

AI Data Fabric: Unifying Exposure Management for Resilient Security Operations and Governance

TONIC BROCHURE

Introduction

Modern security organizations struggle under an onslaught of data. Security teams ingest vast telemetry from endpoints, networks, cloud workloads, identities, and more – yet attackers still slip through. Traditionally, security data pipelines (like data lakes and SIEMs) and tools were not built for this scale or diversity of data. They collect everything in hopes of catching something, leading to bloated data volumes and high costs, while failing to provide valuable context and meaning.

Stifled by immense data volumes, security teams are flying blind without context, making investigation and remediation harder and longer than needed. Vulnerability management and exposure data (asset inventories, attack surface findings) often live in separate silos, leaving security teams with little insight into which alerts matter most. In parallel, exposure teams work to catalog and prioritize assets and weaknesses, but their findings rarely feed into detection and triage logic.

The result? Inefficient processes, higher costs, and threats lingering undetected. It's clear that traditional approaches are no longer sufficient. To stay ahead of adversaries and justify security spend, **we clearly need a new blueprint that ties everything together.**

Security AI Data Fabric is that solution. This architecture treats security data as a unified fabric: weaving together streaming telemetry, logs, asset and vulnerability data, threat intelligence, and more into an integrated layer. Instead of isolated pipes feeding separate point tools, a security data fabric provides a shared data backbone with active context. By fusing previously siloed sources, it empowers more intelligent analytics and response. The AI Data Fabric is a modern architectural framework that leverages AI-driven data integration and correlation techniques to automate normalization, enrichment, deduplication, and aggregation of high-volume, high-diversity exposure and telemetry datasets – enabling security teams to keep pace with the explosive growth of data.


It aims to deliver the right data, with the right context, to the right tools and teams, at the right time.

But what is a data fabric? What role does AI play in it? How is it different from data mesh and Cybersecurity Mesh Architecture (CSMA)? And how can it be leveraged for Exposure Management?

Let's dive in.

What Is a Data Fabric?

At its core, a data fabric is a modern data management design that enables **flexible, reusable, and augmented data integration pipelines** across complex environments. In contrast to rigid, point-to-point data pipelines, a data fabric is an architecture that **stitches together data from multiple sources and delivers it to various consumers through a unified layer**. The fabric isn't a single product, but an architectural approach comprising integration tools, metadata management, and orchestration. It supports **multiple data integration mechanisms** in combination – from bulk batch ETL processes to real-time streaming (message queues), API-driven access, data virtualization, and microservices-based pipelines. This means a data fabric can gather data from anywhere in your hybrid cloud/on-prem estate and make it available in the form needed, whether for big-data analytics, instantaneous lookups, and anything in between.



Source: Gartner

A defining feature of the data fabric is its use of **active metadata and semantics**. The fabric continuously collects and analyzes metadata about the data itself – schemas, ontologies, data lineage, usage patterns, quality stats – contributing to the **knowledge graph** of how data relates across the enterprise. This rich context lets the fabric automate and optimize data integration. For example, the fabric can dynamically decide **whether to move data or query it in place (virtualize), how to transform or join data on the fly, and even recommend integration design changes**. By leveraging machine learning on metadata, the fabric can suggest improvements like caching frequently used datasets or flagging missing quality checks.

In summary, a data fabric is an architectural approach and technology layer that integrates data across diverse platforms, locations, and types. It enables consistent data management, seamless access, governance, and enhanced insights, regardless of data location (on-premises, cloud, or hybrid environments).

This concept is the basis for Tonic’s AI Data Fabric, applied to Exposure Management.

What Role Does AI Play in Data Fabrics?

Artificial intelligence – especially generative AI – supercharges the data fabric by automating and enhancing many of its functions. **AI is woven into the fabric’s metadata management and integration tooling** to tackle tasks that traditionally bog down data teams. One major role of AI is **automation of integration workflows**. For example, modern data fabric platforms use AI/ML to automatically map and transform data from source to target.

AI also enables **augmented analytics and NLP interfaces** on the data fabric. An example is using natural language questions to query data or discover relationships (“Which critical servers has an unpatched vulnerability that is publicly exploited and published in the last month?”). Under the hood, such capability relies on the fabric’s knowledge graph and context (so the LLM knows what “critical servers” and “vulnerability exploited” mean in the enterprise).

This is why **strong metadata discipline is crucial** – without it, AI may produce inaccurate or “hallucinated” answers. AI in the data fabric thus goes hand-in-hand with mature data

cataloging, ontologies, and context to ensure reliable outputs. When done right, an AI-augmented fabric allows analysts and even non-technical users to interact with data through natural language and receive **context-rich insights** rather than raw tables.

AI in the data fabric is only as good as the information it has about your data. If your asset inventory is incomplete or your log data lacks normalization, the AI cannot accurately map relationships or answer questions. In practice, this means building a robust security data catalog (assets, vulnerabilities, threats, etc.), establishing ontologies (e.g., defining what constitutes an “asset criticality” or “finding priority”), and curating training data for ML models on known good patterns. With this foundation, AI becomes a powerful ally – automating integration tasks, enabling conversational analytics, and even optimizing costs – all within the security data fabric.


Data Fabric vs. Data Mesh (or CSMA)

There's some confusion around this that needs clarification. As new architectures emerge, it's important to understand how a **Data Fabric** compares to concepts like **Data Mesh** and **Cybersecurity Mesh Architecture** (CSMA). Despite similar names, these paradigms address different challenges and complement rather than compete with each other:

- **Data Mesh vs. Data Fabric:** Data mesh is a data management operating model. Data mesh advocates a decentralized, domain-driven approach: each business domain (or in a security context, each security domain like endpoint, network, cloud) owns its data as a product, with standardized interfaces and governance. The mesh's goal is to scale data delivery by aligning it with business context and accountability. In essence, **mesh focuses on organizing people and processes around data** (domain ownership, data-as-product, self-service data access, distributed governance), **whereas fabric focuses on the underlying technology to enable seamless data integration**. And they **can work together**: a data mesh can actually be built on top of a data fabric.
- **Cybersecurity Mesh Architecture (CSMA) vs. Data Fabric:** CSMA is a term coined by Gartner, which defines it as “a **composable** and **scalable** approach to extending security controls, even to widely distributed assets, enabling **tools to interoperate** through supportive layers like consolidated **policy management, security intelligence** and **identity fabric**.” In simpler terms, instead of each security tool (firewall, identity provider, endpoint agent, etc.) acting in isolation, a cybersecurity mesh ensures they can share signals, enforce common policies, and collectively

strengthen defense. Now, how does this differ from a data fabric? The **data fabric** is focused on **data integration** and **analytics**, whereas **CSMA** is focused on **control**

- **integration and operational security.** They operate at different layers but are highly complementary. One can think of the data fabric as the **brain (insights)** and the cybersecurity mesh as the **nervous system (coordinated action)** of a modern security program. Neither replaces the other.



Cyber Security Mesh Architecture

In short, **data fabric, data mesh, and cybersecurity mesh architecture address different dimensions: technology integration, organizational operating model, and security control integration, respectively.** Rather than choosing one over another, we treat these concepts as complementary pieces of a cohesive strategy. We leverage the advantages of the data mesh, grounded in the data fabric, while using LLMs and agentic AI to enhance operability and automate remediation. This holistic vision is what the Tonic **AI Data Fabric** is ultimately about.

Leveraging an AI Data Fabric for Exposure Management

A Data Fabric comes to life when applied to cybersecurity use cases, especially Exposure Management. By embedding asset context, vulnerabilities, misconfigurations, control gaps and threat intelligence directly into security data pipelines, we gain a holistic view needed for implementing a **Continuous Threat Exposure Management** (CTEM) program. CTEM is about continuously assessing and improving the organization's security posture; leveraging a data fabric makes this highly feasible by integrating data from different IT and security tools.

A security data fabric can also break down the wall between threat detection and response, and asset and vulnerability management teams. It blends the two data sets so that detection and response are always performed with awareness of what is being protected and how it's at risk. This synergy can yield significant improvements in key security metrics, such as more accurate detection and triage; faster Mean Time to Detect, Respond and Remediate; deeper Attack Path Analysis and blast radius scoping before or during an incident.

The Tonic Data Fabric: Security that Makes Sense


For organizations ready to elevate their security posture, the roadmap is clear. By weaving together your vulnerability, threat, and IT data – and applying AI to make sense of it – you can stay one step ahead of attackers and regulations alike. Those who embrace this strategy will not only bolster their defenses but also unlock operational resilience and clarity that was previously out of reach.

At Tonic Security, we have embraced this vision: our platform is built on a **proprietary AI Data Fabric**, enabling CISOs, security architects, and SOC teams to harness unified data and intelligent automation as a true strategic advantage. This approach is the key to securing today's dynamic enterprises while futureproofing it for the challenges to come.

Our AI Data Fabric leverages LLMs and visual models to extract meaningful context about any type of entity in the organization that affects its security posture. **The more complex your digital environment is, the more our AI Data Fabric can help** make sense of enterprise data – **break siloes, establish coherence**, and **deliver rich context** to enable security teams to **act fast with clarity and confidence**.

Tonic's AI Data Fabric allows users to make their security data:

- **Integrated:** Connecting and managing data across various environments (cloud, on-premises, multi-cloud, hybrid), enabling comprehensive and coherent visibility.
- **Automated:** Leveraging AI and machine learning to automate data management tasks such as discovery, integration, orchestration, and context extraction.
- **Scalable:** Providing the ability to handle large volumes of diverse data across the enterprise.
- **Continuous:** Delivering continuous data access and processing, enabling up-to-date situational awareness and faster response times.



Tonic Security Data Fabric

Tonic's multi-layered Data Fabric begins with **data onboarding** – where we gather data from security tools, collaboration systems, IT tools, and homegrown applications – and store it in a raw data “**Bronze**” repository.

Then it **normalizes the data**, converting and standardizing data, improving data integrity, and removing sensitive data to ensure privacy by design. Normalization ensures that data can be accurately compared and analyzed. Normalized data is stored in the “Silver” repository.

In the next steps, our AI does most of the heavy lifting. The Data Fabric **extracts context** and insights from structured and unstructured data (such as emails, documents and chats) and uses LLMs and other models to identify entities and their relationships. The context extracted includes business, organizational, geographical, operational, temporal, and adversarial.

The data is **harmonized**, resolving conflicts between data sources, considering source reliability, data recency, and consensus among sources. The harmonization process includes:

- **Correlation:** The process of identifying and matching records that represent the same entity but are described differently across disparate data sources. Variations include format, spelling, or structure. Correlation is foundational for both deduplication and aggregation, as it ensures that data referring to the same entities is accurately matched and integrated, leading to more consistent, comprehensive, and high-quality datasets.
- **Deduplication:** The process of identifying and removing duplicate records from a dataset. When data comes from multiple sources, duplicates are common, and deduplication ensures that each unique piece of information is represented only once.
- **Conflict Resolution:** Data from different sources can often be inconsistent or contradictory. The Tonic AI Data Fabric resolves these conflicts by determining the most reliable or representative value - also known as the “winning” value - using techniques such as consensus analysis, maximum/minimum selection, coalescing, and other domain-specific methods. This ensures data integrity, consistency, and trustworthiness across the unified dataset.
- **Aggregation:** Grouping of several entities to perform analysis as a whole. This is done by combining data from multiple sources into a single, coherent dataset. It is often used to create reports, dashboards, and summary stats from large and diverse sets of data.

Ultimately, the harmonized data is organized into well-defined entities and their interrelationships, forming a Security Graph. This graph acts as the “**Gold**” repository – a single, authoritative source of truth that unifies, contextualizes, and preserves critical information across the enterprise. It enables consistent, graph-driven reasoning and supports advanced security analytics, decision-making, and automation.

Finally, the **application** access layer allows end users to query the graph using both structured queries and natural language.

About Tonic Security

Tonic is reshaping Exposure and Vulnerability Management by providing the context security teams need to accelerate prioritization and remediation of vulnerabilities and threats.

Powered by an Agentic AI Data Fabric, Tonic extracts meaningful and actionable context from unstructured organizational knowledge and threat intelligence, empowering security teams with superior visibility, dramatic reduction in false positives, and a sharp focus on findings that matter. Leading organizations, including Fortune 500 companies, rely on Tonic to slash remediation time and reduce risk to key business processes. To learn more, visit

www.tonicsecurity.com.