI) and will be defined by the convergence of multisignal collection and advanced analytical capabilities ... This multisignal approach, combined with the evolution of AI techniques (e.g., ML and NLP), will enable faster, more accurate detection of emerging threats and subtle attack patterns."[22]**
>
> — Gartner

[22] Gartner, "The Future of Threat Intelligence Is Unified Cyber Risk Intelligence," September 15, 2025

Continuous Threat Exposure Management (CTEM) is an operational model for continuously identifying, validating, and prioritizing exposure. Where UCRI describes what needs to be connected, CTEM describes how it must operate: not as a periodic assessment or an annual penetration test, but as a continuous cycle of scoping, discovery, prioritization, validation, and mobilization.

> **[CTEM] is a strategic framework designed to help organizations proactively manage and mitigate security exposures across their digital landscape. CTEM acts like a seasoned navigator, charting the course through the complex and ever-changing seas of cybersecurity threats to ensure safe passage for the organization."[23]**
>
> — Gartner

Together, these frameworks represent the industry catching up to what practitioners have been asking for. UCRI is the blueprint for the architecture. CTEM is the blueprint for the operating rhythm. Both point toward the same conclusion: the chain between the threat landscape, security posture, and business impact must be reconnected — and it must stay connected continuously.

The data validates that convergence. The Dataminr 2026 Cyber Threat Landscape Report found that 2025 marked a clear shift toward fewer total cyber incidents, but incidents that did occur had a materially larger financial impact — multi-vector attacks combining credential theft, data exfiltration, operational disruption, and regulatory exposure into a single event.[24] When single incidents can reach nine figures, the cost of disconnected intelligence isn't measured in efficiency. It's measured on the balance sheet.

Frameworks Aren't Enough

UCRI and CTEM represent genuine progress. They name the right problems. They describe the right architecture. They've given the industry a shared language for what modern cyber defense should look like.

But frameworks describe the destination, not the vehicle.

Most organizations that adopt these models still face the same practical challenge: how do you actually operationalize them? UCRI tells you to unify internal and external intelligence. It doesn't tell you how to fuse real-time threat actor chatter with your specific asset inventory and control posture — at speed, continuously, across millions of signals. CTEM tells you to prioritize exposure continuously. It doesn't tell you how to translate a trending exploit into a dollar-denominated business risk for your environment in time to act on it.

The frameworks define what and how. They don't answer with what.

That's not a criticism. It's the natural boundary of any architectural model. Frameworks set the standard. Systems meet it.

The question is no longer whether these connections need to exist. It's whether there's a system that can actually deliver them — continuously, at scale, and tailored to each organization's reality.

[22] Gartner, "The Future of Threat Intelligence Is Unified Cyber Risk Intelligence," September 15, 2025
[23] Gartner, "Guide to Security Operations Concepts," November 7, 2025
[24] Dataminr, "2026 Cyber Threat Landscape Report"

Three Gaps, Three Challenges

The Chain Doesn't Break in One Place. It Breaks in Three.

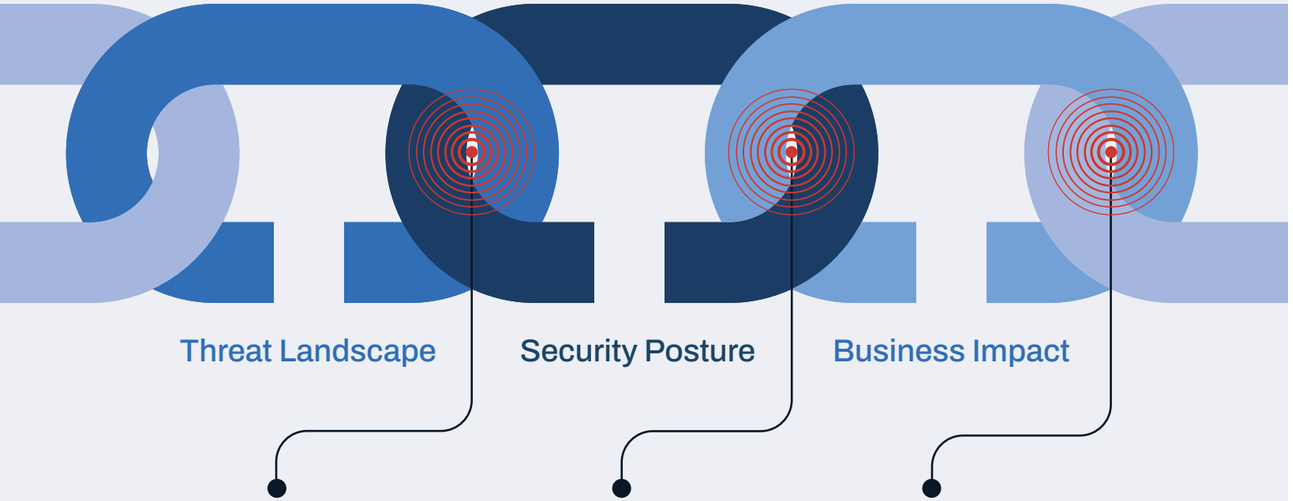

Intelligence is global. Your environment is specific. No automated bridge between them.

Vulnerability data exists. Financial exposure doesn't. Priority is driven by CVSS and gut feel.

Even when risk is understood, it's stranded in dashboards — not flowing into the workflows where decisions are made.

CHAPTER 3

What a Mended Chain Looks Like

Three Requirements for Reconnecting the Chain

The previous chapters named the problem and its cause. The chain between the threat landscape, security posture, and business impact is broken — and it broke because the industry spent a decade optimizing each link without connecting them. Gartner's UCRI and CTEM frameworks describe the architecture for putting it back together.

But what does a mended chain actually look like in daily operations?

It requires three things. Not three tools, not three dashboards — three capabilities that must work together continuously.

Foresight: Know what's coming — early enough to change the outcome.

Foresight is the ability to detect and understand emerging threats while they're still forming. Not just early alerting — early understanding. A generic advisory that a new exploit exists is useful. Knowing that the exploit targets software running in your environment, that it's being actively discussed by threat actors in your industry, and that your specific configuration is vulnerable — that's foresight.

In 2025, Dataminr tracked a **225% increase in monthly threat actor alerts** year-over-year.[25] The volume of signals is accelerating. Foresight compresses the time between "something is happening in the world" and "we know what it means for us" — before the adversary gets there first.

Focus: Know what it costs — grounded in evidence, not severity scores.

Focus is the ability to determine which threats and exposures actually matter to your organization, measured in business impact rather than technical severity. When your scanner surfaces dozens of critical vulnerabilities in a week, focus is what tells you which one threatens a nine-figure loss and which one sits behind four layers of compensating controls on a system that processes nothing sensitive.

The Dataminr 2026 Cyber Threat Landscape Report found that technical severity does not always equal business risk — a consistent observation across the vulnerability landscape in 2025.[26] A CVSS 7.9 with no authentication required and broad industry exposure can carry catastrophic financial risk. A CVSS 9.8 on a system behind segmentation and zero-trust architecture may carry very little. Focus is what makes that distinction visible and measurable.

Action: Know what to do — without rebuilding context from scratch.

Action is the ability to turn intelligence and risk insight into consistent, timely response. Not "generate an alert and hope the right person sees it" — but ensuring that understanding actually reaches the people and workflows where decisions are made, with enough context to act immediately.

In 2025, 98% of security teams said they need to be faster at implementing changes to their cybersecurity strategy based on the latest threat intelligence.[27] The bottleneck isn't willingness. It's that the intelligence, the exposure data, and the business context arrive separately, at different times, in different formats — and someone has to assemble them by hand before anything happens.

[25] Dataminr, "2026 Cyber Threat Landscape Report"
[26] Dataminr, "2026 Cyber Threat Landscape Report"
[27] Mandiant, Global Perspectives on Threat Intelligence, 2023

Why These Must Work Together

Foresight without Focus creates noise. You know something's coming, but you can't tell whether it matters to you — so everything gets treated as urgent, and nothing gets treated well.

Focus without Action creates delay. You know exactly which exposure threatens the business, but the insight sits in a risk model while the SOC works a different queue.

Action without Foresight and Focus. You move fast — but only after someone else tells you there's a problem, because you didn't see the threat forming early enough to get ahead of it.

A mended chain isn't three separate improvements bolted onto the existing stack. It's a system where intelligence, posture, and impact inform each other continuously. When a new threat emerges, the organization should immediately know whether it applies, what it would cost, and what to do — without anyone rebuilding context by hand.

The Mended Chain in the Wild

What does this look like in practice?
Three scenarios — each showing what changes when the link in the chain holds.

Foresight

Foresight in the Wild: Detecting an Emerging Exploit Before It Lands

A threat actor begins circulating proof-of-concept code for a new deserialization vulnerability on a dark-web forum. Within the hour, scanning activity spikes globally as attackers probe for exposed systems.

In a broken chain, this surfaces as a generic advisory — "new zero-day, CVSS 9.8, patch when available." The SOC adds it to the backlog. Somebody will look at it when they get to it.

In a mended chain, the response is immediate and specific. The system correlates the exploit's prerequisites against your internal telemetry and identifies that a subset of your customer-facing API gateways use the vulnerable library version. Instead of a generic alert, the team gets foresight: these three gateways are exposed, they are reachable from the internet, here is the observed exploitation behavior globally, and here's what it would mean operationally if the exploit lands here. The distance between "something is happening" and "here's what it means for us" collapses from hours to seconds.

Focus

Focus in the Wild: Prioritizing by Expected Loss, Not Noise

Your vulnerability scanner surfaces dozens of new critical findings. Two appear identical on paper — both CVSS 9.8, both externally referenced, both generating noise in the feed.

In a broken chain, both get the same severity label. Both compete for the same resources. The patch team works top-down by CVSS score and hopes for the best.

In a mended chain, focus reveals their true difference. Vulnerability A affects a development VM behind four layers of compensating controls. Vulnerability B affects the payment-card microservice cluster and is trending globally with confirmed exploitation. The system quantifies the expected loss: $2.3 million for Vulnerability B given the asset value, internal pathways, and business mission — versus $50,000 for Vulnerability A (Note: These are modeled figures, illustrative of risk quantification output). Security leadership immediately sees the priority: "This one is worth fixing today. That one can wait."

Action

Risk-Informed Action in the Wild: Knowing When Not to Panic

A critical vulnerability begins trending. Headlines claim widespread exploitation is imminent. In most organizations, the outcome is predictable: emergency patch calls, broad scanning, and a flood of tickets — before anyone can answer whether the risk is real for them.

In a mended chain, the response starts with the organization's own reality, not the headlines. The system evaluates your environment and finds: you have several instances of the affected software, but compensating controls — a WAF rule and egress restrictions — mitigate most of them. Except one. A legacy application that sits outside the control coverage, supports a mission-critical business workflow, and is approaching a peak period where downtime carries outsized financial impact.

Instead of mobilizing everyone against everything, the team gets a risk-informed decision package: these assets are effectively mitigated — monitor, don't mobilize. This specific legacy instance is truly exposed — here's the likely exploitation path. These two teams need to be notified first, in this order. Here is the recommended mitigation sequence: control-first, patch second.

This is the scenario most security tools can't deliver — and the one practitioners want most. Not just "act now" but "here's exactly what warrants action, and here's what doesn't." The ability to tell your CISO "we're covered here, but we have one real exposure and here's the plan" is worth more than any dashboard.

Foresight, Focus, and Action — The Three Forces That Hold The Chain Together

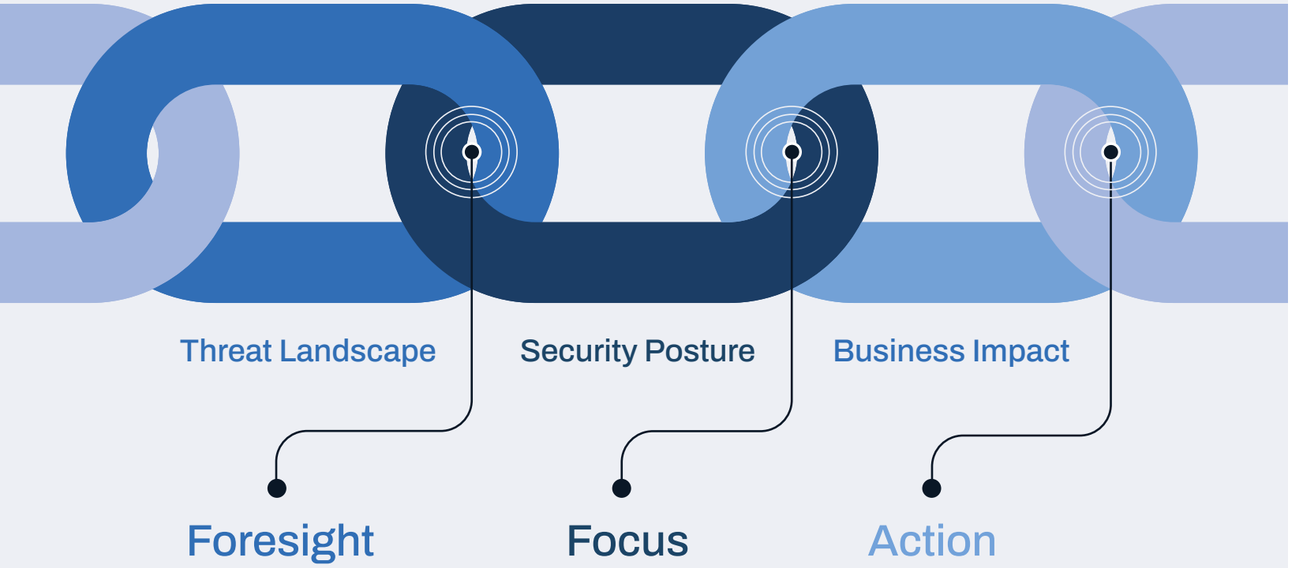

CHAPTER 4

How Dataminr for Cyber Defense Mends the Chain

From Framework to Operating System

UCRI and CTEM define the architecture. Foresight, Focus, and Action define the requirements. What's been missing is a system that delivers all of it — end to end, continuously, and tailored to each organization's reality.

That's what Dataminr for Cyber Defense was built to do.

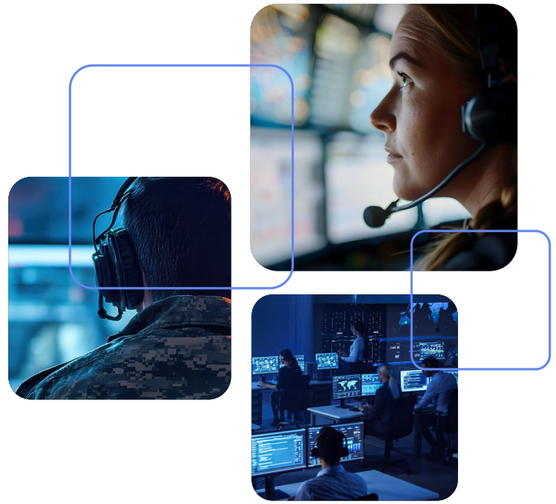

Dataminr for Cyber Defense is an intel-driven threat and exposure management solution suite that assembles and operationalizes real-time, client-tailored intelligence from first signal to risk-prioritized action — continuously and at scale.

In practice, that means a unified system — four capabilities feeding a shared intelligence layer powered by agentic AI — designed to reconnect the chain between what's happening externally, what's true inside your environment, and what it means to your business.

Everything the previous chapters described as the future state — the foresight, the financial clarity, the risk-informed action — this is the system that operationalizes it.

The system also operationalizes CTEM's continuous cycle. Where CTEM calls for continuous scoping and discovery, Threat Intelligence and Investigation Insights provide it — monitoring the external landscape and mapping it to internal exposure in real time. Where CTEM calls for prioritization and validation, Continuous Control Monitoring with Risk Quantification delivers it — translating exposure into financial impact and continuously validating whether controls are actually working. And where CTEM calls for mobilization, the TI Platform's automated workflows ensure intelligence reaches the right teams with enough context to act. The cycle doesn't run quarterly or monthly. It runs continuously — which is what CTEM demands but few organizations achieve.

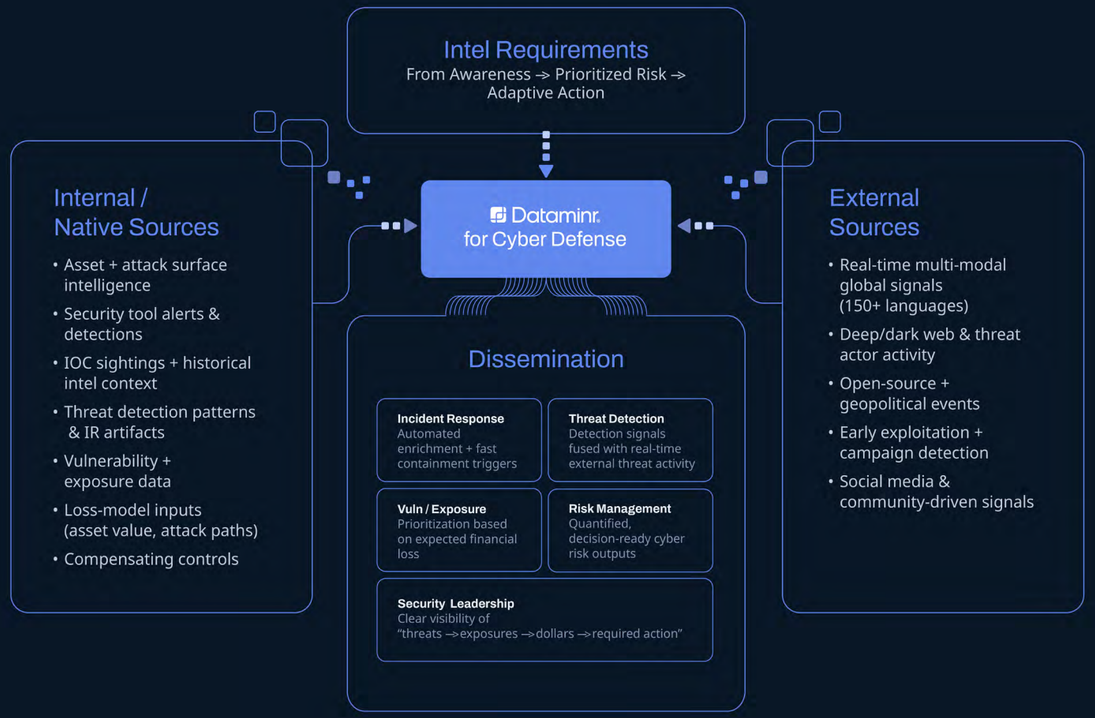

Unified Cyber Risk Intelligence

Dataminr for Cyber Defense maps directly to UCRI's architecture — and extends it. UCRI calls for broad collection across internal and external sources. The system ingests both: internal telemetry including asset inventories, security tool alerts, vulnerability data, compensating controls, and loss-model inputs alongside external signals from 1M+ sources across dark web activity, exploit chatter, geopolitical events, and community-driven intelligence. UCRI calls for contextual fusion into a corroborated fact base. Intel Agents perform that fusion continuously and autonomously. And UCRI calls for dissemination to all security and business stakeholders. The system delivers intelligence directly into incident response, threat detection, vulnerability prioritization, risk management, and security leadership workflows — not as a report to be interpreted, but as decision-ready output tailored to each function's needs.

Four Capabilities
That Reconnect the Chain

Each capability contributes something the system can't function without. Individually, they solve real problems. Together, they close the gaps.

Threat Intelligence: What's emerging in the world.

This is the system's external awareness — the earliest signal that something is changing. Dataminr's Multi-Modal Fusion AI detects emerging threats from more than 1 million public data sources, synthesizing text in 150 languages alongside image, video, and audio signals. It's monitoring the places threat actors actually operate — dark-web forums, Telegram channels, breach sites, data dumps — and surfacing activity often before formal disclosure. In 2025, this capability tracked more than 5,000 threat actors and generated over 6.3 million external threat alerts.[28] In practice: when a threat actor begins advertising a new exploit kit targeting financial services infrastructure on a Russian-language forum, the AI detects the chatter, correlates it with related activity across multiple sources, and delivers a structured intelligence brief to affected customers — often hours or days before any vendor advisory exists. It answers the first question: "What's coming?"

Investigation Insights: What's actually happening in your environment.

This is the system's internal ground truth — and it solves one of the most persistent productivity killers in security operations: the swivel chair. Analysts today spend investigations bouncing between tabs — TI platform, SIEM, vulnerability scanner, CMDB, ticketing system — copying and pasting context between tools that don't talk to each other. Investigation Insights eliminates that pivot. It operates as an intelligent overlay on top of the tools analysts already use, surfacing alerts, asset exposure, prior investigations, and cross-tool context directly in the analyst's workflow — without requiring them to leave their current screen. When a new IOC appears in a SIEM alert, the AI-powered overlay instantly enriches it with threat actor attribution, related campaigns, internal asset matches, and historical context — all without the analyst leaving the screen they're already working in. It answers the question that 83% of teams struggled with before: *"Does this apply to us?"* [29]

The overlay extracts indicators, vulnerabilities, and entity references directly from whatever is on the analyst's screen — a screenshot in a Slack thread, a PDF advisory, a browser tab with a threat report — and retrieves relevant context automatically.

Agentic Threat Intelligence Platform: What you've learned — and how it moves without manual effort.

This is the system's institutional memory and automation engine. It transforms early signals into structured, reusable intelligence that flows consistently into detection engineering, incident response, threat hunting, and executive reporting — automatically. That last word matters. Most TI programs still rely on analysts to manually enrich indicators, build detection rules, distribute intelligence to downstream tools, and update reporting. The TI Platform automates those workflows end to end. AI-driven playbooks trigger enrichment, create cases, push indicators to SIEMs and EDRs, notify the right teams, and update dashboards — without an analyst touching each step. Intelligence doesn't get rebuilt from scratch every time. It compounds — each investigation enriching the knowledge base, each automated workflow informing future prioritization. In practice, when a new ransomware campaign is attributed to a

[28] Dataminr, "2026 Cyber Threat Landscape Report"
[29] Internal Dataminr customer survey fielded in 2025

known actor, the platform automatically pulls in historical TTPs, maps them to existing detections, identifies gaps, and queues detection rule updates — in the time it used to take an analyst to open a second browser tab. It answers the question: "What have we seen before, and what should we do next?"

When the platform handles collection, processing, enrichment, and distribution automatically, analysts stop spending their days assembling intelligence and start spending them analyzing it. The system moves the team from the mechanical left side of the intelligence cycle to the analytical right side — where human judgment actually changes outcomes.

Continuous Control Monitoring with Risk Quantification: What exposure actually costs — and which controls are working.

This is the system's business impact layer. It continuously validates control effectiveness using live telemetry and translates threat exposure into financial risk — not static severity scores, but dynamic, dollar-denominated assessments grounded in your actual asset values, compensating controls, and operational dependencies. The AI continuously models how changes in the threat landscape, control posture, and asset configuration shift your financial exposure — so the risk picture is always current, not a snapshot from last quarter's assessment. In practice: when a new vulnerability is disclosed in software running on your payment infrastructure, Continuous Control Monitoring with Risk Quantification doesn't just flag it as "Critical." It models the expected financial loss given your specific control coverage, compares it against other open exposures competing for the same remediation resources, and delivers a prioritized recommendation the CISO can take to leadership with evidence attached. It answers the question the 80% struggled with: *"What does this cost us, and where will action actually reduce risk?"* [30] Continuous Control Monitoring with Risk Quantification is built on more than 20 years of historical cyber loss data, so the financial models aren't theoretical — they're calibrated against what breaches actually cost in the real world.

[30] Internal Dataminr customer survey fielded in 2025

Six Minutes:
From Signal to Decision

What does the system actually produce? Here's a composite scenario showing all four capabilities working in sequence. (Illustrative scenario based on representative system behavior. Timestamps and financial figures are modeled.)

8:04 AM — External signals detected. Threat Intelligence picks up converging activity: a new proof-of-concept exploit shared on a dark web development forum, related chatter spiking on a Telegram channel, and early scanning activity reported targeting internet-facing CRM APIs. Signal confidence: medium and rising.

8:06 AM — Internal exposure mapped. Investigation Insights automatically correlates the external signal against the customer's environment. Two production CRM API servers are internet-facing and running the vulnerable library version. A WAF rule is in place but egress restrictions are missing. No analyst pivoting required — the correlation is automatic.

8:08 AM — Risk assessed and quantified. Continuous Control Monitoring with Risk Quantification applies tailoring factors: asset business value (high — revenue-impacting), internet exposure (confirmed), compensating controls (partial), exploit prerequisites (met). The system models expected financial loss across all affected assets.

8:09 AM — Prioritization delivered.

Aggregate risk exposure: $1.64M – $2.33M. Priority reflects exploit likelihood × asset value × control gaps — not severity scores. (Modeled figures, illustrative of risk quantification output.)

8:10 AM — Client-tailored alert delivered. The TI Platform structures the intelligence, triggers automated workflows, and delivers a prioritized alert with recommended actions: validate exploitability on the highest-risk asset, apply egress restrictions, monitor for follow-on identity abuse.

This alert did not exist at 8:03 AM. It exists at 8:10 AM because intelligence was detected, correlated, tailored, quantified, and delivered — automatically.

PRIORITY	ASSET	EXPOSURE CONTEXT	EST. FINANCIAL IMPACT
#1	crm-api-prod-01	Internet-facing, no egress control, exploit prereqs met	$1.2M – $1.6M
#2	crm-api-prod-02	Internet-facing, partial WAF mitigation	$400K – $650K
#3	crm-api-stage-01	Non-prod, limited exposure, lower asset value	$40K – $80K

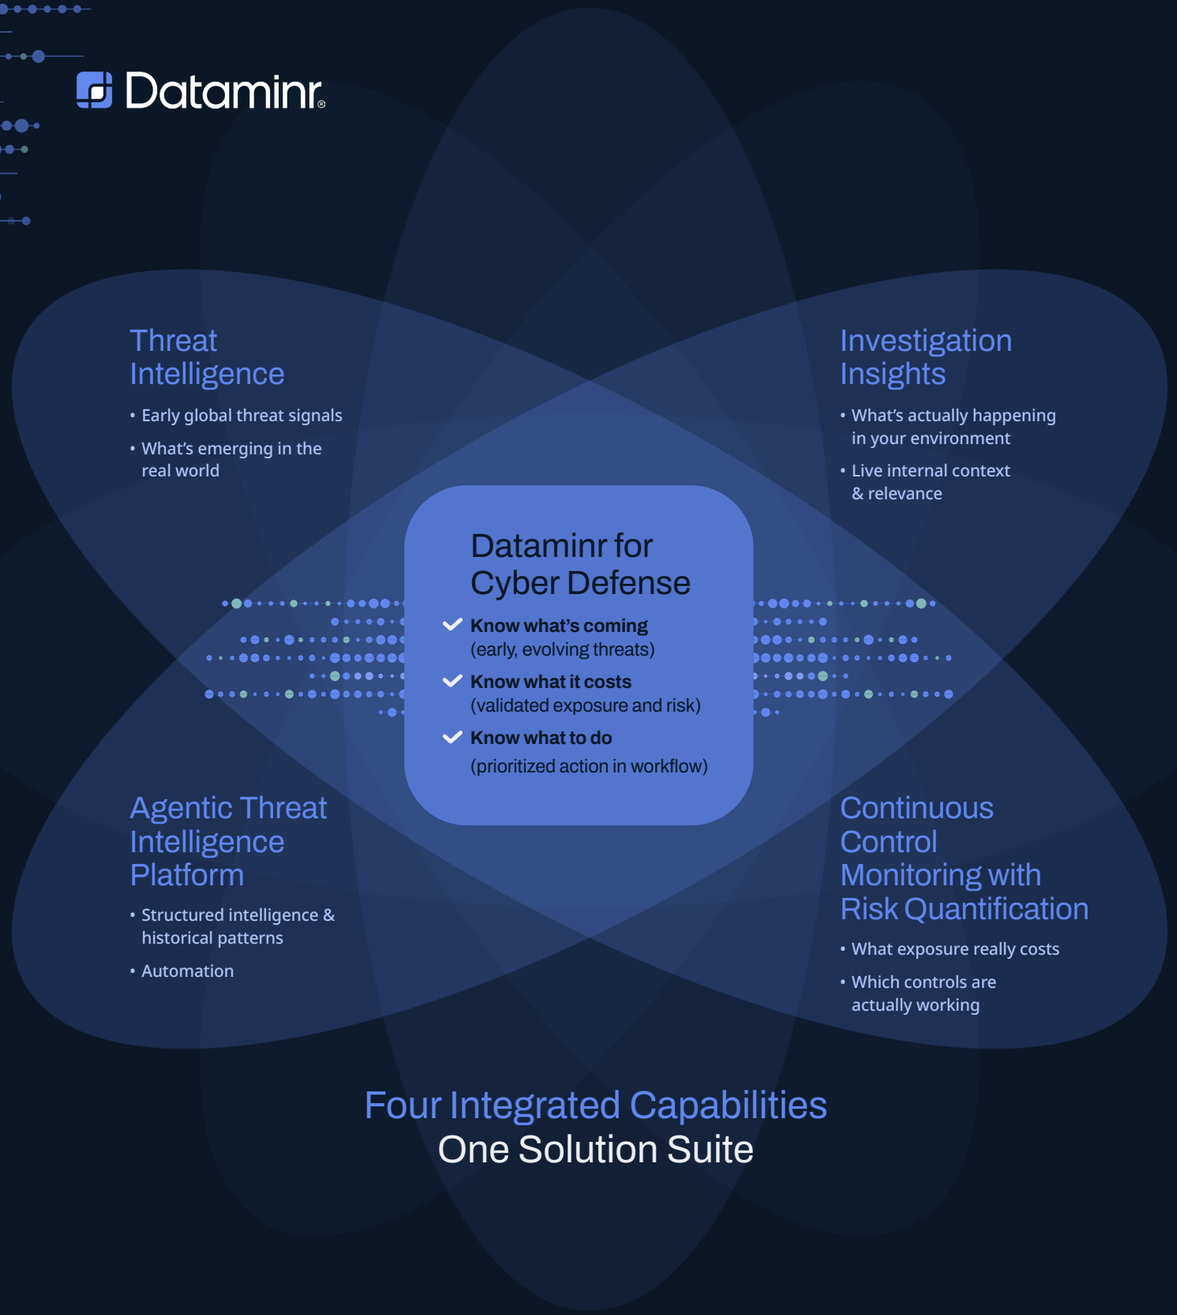

Agentic AI:
The Connective Tissue

The four capabilities describe what the system sees. The AI is how it thinks.

The architecture below shows how. Signals flow in from both sides — public sources including OSINT, deep and dark web, geopolitical, brand, and third-party intelligence on one side, and internal telemetry from your SIEM, EDR, XDR, and risk tools on the other. The agentic core fuses those inputs, contextualizes them against your assets and controls, models business impact, prioritizes by business consequence, and orchestrates the response. Operational actions flow out to the teams and workflows that need them — alert triage, hunt automation, detection engineering, incident response, vulnerability prioritization, and exposure management. And every analyst decision feeds back into the intelligence layer, sharpening the next cycle. The result isn't a pipeline with an endpoint. It's a closed loop.

Powering that loop: more than 55 proprietary LLMs and agentic AI solutions alongside thousands of detection models, all purpose-built for specific security tasks and working in parallel across more than 43 terabytes of event, threat, and risk signals ingested every day [Dataminr 2026 Cyber Threat Landscape Report]. These detections are enriched through a proprietary knowledge graph of over 3 billion entities — including unique indicators of compromise, threat actor profiles, vulnerability records, and asset relationships [Dataminr 2026 Cyber Threat Landscape Report].

Most vendors have responded to the AI moment by wrapping a single large language model around their existing data — a chatbot for your security stack. Here's a simple test for whether AI is a feature or an engine: remove the AI, and does the product still work? For most vendors, the answer is yes. The AI summarizes what the product already found. It's a feature. Remove Dataminr's AI, and there is no product. The AI detects, assembles, tailors, and predicts — autonomously and dynamically. It's the engine. Dataminr took a fundamentally different approach — and has been building it for longer than most vendors have been talking about AI.

How the System Thinks

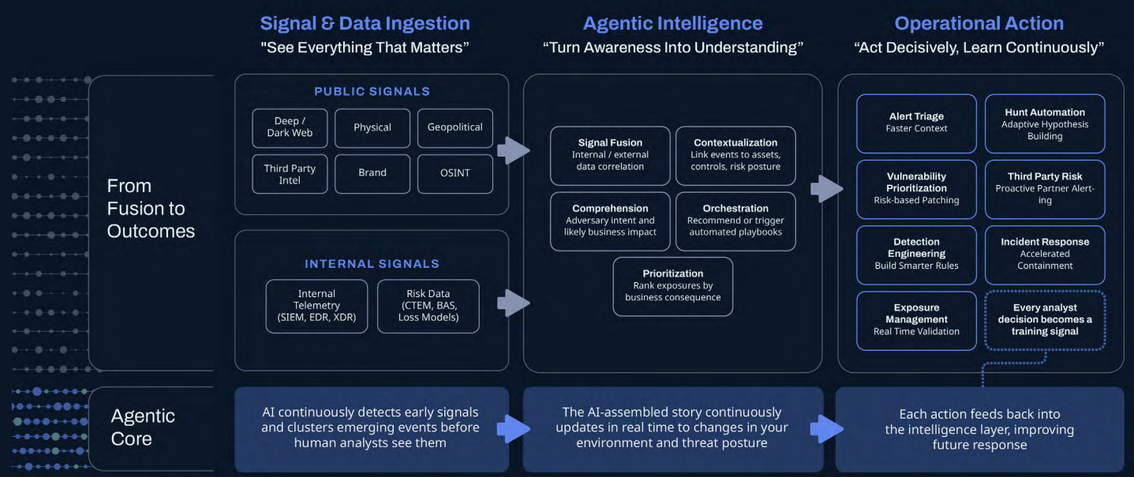

The Closed Loop

This architecture wasn’t built in response to the AI hype cycle. Dataminr deployed its first foundation model in production in 2020 — years before most security vendors started talking about AI. Generative AI has been in production for more than five years. This isn’t a pivot. It’s a 12+ year head start, built by more than 20 PhDs in NLP, computer vision, and knowledge graph construction. And the AI runs on a data moat no competitor can replicate: more than 12 years of detected, fully-labeled security events. More than 100 million events processed and structured by Dataminr’s own AI. Trillions of public data points from over 1 million sources. No other AI platform in cybersecurity has access to this archive — and AI is only as good as what it’s trained on.

Most AI-powered security products process signals sequentially — one model detects, then another enriches, then another contextualizes. Dataminr’s architecture works differently.

More than 100 specialized AI models operate in parallel: multi-modal fusion models process text, images, video, and audio simultaneously. Entity extraction models identify actors, infrastructure, and TTPs in real time. Client-tailoring models filter for relevance while correlation models build the complete picture. This parallel processing compresses what would take analysts hours into seconds — delivering both awareness and understanding at the same time, not one after the other.

Client-Tailored and Predictive by Design

What makes this architecture different from bolting a chatbot onto an existing security tool isn’t speed alone. It’s that the AI delivers two capabilities most vendors can’t credibly claim — and that together define a fundamentally different kind of intelligence.

Client-tailored intelligence answers a different question than traditional threat intelligence. Not “Is this threat real?” but “Does this threat matter to us?” Intel Agents automatically filter and prioritize every signal using what the system knows about the customer — exposed assets, technologies in use, observed alerts, historical activity, and control posture. The same exploit produces a different intelligence output for a healthcare system than for a financial services firm, because the exposure, the regulatory implications, and the business impact are different. This is what moves intelligence from a feed that teams have to interpret to a sightline they can act on.

Predictive intelligence provides forward-looking insight into how a situation is likely to evolve. Rather than freezing intelligence at the moment of first detection, Intel Agents assess patterns of change over time — acceleration of threat actor activity, maturation of exploit tooling, shifts in targeting, convergence with internal exposure. The system doesn’t just tell you what happened or what’s happening — it anticipates which vulnerabilities are most likely to be exploited in your environment, which campaigns are trending toward your industry, and where control degradation is creating new exposure before it’s tested. This is what separates early warning from early advantage.

Together, these capabilities create a system that compounds over time. Every analyst action — validating intelligence, closing an investigation, adjusting a detection — feeds back into the system. The AI learns what matters to your organization, what your team acts on, and how your environment changes. Intelligence gets sharper. Prioritization gets more accurate. The system doesn’t just run — it improves.

Intel Agents don’t produce intelligence once and move on. They maintain it. When new signals emerge — a scanning spike, a forum post claiming exploitation, a related CVE disclosure — the agents reassess what they’ve already assembled: updating confidence levels, adjusting relevance scores, and re-evaluating whether the threat applies to your environment given what’s changed. Intelligence stays current rather than decaying from the moment it’s created.

The AI doesn’t replace analysts. It relocates them — from manually collecting and processing intelligence to evaluating, deciding, and directing action. Every workflow the system automates is an hour returned to the work that actually requires human expertise.

Growing Into the System

Organizations don't have to adopt everything at once. Dataminr for Cyber Defense meets teams where they are and scales as their program matures. Three solutions provide natural entry points based on where an organization's needs are greatest:

Client-Tailored Threat Intelligence — for teams that need earlier, more relevant intelligence delivered directly into their existing workflows. Combines Threat Intelligence and Investigation Insights so analysts see emerging threats and know immediately whether they apply.

Agentic TI Ops — for teams ready to produce and operationalize intelligence at scale. Adds the Agentic Threat Intelligence Platform so that intelligence becomes structured, reusable, and automatically distributed to detection, response, and reporting functions.

Predictive Threat Exposure Management — for risk teams that need continuous control validation and financial risk quantification. This solution stands on its own — a risk or vulnerability management team can adopt it independently to replace static assessment cycles with live, financially grounded prioritization. It doesn't require the other solutions to deliver value on day one. But it also serves as the final piece of the full system. When combined with Client-Tailored TI and Agentic TI Ops, Predictive TEM completes the chain — connecting the earliest external threat signals to internal posture assessment to dollar-denominated business impact in a single, continuous loop. That's when the system delivers what the previous chapters described: foresight, focus, and action working together without a manual handoff in between.

Each solution delivers standalone value. Combined, they deliver the full system — and the AI gets smarter with every capability added, because each one contributes data the others can't see. An organization can start with threat intelligence and grow into operationalization. A risk team can start with Predictive TEM and gain immediate financial clarity. Over time, the paths converge into one connected system.

CHAPTER 5

The Mended Chain in Action

From First Signal to Decisive Action

Everything in this book — the broken chain, the three gaps, the frameworks, the system — comes down to one question: does it actually change outcomes when a real threat is moving?

In October 2025, it did.

On October 7th at 12:50 PM ET, Dataminr's AI detected an unknown Fortinet exploit — possibly a variant of CVE-2022-40684 — originating from a specific IP address, with a JWT payload containing admin credentials and attempting to access a Fortinet device. This wasn't a published CVE. There was no vendor advisory. No entry in the Known Exploited Vulnerabilities catalog. The signal came from where threat actors actually operate, detected by AI models monitoring exploit activity across the deep and dark web in real-time.[31]

That was Day 1. Here's what happened over the next 38 days — and what the rest of the industry missed.

Foresight: The signal is detected.

Dataminr's Threat Intelligence capability picks up the exploit activity and begins correlating it with related signals — scanning patterns, forum chatter, infrastructure shifts. Within hours, customers receive tailored intelligence: not a generic "Fortinet vulnerability exists" advisory, but specific context about the attack vector, the targeted configurations, and the indicators to hunt for immediately.

Focus: Exposure is mapped and prioritized.

Here's where a connected system changes the calculus. Once the signal arrives, Investigation Insights overlays it against the customer's environment — identifying which Fortinet appliances are deployed, which are internet-facing, and which have the vulnerable configuration. No swivel chair. No manual CMDB lookup. The TI Platform structures the intelligence into reusable records and automatically routes it to detection engineering and incident response teams so they can begin hunting immediately.

For organizations running Continuous Control Monitoring with Risk Quantification, the system goes further: modeling the financial exposure of the specific assets at risk, factoring in compensating controls and business criticality. Instead of "we have Fortinet appliances and there's a new exploit," leadership gets "here are the three systems that carry material financial risk if this exploit lands, and here's the dollar exposure."

Action: Defenses are hardened — weeks before anyone else knows.

Armed with specific, prioritized intelligence, security teams investigate for early signs of intrusion, harden defenses against vulnerable systems, and adjust detection rules — all while the exploit remains unknown to the broader security community.

November 13 — 37 days later. watchTowr Labs independently publishes their analysis of the actively exploited Fortinet FortiWeb vulnerability. The next day, they release a detection script.

November 14 — 38 days later. US CISA formally adds CVE-2025-64446 to the Known Exploited Vulnerabilities catalog. Fortinet issues a patch.

[31] Dataminr, "2026 Cyber Threat Landscape Report"

For most organizations, this is day one — the beginning of triage, scoping, and prioritization. For Dataminr customers, it was day 38 — and defenses were already in place.[32]

Put a dollar figure on that gap. Based on proprietary Dataminr cyberloss data, every 10 days a breach goes undetected increases expected financial loss by $2,209,100.[33] At 38 days of lead time, that translates to more than $8.4 million in potential risk exposure avoided (illustrative, based on applying the $220,910/day figure to the documented lead time). That's the financial value of a connected system versus a disconnected one.

This wasn't an isolated case. In the PowerSchool breach — one of the largest education-sector incidents of 2025 — Dataminr alerted on compromised PowerSchool credentials more than six months before the breach occurred.[34] Six months of lead time. The signal was there. The question was whether it reached someone who could act on it, in a system that connected the signal to the exposure to the business impact.

That's what a mended chain delivers. Not just earlier detection — earlier understanding, faster prioritization, and decisive action, with the full context of what's at stake.

FORTINET FORTIWEB | CVE-2025-64446

38 Days Earlier

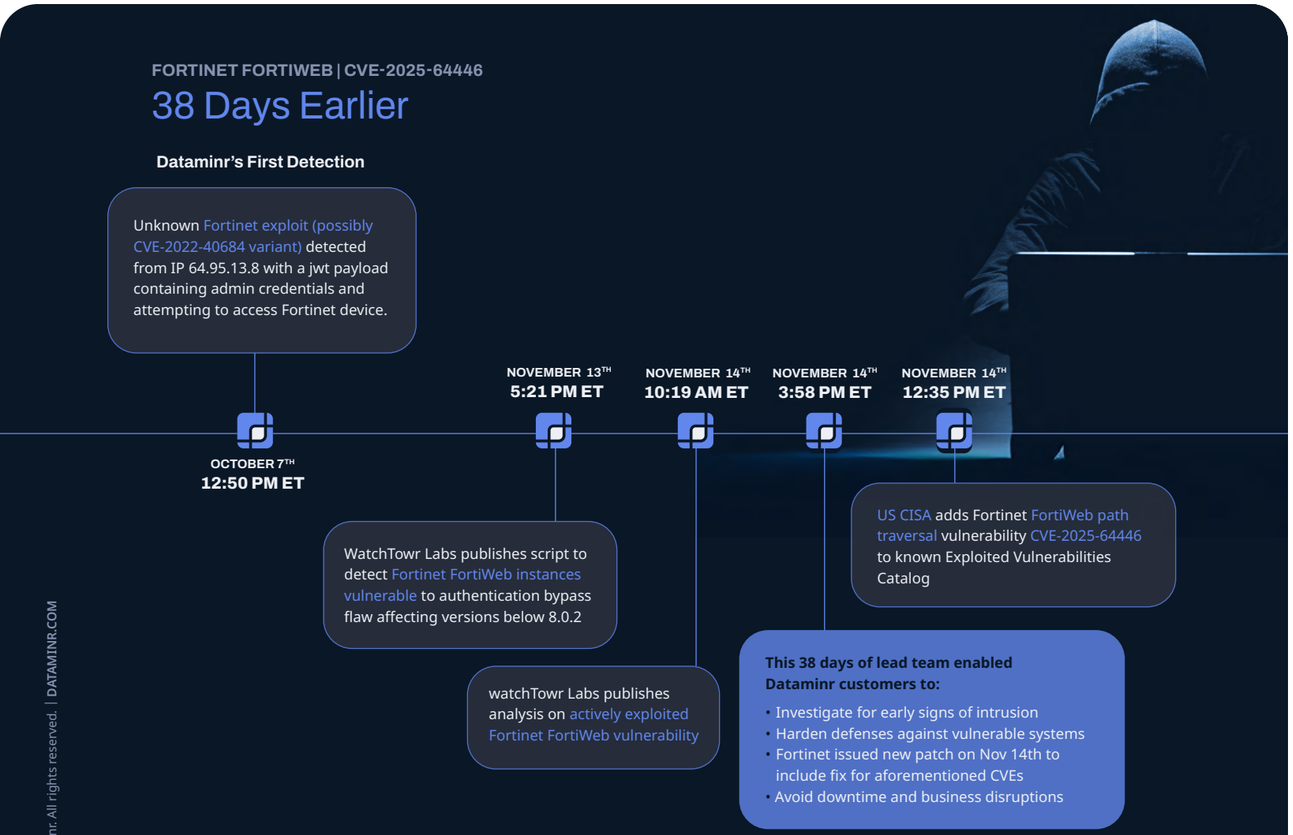

[32] Dataminr Cyber Intel Brief, "CVE-2025-64446: Fortinet FortiWeb Zero-Day," November 18, 2025
[33] Dataminr Continuous Control Monitoring with Risk Quantification proprietary cyber loss data
[34] Dataminr, "2026 Cyber Threat Landscape Report"

Proof Points

Real-world outcomes from organizations using the system today.

7 hours → 37 minutes
Incident response time, still decreasing
Forbes 2000 Healthcare System

2-5 min → 2 seconds
Time per investigation task, after automation
Fortune 500 Retailer

200M → 12
SIEM events/month narrowed to 12 actionable incidents
Fortune 300 Financial Institution

$1.3M saved/year
60+ workflow automations replacing manual effort
Fortune 200 Healthcare Org

74% reported ≥25% false positive reduction
42% reported reductions of 50%+
2025 Dataminr Customer Survey

80% reported mending of threat, asset, and business data
As a Dataminr differentiator
2025 Dataminr Customer Survey

Afterword

I co-founded ThreatConnect. I co-authored the Diamond Model of Intrusion Analysis. I've spent the better part of 15 years building threat intelligence tools — the very category this book argues isn't enough on its own. So let me say something that might surprise you coming from me: the book is right.

When we started ThreatConnect, analysts had no shared system to collect, structure, and act on intelligence. We built one. The Diamond Model gave them a framework for understanding intrusions. ThreatConnect gave them a platform to operationalize it. For a long time, that was the most urgent thing to solve.

But the threat landscape didn't wait for us to finish. Somewhere along the way, the question shifted. It went from "how do we operationalize intelligence?" to "how do we connect intelligence to everything else — posture, exposure, financial impact, business decisions — fast enough to actually matter?" That's a fundamentally different problem. And it's the one Dan lays out in these pages.

I'll be honest: as a builder, it's tempting to believe that better features solve structural gaps. They don't. The gap between a threat intel platform and a vulnerability scanner isn't a feature gap. It's an architectural one. Closing it required rethinking what the system even is — not another tool in the stack, but the connective layer across all of them.

That's what Dataminr for Cyber Defense is becoming.

— Andy Pendergast

Andy Pendergast is EVP of Product at Dataminr and co-founder of ThreatConnect. He co-authored the Diamond Model of Intrusion Analysis and has spent 15 years building threat intelligence platforms from inside the U.S. Army, the DoD, and the private sector.

The Path Forward

For years, the cyber defense industry has operated on a simple assumption: if you buy enough good tools, the defense will hold. But the evidence says otherwise. The tools are good. The chain between them is broken.

This book laid out the problem — three links that should hold together but don't, because threat intelligence, security posture, and business impact were built, bought, and managed in isolation. It named the structural cause — a point-tool era that optimized each link without ever connecting them. And it showed what changes when the chain is mended: foresight that compresses hours to seconds, focus that replaces severity scores with financial clarity, and action grounded in your actual exposure rather than headlines.

The industry frameworks are here. UCRI and CTEM describe the architecture. The technology is here. Agentic AI makes it operational at a scale and speed that manual processes never could.

And the evidence is here. Organizations running Dataminr for Cyber Defense are detecting threats 38 days before formal disclosure. They're reducing incident response times from hours to minutes. They're replacing millions of raw alerts with a handful of actionable, prioritized incidents — and quantifying the risk in language their leadership can act on.

In our 2025 customer survey, 79% of respondents rated the platform as critical or very important to their operations. That's not adoption. That's reliance. It's the signal that intelligence has moved from a tool teams use to an operating system teams depend on.

The question isn't whether the chain needs to be mended. It's how fast your organization can get there.

Quick Reference

Dataminr for Cyber Defense: Four Capabilities, One System

Dataminr for Cyber Defense is an intel-driven threat and exposure management solution suite that preemptively assembles and operationalizes real-time, client-tailored intelligence from first signal to risk-prioritized action. Four capabilities feed a shared intelligence layer powered by agentic AI that is client-tailored to each customer's environment and predictive of how threats will evolve — reconnecting the chain between what's happening externally, what's true inside your environment, and what it means to your business.

Four Capabilities

CAPABILITY	WHAT IT DOES	THE QUESTION IT ANSWERS
Threat Intelligence	Detects emerging threats from 1M+ sources — text, image, video, audio, sensor data — across the deep, dark, and open web. Often before formal disclosure.	**"What's coming?"**
Investigation Insights	Overlays cross-tool context directly in analyst workflows. No swivel chair. No tab-switching. Computer vision extracts indicators from whatever's on screen.	**"Does this apply to us?"**
Agentic Threat Intelligence Platform	Structures intelligence into reusable records. Automates enrichment, distribution to SIEMs and EDRs, detection engineering, and reporting. Intelligence compounds — it doesn't get rebuilt from scratch.	**"What have we seen before, and what do we do next?"**
Continuous Control Monitoring with Risk Quantification	Continuously validates control effectiveness and translates exposure into dollar-denominated financial risk — grounded in your actual assets, controls, and 20+ years of historical cyber loss data.	**"What does this cost us, and are our controls working?"**

Key Use Cases

Alert Triage
Cut noise and focus analyst time on what matters

Detection Engineering
Strengthen coverage and reduce false positives

Vulnerability Prioritization
Fix what's costliest, not what's loudest

Control Assessment
Know instantly if defenses are performing as intended

Hunt Automation
Find evolving threats sooner with AI-driven context

Incident Response
Shorten MTTR, stop threats before impact spreads

Third Party Risk
See supplier & partner exposures before they affect you

Exposure Management
See and prioritize exposures in real time - before they become incidents

Three Solutions — Two Entry Points

1 **Starting from threat intelligence:**
Client-Tailored TI ⟶ Earlier, more relevant intelligence inside existing workflows.
Threat Intelligence + Investigation Insights

Agentic TI Ops ⟶ Structured, automated, scalable intelligence operations.
Adds Agentic Threat Intelligence Platform

2 **Starting from risk:**
Predictive Threat Exposure Management ⟶ Continuous control validation and financial risk quantification — standalone or as the capstone of the full system.

What the Mended Chain Delivers

Foresight — Know what's coming, early enough to change the outcome.

Focus — Know what it costs, grounded in evidence.

Action — Know what to do, without rebuilding context from scratch.

By the Numbers: Real Customer Outcomes

METRIC	RESULT	SOURCE
Detection lead time	**38** days before formal CVE disclosure	Dataminr 2026 Cyber Threat Landscape Report
Early warning	**6+** months before the PowerSchool breach	Dataminr 2026 Cyber Threat Landscape Report
Incident response	**7** hours ⟶ **37** minutes	Forbes 2000 Healthcare System
Investigation speed	**2–5** minutes⟶ **2** seconds per task	Fortune 500 Retailer
Alert reduction	**200M** SIEM events ⟶ **12** actionable incidents/month	Fortune 300 Financial Institution
Cost savings	**$1.3M**/year in automated workflow savings	Fortune 200 Healthcare Organization
False positive reduction	**74%** reported ≥**25%** reduction; **42%** reported ≥**50%**	2025 Dataminr Customer Survey
Tool effectiveness	**70%** of respondents confirmed improved SIEM/ EDR/XDR/SOAR performance	2025 Dataminr Customer Survey

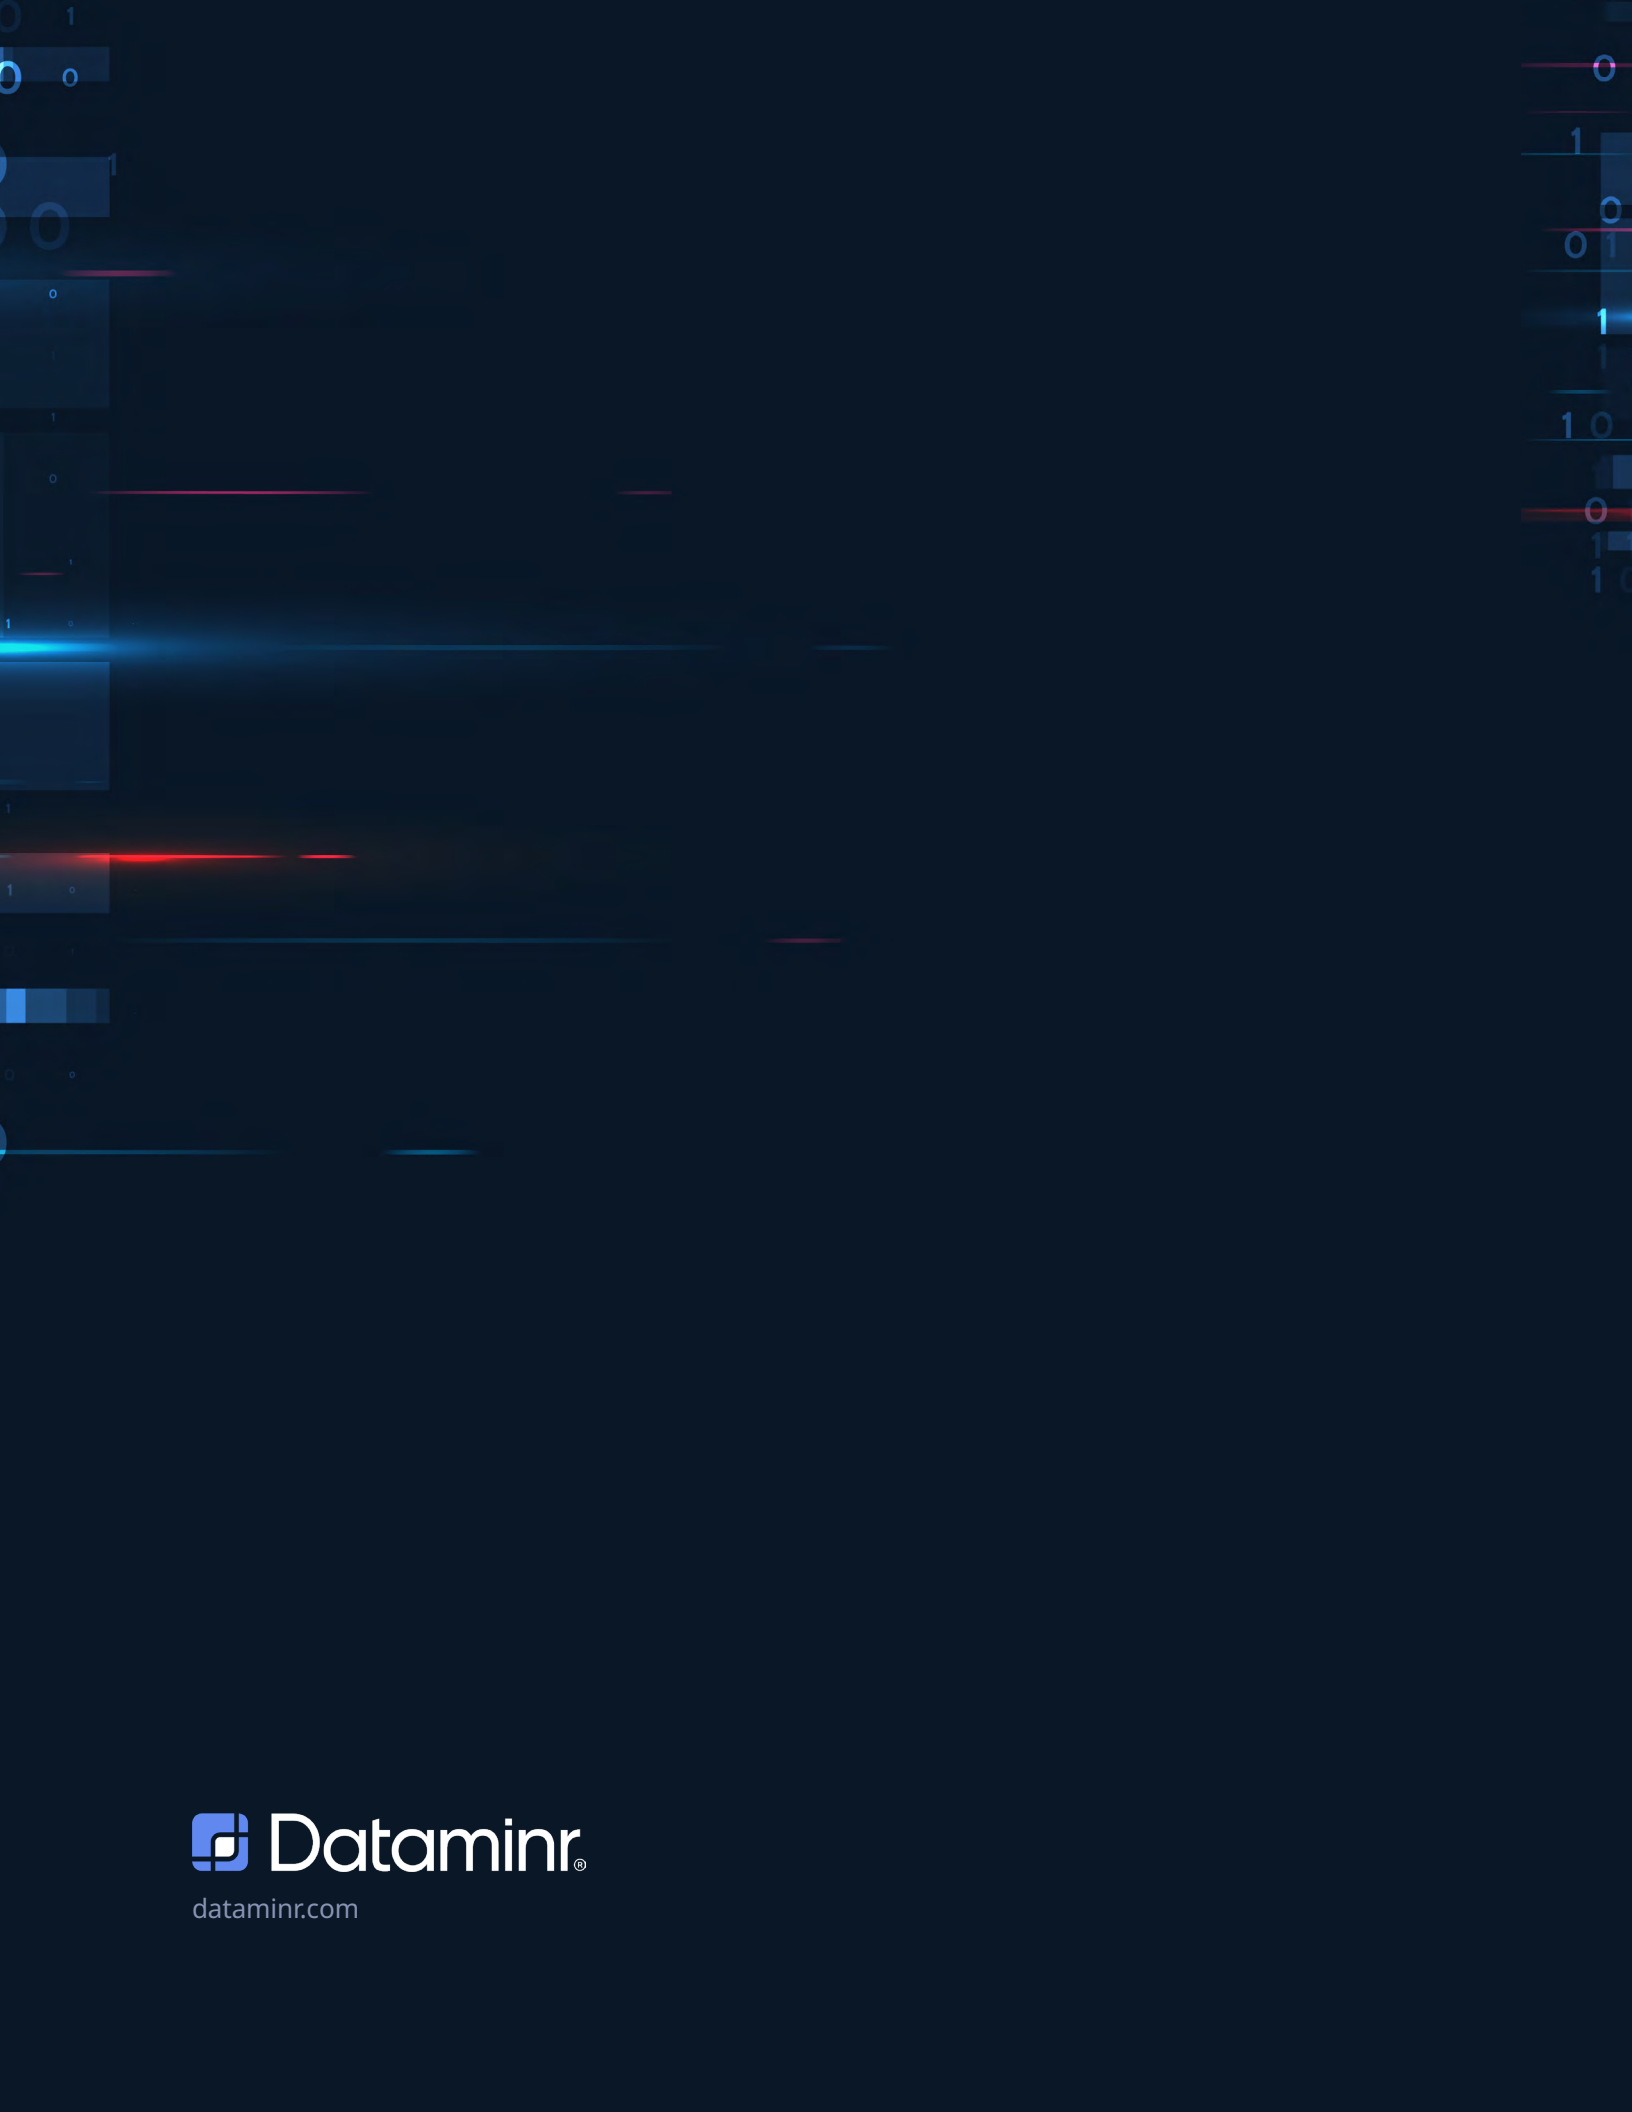

Dataminr®

dataminr.com