databricks

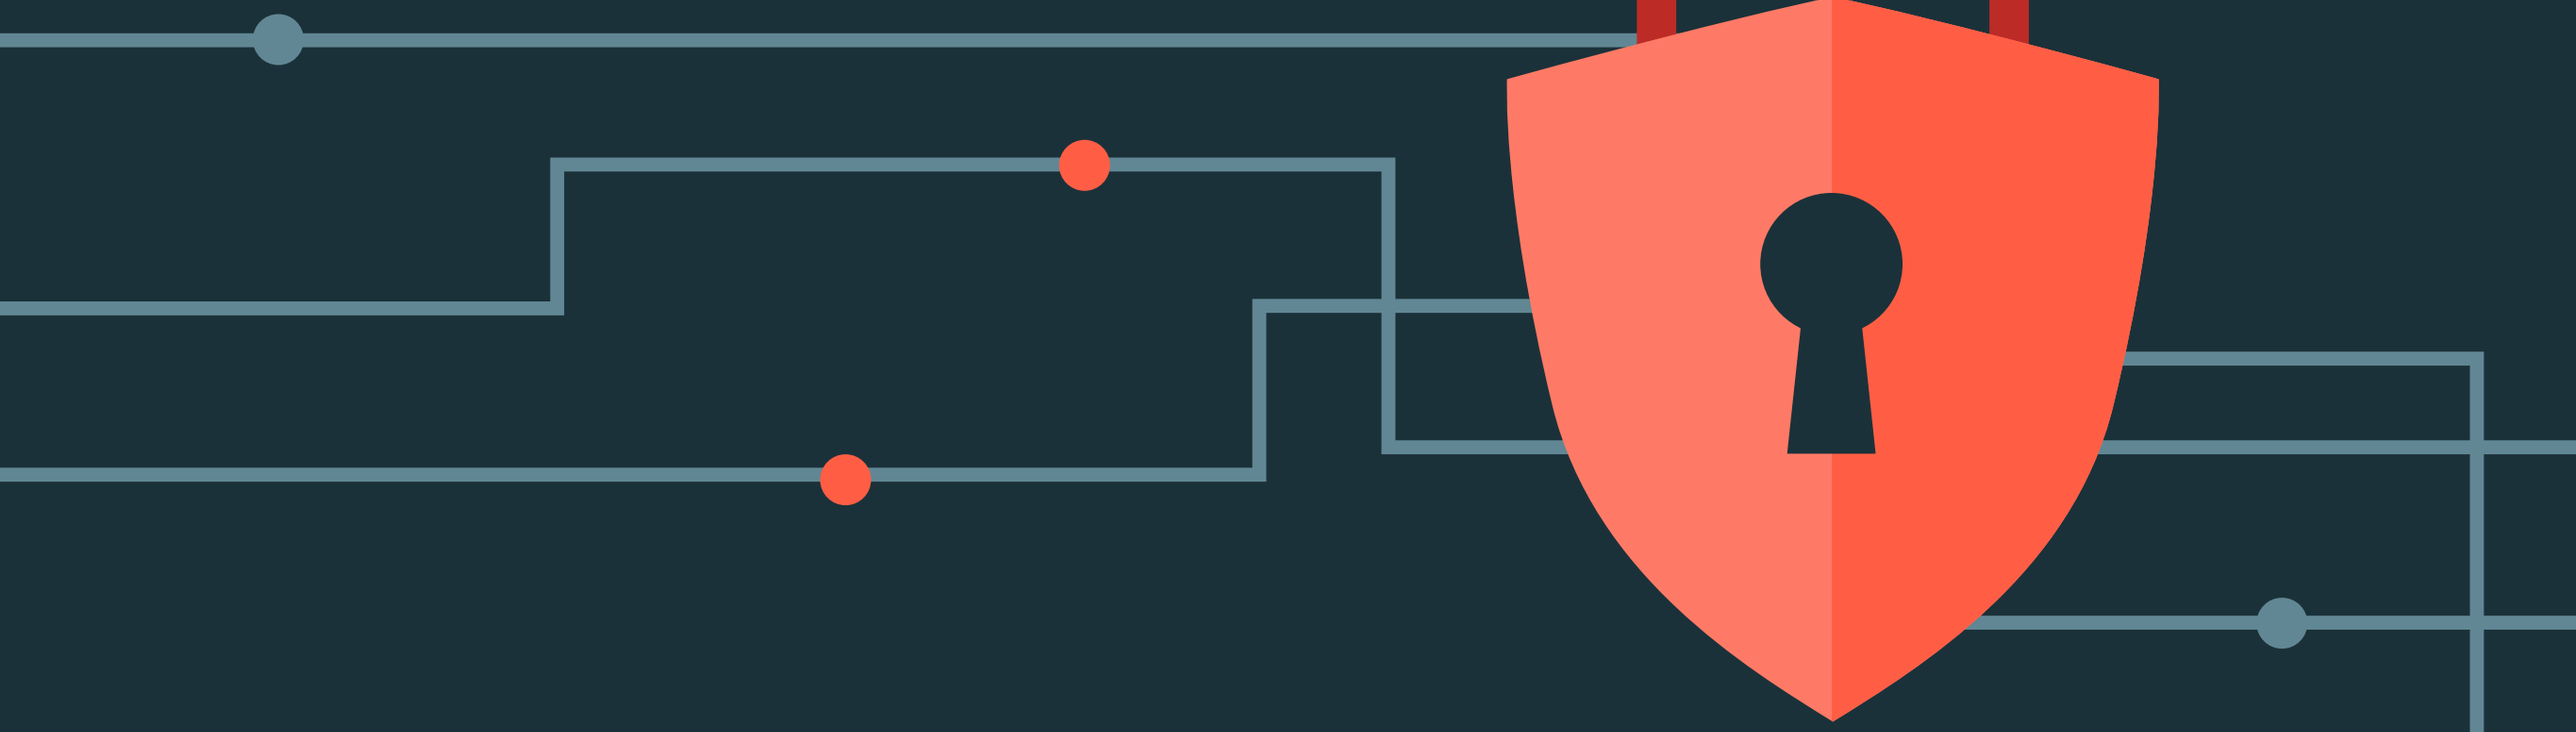

eBook

How Security Leaders Are Shaping the Future of AI-Driven Defense With Databricks

Contents

Introduction: A New Era in Cybersecurity

Cybersecurity leaders today face a perfect storm. As organizations digitize at a record pace, their attack surface explodes and cyberthreats become more advanced, amplified by adversarial AI, nation-state actors and the rapid adoption of multicloud and software as a service (SaaS). Security operations centers (SOCs) drown in fragmented telemetry, unmanageable alert volume and a persistent talent crisis. Despite heavy investment in tools, most organizations remain hamstrung by vendor and data silos, sky-high security information and event management (SIEM) costs and slow, reactive processes.

The scale of the challenge is staggering:

- **Cybercrime costs:** Expected to hit $10.5–12 trillion per year by 2025 (Barracuda, Cybersecurity Ventures, World Economic Forum / Accenture)
- **Average data breach cost:** Now $4.5M globally and $10.2M in the U.S. (IBM Cost of a Data Breach)
- **Time to identify and contain a breach:** 204–277 days on average (IBM Cost of a Data Breach)
- **Security tool sprawl:** Enterprises now run 50–83 security tools from 29+ vendors (Cybersecurity Dive, VentureBeat / Forrester)
- **Breach prevalence:** Up to 67% of large enterprises have had a breach in the past two years (WEF), with 70% of siloed-data organizations suffering a breach (BlinkOps, Infoverity)
- **Skills shortage:** Over 3.5 million unfilled cyber jobs globally, compounding all these challenges (ISC2 Cybersecurity Workforce Study)

Cybersecurity is now a data and AI problem. To outpace attackers, security teams must unify, govern and harness mountains of enterprise data — transforming it into actionable insight, not just more alerts. Databricks is enabling this transformation for thousands of top global brands with Data Intelligence and AI.

Data Intelligence for Cybersecurity empowers your team with advanced AI, custom detections and seamless integrations, enabling rapid, predictive threat detection and response. Databricks modernizes your SOC, reduces operational overhead and helps you avoid vendor lock-in, delivering timely, actionable insights across your entire data ecosystem.

In this eBook, we'll explore how modern organizations are supercharging cybersecurity efforts with data intelligence and agentic AI. Learn how, with Databricks, your data becomes your strongest defense.

Cybersecurity's Biggest Challenges: Why SOC Teams Struggle

Security executives and their teams are united by common pain points. Every day, SOCs contend with a landscape that is more volatile, complex and fast-moving than ever before. The explosion of threats — ranging from ransomware to supply chain attacks — means defenders are on constant alert, yet are too often hindered by the complexity of their own environments. Despite significant investments in security tools, teams are frequently blocked by barriers that limit visibility, slow response times and stifle innovation. Legacy architectures, tool sprawl and a persistent shortage of skilled talent only aggravate their struggle to stay proactive. As a result, many organizations find themselves reacting to crises instead of operating with confidence and clarity. The following challenges illustrate the root causes behind these struggles:

- **Fragmented data silos:** 72% of security executives say their security and IT data is siloed, slowing response (Ivanti State of Cybersecurity Report, WEF Cyber Report)
- **Escalating SIEM costs:** Legacy "store everything" SIEM architectures are unsustainable, especially as average breach costs surge and the global financial toll of cybercrime continues to grow
- **Alert fatigue and manual investigations:** SOC analysts chase thousands of low-fidelity alerts, with 62% reporting that data silos specifically slow their ability to respond to threats (Ivanti)
- **Vendor lock-in and limited flexibility:** With an average of 75 security solutions in play, tools often lack integration, hindering the ability to see and act on threats holistically (The Hacker News)
- **AI and automation readiness:** Legacy platforms can't handle the data or compute required to bring modern, agentic analytics to production
- **Compliance and governance pressure:** Growing regulatory scrutiny and class-action settlements (outpacing regulatory fines by 50% in 2025) make auditability and data lineage nonnegotiable (Polymer)

> "SOC teams remain one step behind not because they lack tools, but because their data is locked away, their processes are slow and their platforms weren't built for the reality of today's threats."
>
> — Dave Herrald, Head of Cyber GTM, Databricks

The North Star of Cybersecurity

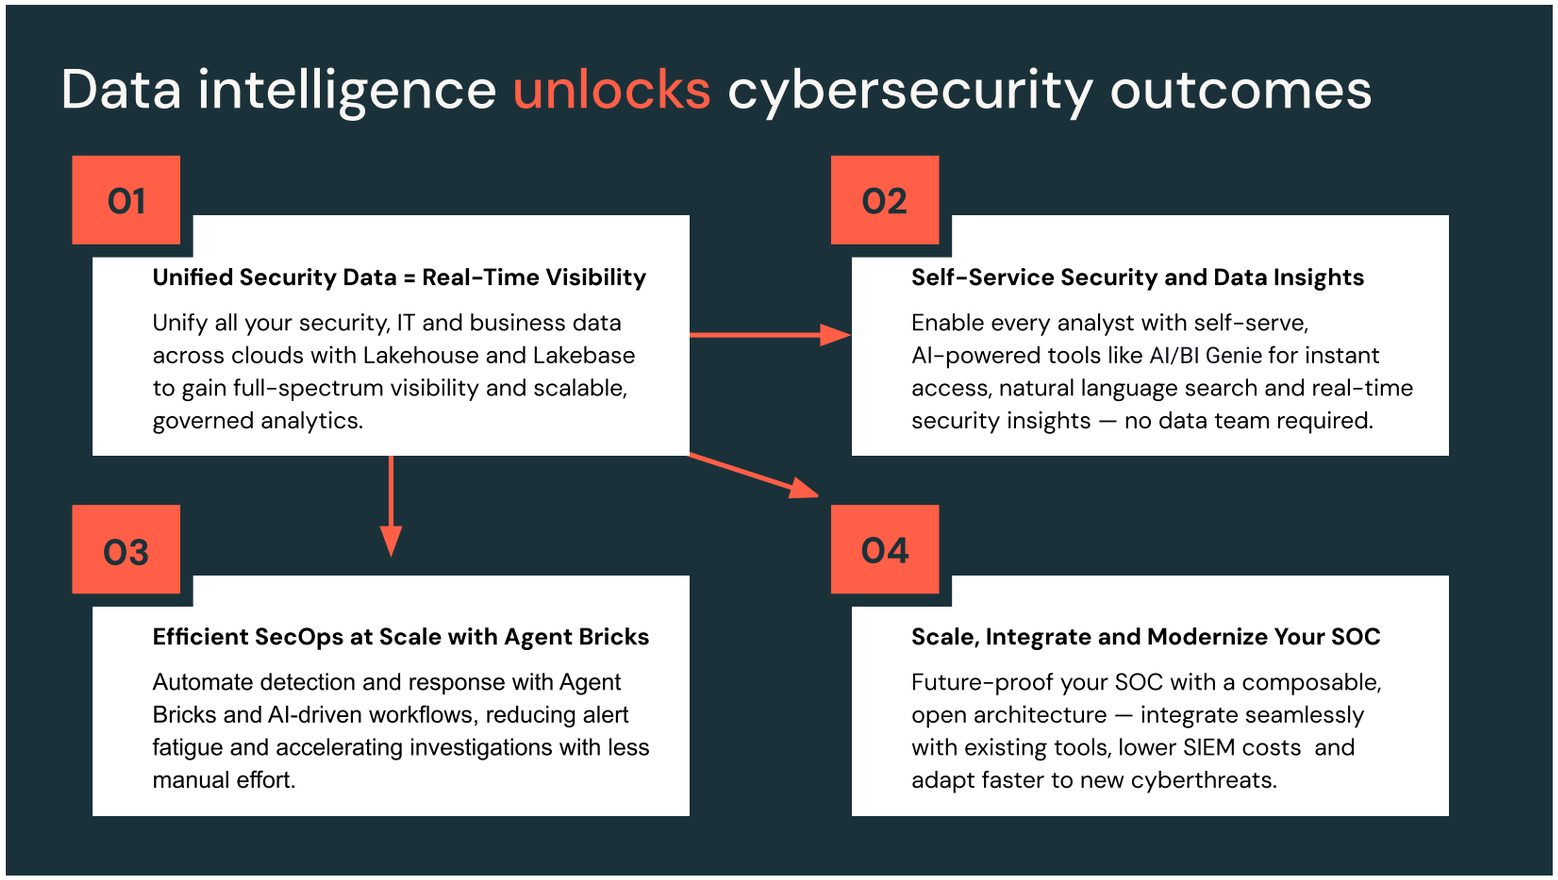

Databricks stands apart by treating security first and foremost as a data and AI challenge. The Databricks Data Intelligence Platform is an open, unified data lakehouse specifically designed to:

- **Empower AI-driven detection:** Native support for advanced machine learning (ML), automation workflows and natural language queries — in seconds, at petabyte scale
- **Unify all security, IT and business data:** Structured, semi-structured and unstructured data from any tool, cloud or SaaS platform, now in one governed platform
- **Modernize without rip and replace:** Integrate with your existing SIEM, SOAR (security orchestration, automation and response) and security stack partners to immediately optimize costs and operations while future-proofing your investments
- **Enable end-to-end governance:** Full auditability with Unity Catalog, supporting the most stringent compliance needs and AI security frameworks
- **Avoid vendor lock-in:** Open standards (Delta Lake, OCSF) and multicloud flexibility give you back control. Make your security architecture as agile as your business.

Databricks customers are upgrading their cyber defense — moving from reactive alert-chasing to intelligent, proactive, AI-powered security operations (SecOps).

Data Intelligence and the Three Pillars of Modern Security

Security threats, operational demands and business risks increase exponentially every year, so chief information security officers (CISOs) and security leaders need a data platform built for the realities of AI, cloud and always-on digital transformation. Databricks structures their Data Intelligence for Cybersecurity solution around three foundational pillars:

- More efficient SecOps at scale with Agent Bricks
- Self-service security and data insights
- Unified security data foundation

These pillars address the deepest issues facing modern security teams: fragmented telemetry, operational silos, alert fatigue and the need to harness AI for defense — not just analytics. When orchestrated together, they deliver end-to-end situational awareness, real-time agility and transformational gains in efficiency.

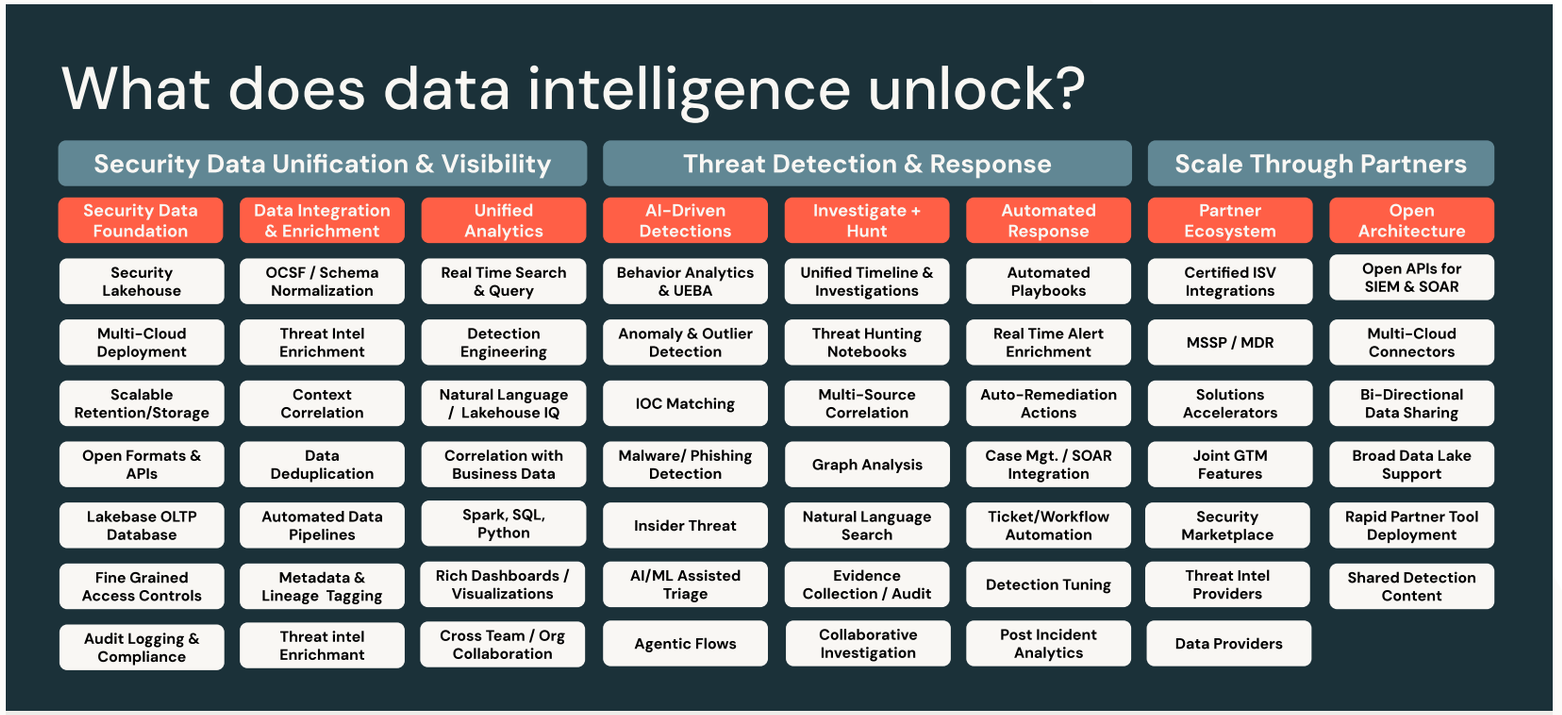

Powered by the Databricks Data Intelligence Platform

Data Unification	Data Governance	Data Engineering	Data Warehousing	DS/ML/LLM	Security Analytics
Delta / Iceberg/ Parquet, Lakebase OCSF	Unity Catalog, DASF 2.0, DAGF Framework	Lakeflow Connect, Declarative Pipelines	DBSQL, Serverless, Spark, Python, Notebooks	Lake Base, Agent Bricks MLflow	Solutions Accelerators, Brickbuilders, Connectors

Data intelligence, from the Databricks perspective, is the ability to unify, understand and act on all your enterprise data — no matter the type, source or scale — in real time. Rather than treating security data as an isolated or static resource, Databricks views it as the critical foundation for truly modern, AI-powered defense. By harnessing a unified lakehouse architecture and open standards, data intelligence enables security teams to break down silos, eliminate blind spots and bring together security, IT and business data on a single governed platform.

For cybersecurity, this unlocks a new era of visibility and control, enabling teams to detect threats faster, automate investigations and respond proactively instead of reactively. Analysts gain self-service access to context-rich insights, custom analytics and agentic AI, which allows them to focus on high-impact threats rather than repetitive manual work. Ultimately, data intelligence empowers organizations to transform their security operations — scaling detection, accelerating response, reducing costs and future-proofing their cyber resilience in a rapidly evolving threat landscape.

More efficient SecOps at scale with Agent Bricks

Modern threats move at machine speed, and adversaries leverage automation and AI to bypass traditional controls. To keep pace, security operations need more tools, but they also need integrated intelligence and agentic automation. Databricks uniquely delivers these capabilities at scale with Agent Bricks — reducing costs, accelerating response and amplifying human expertise.

- **AI agents for triage and detection:** Agent Bricks enables secure, production-ready AI agents for cyber defense — automating alert triage, investigation and response so SOC teams can act faster, reduce noise and adapt to emerging threats in real time — freeing analysts for higher-impact work
- **Automated contextual enrichment:** Seamlessly integrate with SOAR, SIEM, EDR and other platforms, orchestrate cross-channel response actions, and summarize incidents with precision and consistency
- **Operational resilience at scale:** Cut SIEM costs by up to 80%, reduce mean time to detect/respond (MTTD/MTTR) by up to 90% and eliminate redundancy by automating repetitive tasks across the SOC
- **Continuous optimization and governance:** Unity Catalog, Databricks AI Security Framework (DASF) and adaptive agent frameworks ensure all automation is secure, audit-ready and continuously refined for quality and compliance

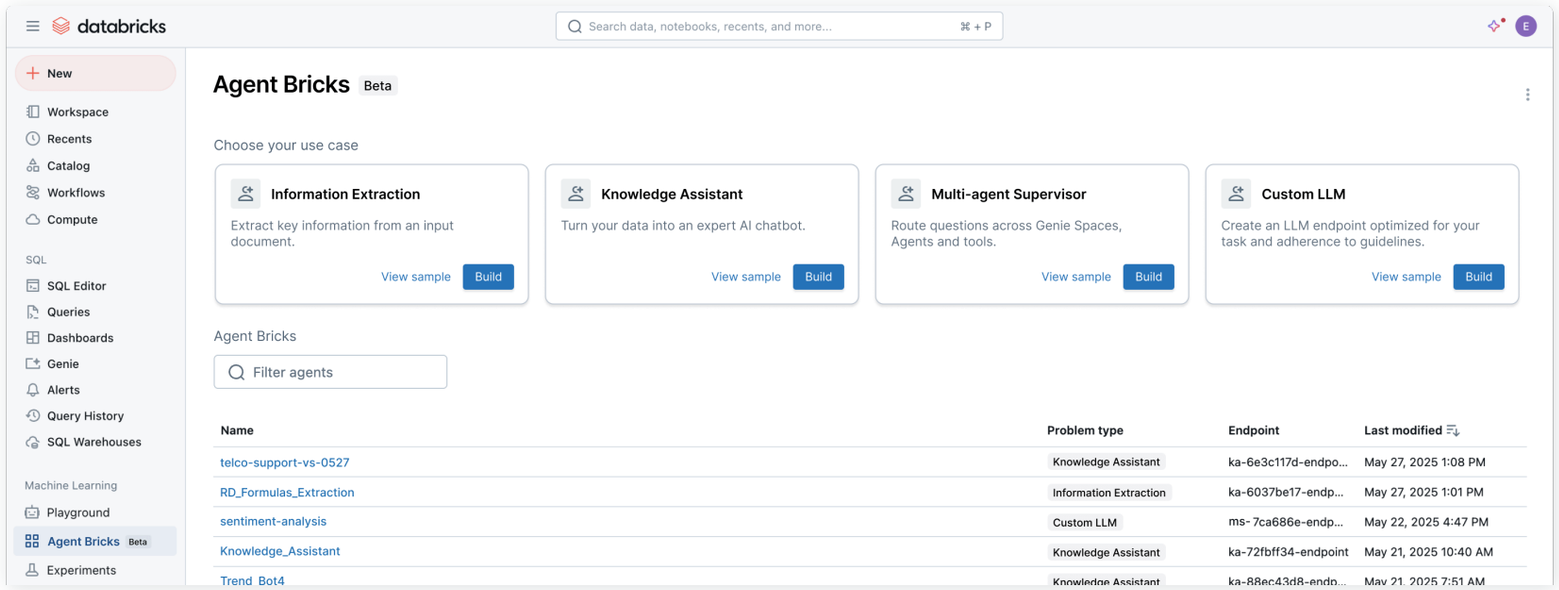

Agent Bricks dashboard

> "Databricks enables us to unify massive volumes of security data and build AI directly into our workflows—empowering our experts to deliver faster, more reliable threat detection for the thousands of organizations we protect."
>
> — Justin Lai, Distinguished Data Architect, Arctic Wolf

With Databricks, organizations modernize operations, enabling security teams to work smarter — not just harder — no matter how adversaries, data or regulations evolve.

Unifying these three pillars, Databricks provides the secure, AI-powered engine that modern SOCs and CISOs need to unlock true cyber resilience and future-proof their operations.

Self-service security and data insights

True cyber resilience hinges on empowering every analyst, detection engineer and security leader to access and act on data — without waiting in line for data engineering support. In an era of growing attack volume and skilled-labor shortages, democratizing analytics is the difference between proactive defense and reactive firefighting.

- **Databricks One and AI/BI Genie bring real-time, AI-powered security insights to everyone:** Both security and business users can access unified dashboards, monitor risk and ask natural language questions — all in a simple, code-free interface governed for compliance and privacy
- **Self-service analytics:** Analysts use SQL, Python, natural language and visual tools to access, harmonize and investigate unified data on demand — no more gatekeeping or delayed investigations
- **Federated access, instant enrichment:** Enable lightning-fast pivots, root cause analysis and threat hunting with context-rich linking of user, asset and business data at any scale
- **Actionable dashboards and collaboration:** Build and share real-time dashboards, notebooks and detections — fueling faster decision-making from the SOC to the C-suite
- **No bottlenecks, only agility:** Security professionals roll out, test and tune detections and analytics directly — removing operational bottlenecks and improving overall team velocity

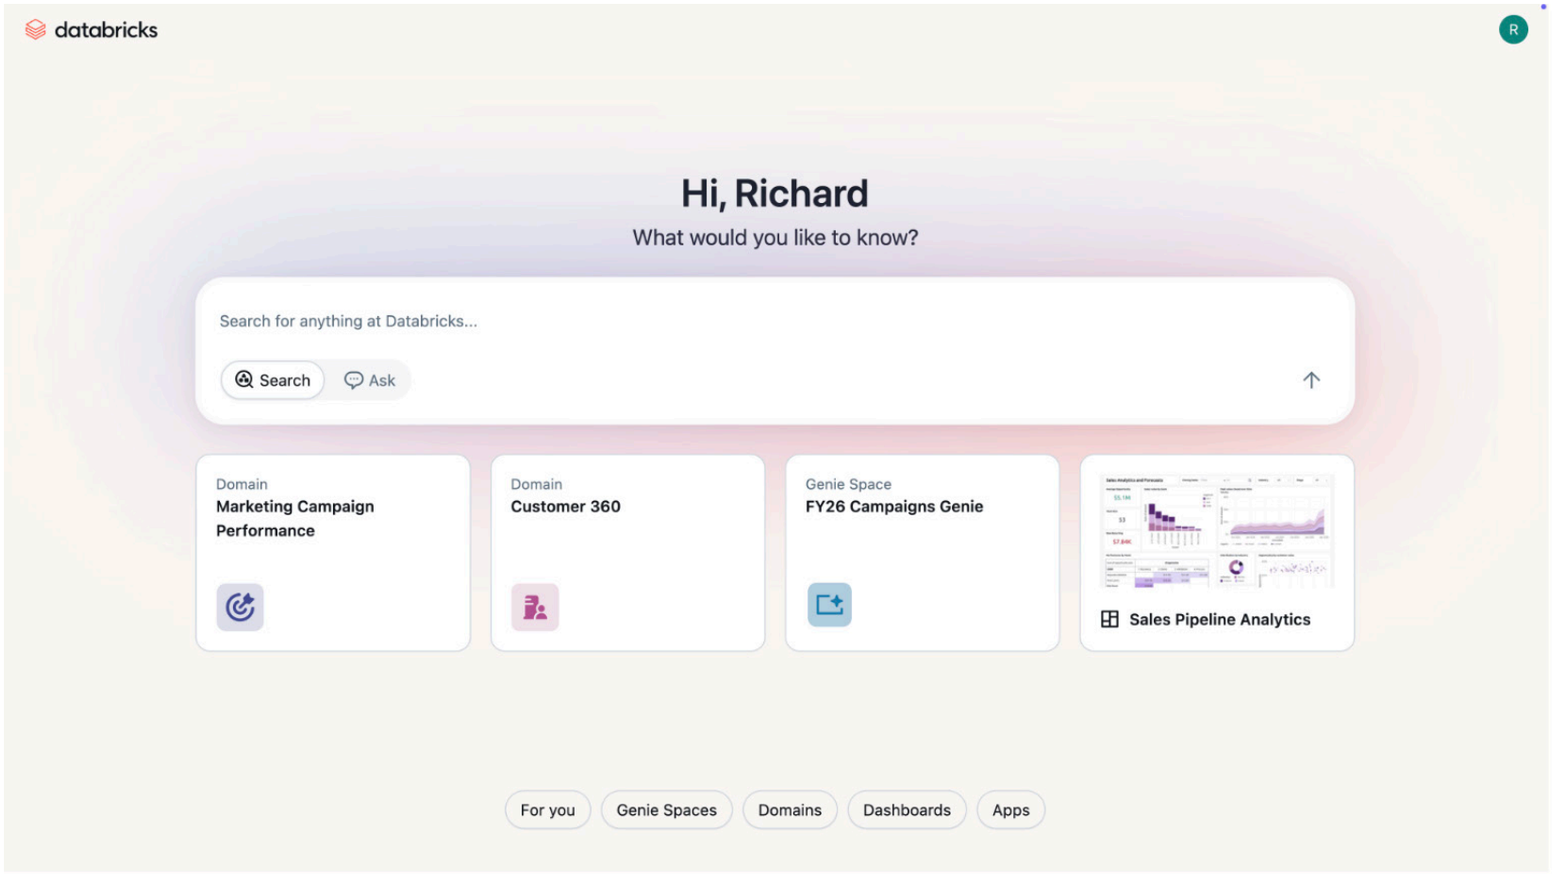

> “Every alert, with full processing context, is centrally logged and delivered directly to customer dashboards so teams can act fast.”
>
> — Merium Khalid, Director, SOC Offensive Security, Barracuda Networks

Self-service security insights and actionable dashboards drive productivity, reduce alert fatigue and transform security teams into bionic, data-driven defenders.

Unified security data foundation

The modern enterprise operates in a dispersed, multicloud, multivendor environment where data fragmentation is the enemy of security. Without a unified foundation, threats go undetected, data remains siloed and compliance risks pile up.

A unified security data foundation grants teams the full-spectrum visibility needed for advanced detection and rapid, accurate investigations. A unified security data foundation with Databricks offers:

- **Real-time data unification:** Lakebase is a fully managed PostgreSQL database in the cloud. It combines the low latency and reliability of a transactional database with the scale and query performance of the Databricks lakehouse architecture — supporting security use cases such as threat intelligence curation, case management and vulnerability management.
- **Centralized telemetry:** Aggregate all security, IT and business data — structured, semi-structured and unstructured — across cloud, endpoint, SaaS and legacy sources, removing blind spots and enabling a "single source of truth"
- **Open and scalable architecture:** Databricks leverages lakehouse principles and open standards (Delta Lake, OCSF) to support petabyte-scale analytics, flexible ingest and future-proof integration
- **Governed and compliant by design:** Built-in Unity Catalog provides fine-grained access controls, lineage and audit trails, meeting the most rigorous regulatory and operational requirements
- **Accelerated context and readiness:** Contextualize security telemetry instantly with business and asset data — giving every analyst richer context for threat hunting, response and compliance

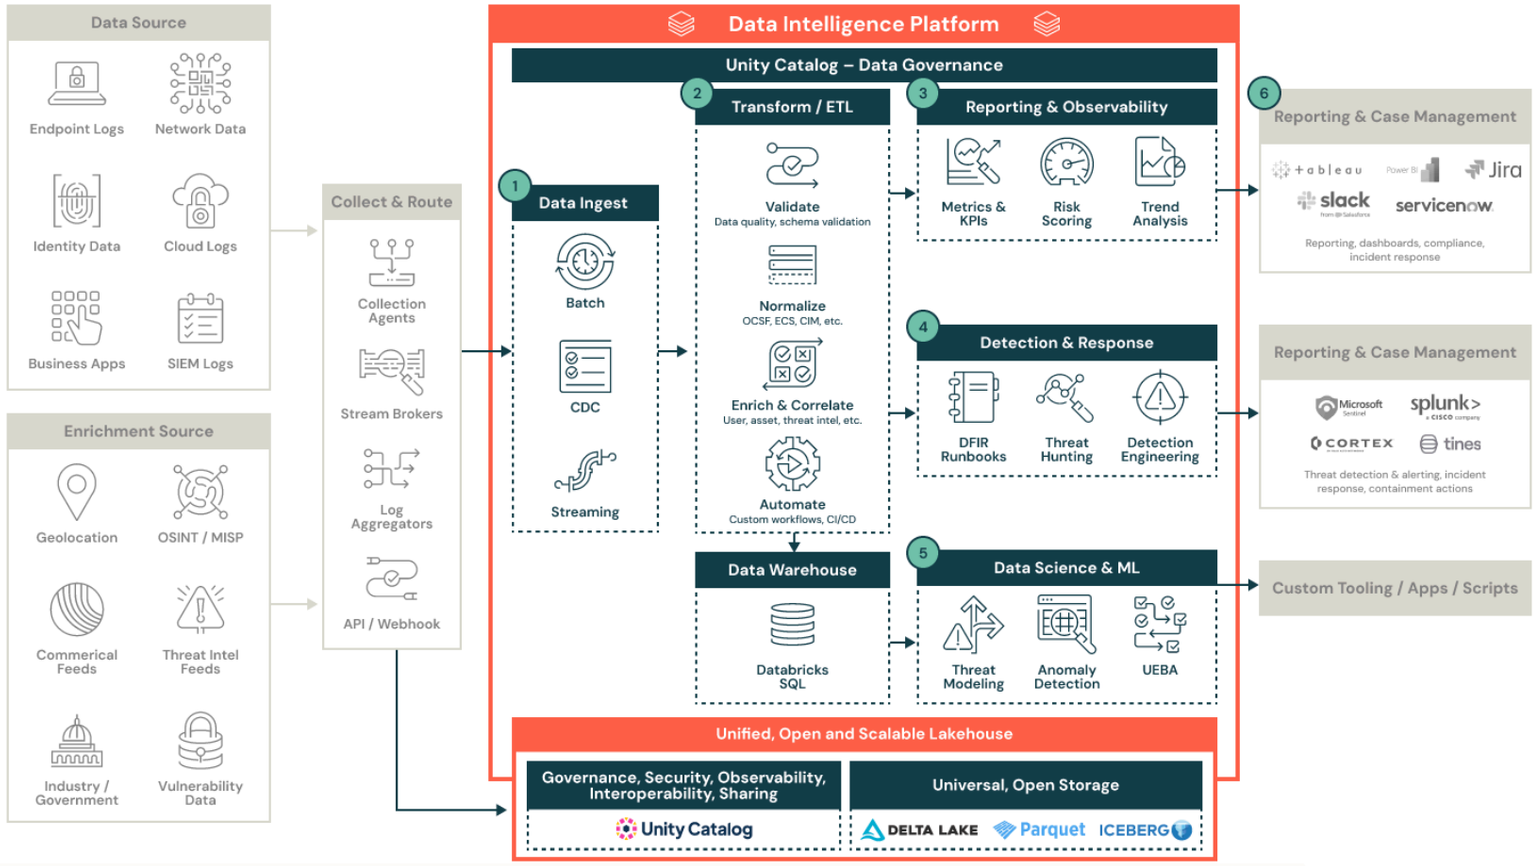

Databricks security reference architecture

> "With Anvilogic on the Databricks Platform, we process threats faster than ever — reducing engineering time by 80% and increasing rule deployment speed 5–6x. We've gained speed, visibility, cost savings and full control over our data."
>
> — Roland Costea, Chief Information Security Officer (CISO), Enterprise Cloud Services (ECS), SAP

With Databricks, your organization's data lakehouse becomes your security lakehouse: a foundation that delivers the breadth and depth modern defense demands.

The Databricks Advantage: Open, Flexible, AI-Driven

The Databricks Data Intelligence Platform is the clear choice for modern cybersecurity.

- **Proven at scale:** The Databricks solution is trusted by security leaders at Abnormal AI, Rivian, Akamai, Palo Alto Networks, Arctic Wolf, Deloitte, Accenture and many others
- **Automated defense with Agent Bricks:** Instantly build, deploy, and optimize secure, production-ready AI agents tailored to your data, automating repetitive detection, triage, and investigation so your team can focus on the threats that matter most
- **Faster, smarter, cheaper:** Customers cut SIEM storage costs by 60–80%. Analysts run 5x more threat hunts. Security executives unlock 2–3x more value per dollar spent.
- **Composable security stack:** Whether you use the Databricks Platform as an offload analytics engine, a parallel detection platform or the unified backbone, you stay in control — integrate at your pace and own your data
- **Continuous innovation:** Regular updates, new connectors, partner integrations and rapid support for evolving standards (e.g., Databricks AI Security Framework 2.0 and Databricks AI Governance Framework for secure AI governance) keep you ahead of attackers and regulators
- **Secure by default:** Fine-grained access, multicloud deployment, AI risk management and total auditability from raw data to alert

Real Security Outcomes: Transformation in Action

The world's most forward-thinking security organizations — including top Fortune 500 companies — are already harnessing the Databricks Platform to infuse data intelligence and AI-driven automation into every corner of their security operations. In this section, we showcase real transformation in action, offering tangible proof that Databricks empowers security teams to break free from legacy barriers, unify fragmented data and deliver measurable outcomes across speed, efficiency and resilience. These stories highlight how leaders in banking, healthcare, manufacturing and technology have rebuilt or leveled up their security operations on the Databricks Platform — unlocking new levels of visibility, cost reduction and protection against modern threats. The following examples provide real-world evidence from customers who've reimagined cybersecurity with the Databricks Data Intelligence Platform at the core, demonstrating what's truly possible when data becomes your strongest defense.

Arctic Wolf

Arctic Wolf processes over 8 trillion security events weekly and protects 10,000 organizations using AI-driven security operations on Databricks, empowering thousands of experts to deliver faster detection and response while reducing operational overhead and unifying data for more effective threat protection.

> "We support thousands of organizations and handle over 8 trillion security events per week. AI and automation let us do more with the same team—reducing mean time to detect, triage and resolve, but always emulating and amplifying our domain experts"
>
> — Michael Mylrea, AI Fellow & Architect, Arctic Wolf

SAP Enterprise Cloud Services and Anvilogic

SAP Enterprise Cloud Services (ECS) runs one of the world's largest private clouds, managing over 200,000 virtual machines while facing tripling data volumes and escalating customer demand for log transparency. Legacy SIEM tools became a barrier to efficient, accurate threat detection — driving up costs and limiting visibility across their sprawling infrastructure. By deploying Anvilogic's AI-powered detection engineering on the Databricks Data Intelligence Platform, SAP ECS automated critical workflows, enhanced MITRE ATT&CK coverage and dramatically reduced engineering time, all while enabling real-time, shareable insights for both internal teams and customers.

> "With Anvilogic on the Databricks Platform, we process threats faster than ever — reducing engineering time by 80% and increasing rule deployment speed 5–6x. We've gained speed, visibility, cost savings and full control over our data."
>
> — Roland Costea, CISO, SAP Enterprise Cloud Services

Abnormal AI

Abnormal AI built its email security analytics and ML workflows on the Databricks Data Intelligence Platform, migrating from legacy Hadoop/EMR to unified Spark and SQL. Today, Databricks lets us process large volumes of email telemetry, deploy models faster, and power governed analytics used to detect and remediate sophisticated email threats in near real time.

> "Databricks unifies our data and compute so we can move faster—building analytics and ML that protect customers from evolving email threats."
>
> — Erin Ludert, Data Lead, Abnormal AI

Customer story

Rivian

Rivian's mission to "keep the world adventurous forever" extends to cybersecurity — protecting a fast-growing digital footprint across IT, manufacturing and product security. With daily data volumes exploding from 150GB to 10TB and SIEM migrations doubling costs each time, Rivian needed to escape legacy limitations. Building the TRAILS platform on the Databricks Data Intelligence Platform, Rivian achieved a 60% SIEM cost reduction, full control and real-time detection across more than 100 data sources — all completed in less than four months — empowering the security team to respond quickly, scale efficiently and future-proof their operations.

> "Lowering our SIEM costs by 60% while migrating 7–10TB of daily data from over 100 sources in under four months was only possible because Databricks gave us full control, scalability and real-time detection."
>
> — Chris Mandich, Director, Cybersecurity Operations, Rivian

Barracuda Networks

Barracuda Networks reduced daily processing and storage costs by 75% compared to legacy systems and delivered customer alerts within five minutes, unifying 100 detection rules and 50 data sources on Databricks to enable rapid, scalable, automated cyber threat defense at lower cost.

> "Databricks transformed our detection engineering—from real-time alerting to cost savings, and, most importantly, the ability to defend customers at scale. We're able to standardize and automate detection and response, so our customers always benefit from the latest intelligence."
>
> — Merium Khalid, Director, SOC Offensive Security, Barracuda Networks

Akamai

Akamai reduced security event data ingestion time from 15 minutes to under 1 minute and now achieves over 85% of customer queries with responses under 7 seconds, enabling real-time analytics at scale for 30% of the internet's traffic using Databricks SQL and Delta Lake.

> "Delta Lake allows us to not only query the data better but to also acquire an increase in the data volume. We've seen an 80% increase in traffic and data in the last year, so being able to scale fast is critical."
>
> — Tomer Patel, Engineering Manager, Akamai

Palo Alto Networks

Palo Alto Networks, a global leader in cybersecurity, needed to tackle challenges from data fragmentation, siloed analytics and rising cloud complexity across their Prisma Cloud platform. By adopting the Databricks Data Intelligence Platform — including Unity Catalog, Delta Lake, Databricks SQL and Lakeflow Declarative Pipelines — Palo Alto Networks dismantled data silos and modernized governance. This unified foundation empowered teams to develop, deploy and iterate AI/ML-powered threat detection features 3x faster than before, reducing operational costs and boosting visibility. With Databricks technology fueling their Precision AI initiative, Palo Alto Networks delivers real-time, AI-driven insights and proactive defenses to customers worldwide.

> "With Databricks, we've turned data complexity into an advantage. Their platform's scalability has reshaped our approach, accelerated our security innovation and helped us to deliver impactful features to customers with unprecedented speed."
>
> — Krishnan Narayan, Senior Distinguished Engineer, Palo Alto Networks

Customer story

ARCTIC WOLF

ECS Enterprise Consultancy Services

Abnormal

RIVIAN

Barracuda

paloalto NETWORKS

Empowering a New Era in Cyber Defense

In today's hyperconnected landscape, where data volume and cyber risks outpace legacy security tools, the path forward is clear: unify, automate and elevate your defenses with Databricks Data Intelligence for Cybersecurity. Leading organizations worldwide — from banks and manufacturers to cloud-first enterprises — are transforming security outcomes by embracing Databricks as the foundation for modern SOCs.

Security teams no longer have to settle for slow, reactive workflows, siloed data and escalating SIEM costs. With the Databricks Platform, your team can leverage a unified security lakehouse, real-time analytics and production-ready AI agents — empowering analysts to move faster, investigate deeper and respond smarter to ever-evolving threats. The results speak for themselves: faster detection, smarter investigations, improved compliance and significant cost savings.

As we've seen in customer stories from Arctic Wolf, Palo Alto Networks, Barracuda Networks, SAP, Rivian and many more, the Databricks Platform enables dramatic gains in detection speed, team productivity, cost savings and cyber resilience — all while giving organizations full data control, seamless integration and future-proof flexibility.

This is your moment to break with fragmented legacy tools, own your data, democratize insights across the SOC and use your enterprise lakehouse to fight advanced threats.

Taking the Next Step

Unify every security signal, empower your teams with AI agents and lead the next wave of resilient, data-driven cyber operations.

- Schedule a hands-on architecture session or a tailored demo
- Pilot Databricks Data Intelligence for Cybersecurity modularly — sidecar, parallel or as your SOC backbone
- Unlock the full value of open standards, agentic automation and real-time analytics

With Databricks, your data and AI platform becomes your strongest defense — future-proofing your cyber strategy for what comes next. Contact us to learn more.

Let's build resilient, intelligent security — together.

databricks

About Databricks

Databricks is the data and AI company. More than 15,000 organizations worldwide — including Block, Comcast, Condé Nast, Rivian, Shell and over 60% of the Fortune 500 — rely on the Databricks Data Intelligence Platform to take control of their data and put it to work with AI. Databricks is headquartered in San Francisco, with offices around the globe and was founded by the original creators of Lakehouse, Apache Spark™, Delta Lake, MLflow and Unity Catalog. To learn more, follow Databricks on LinkedIn, X and Facebook.

databricks