Tonic

Agentic Security Data Fabric

The Context Layer Powering Continuous Exposure Reduction

March 2026

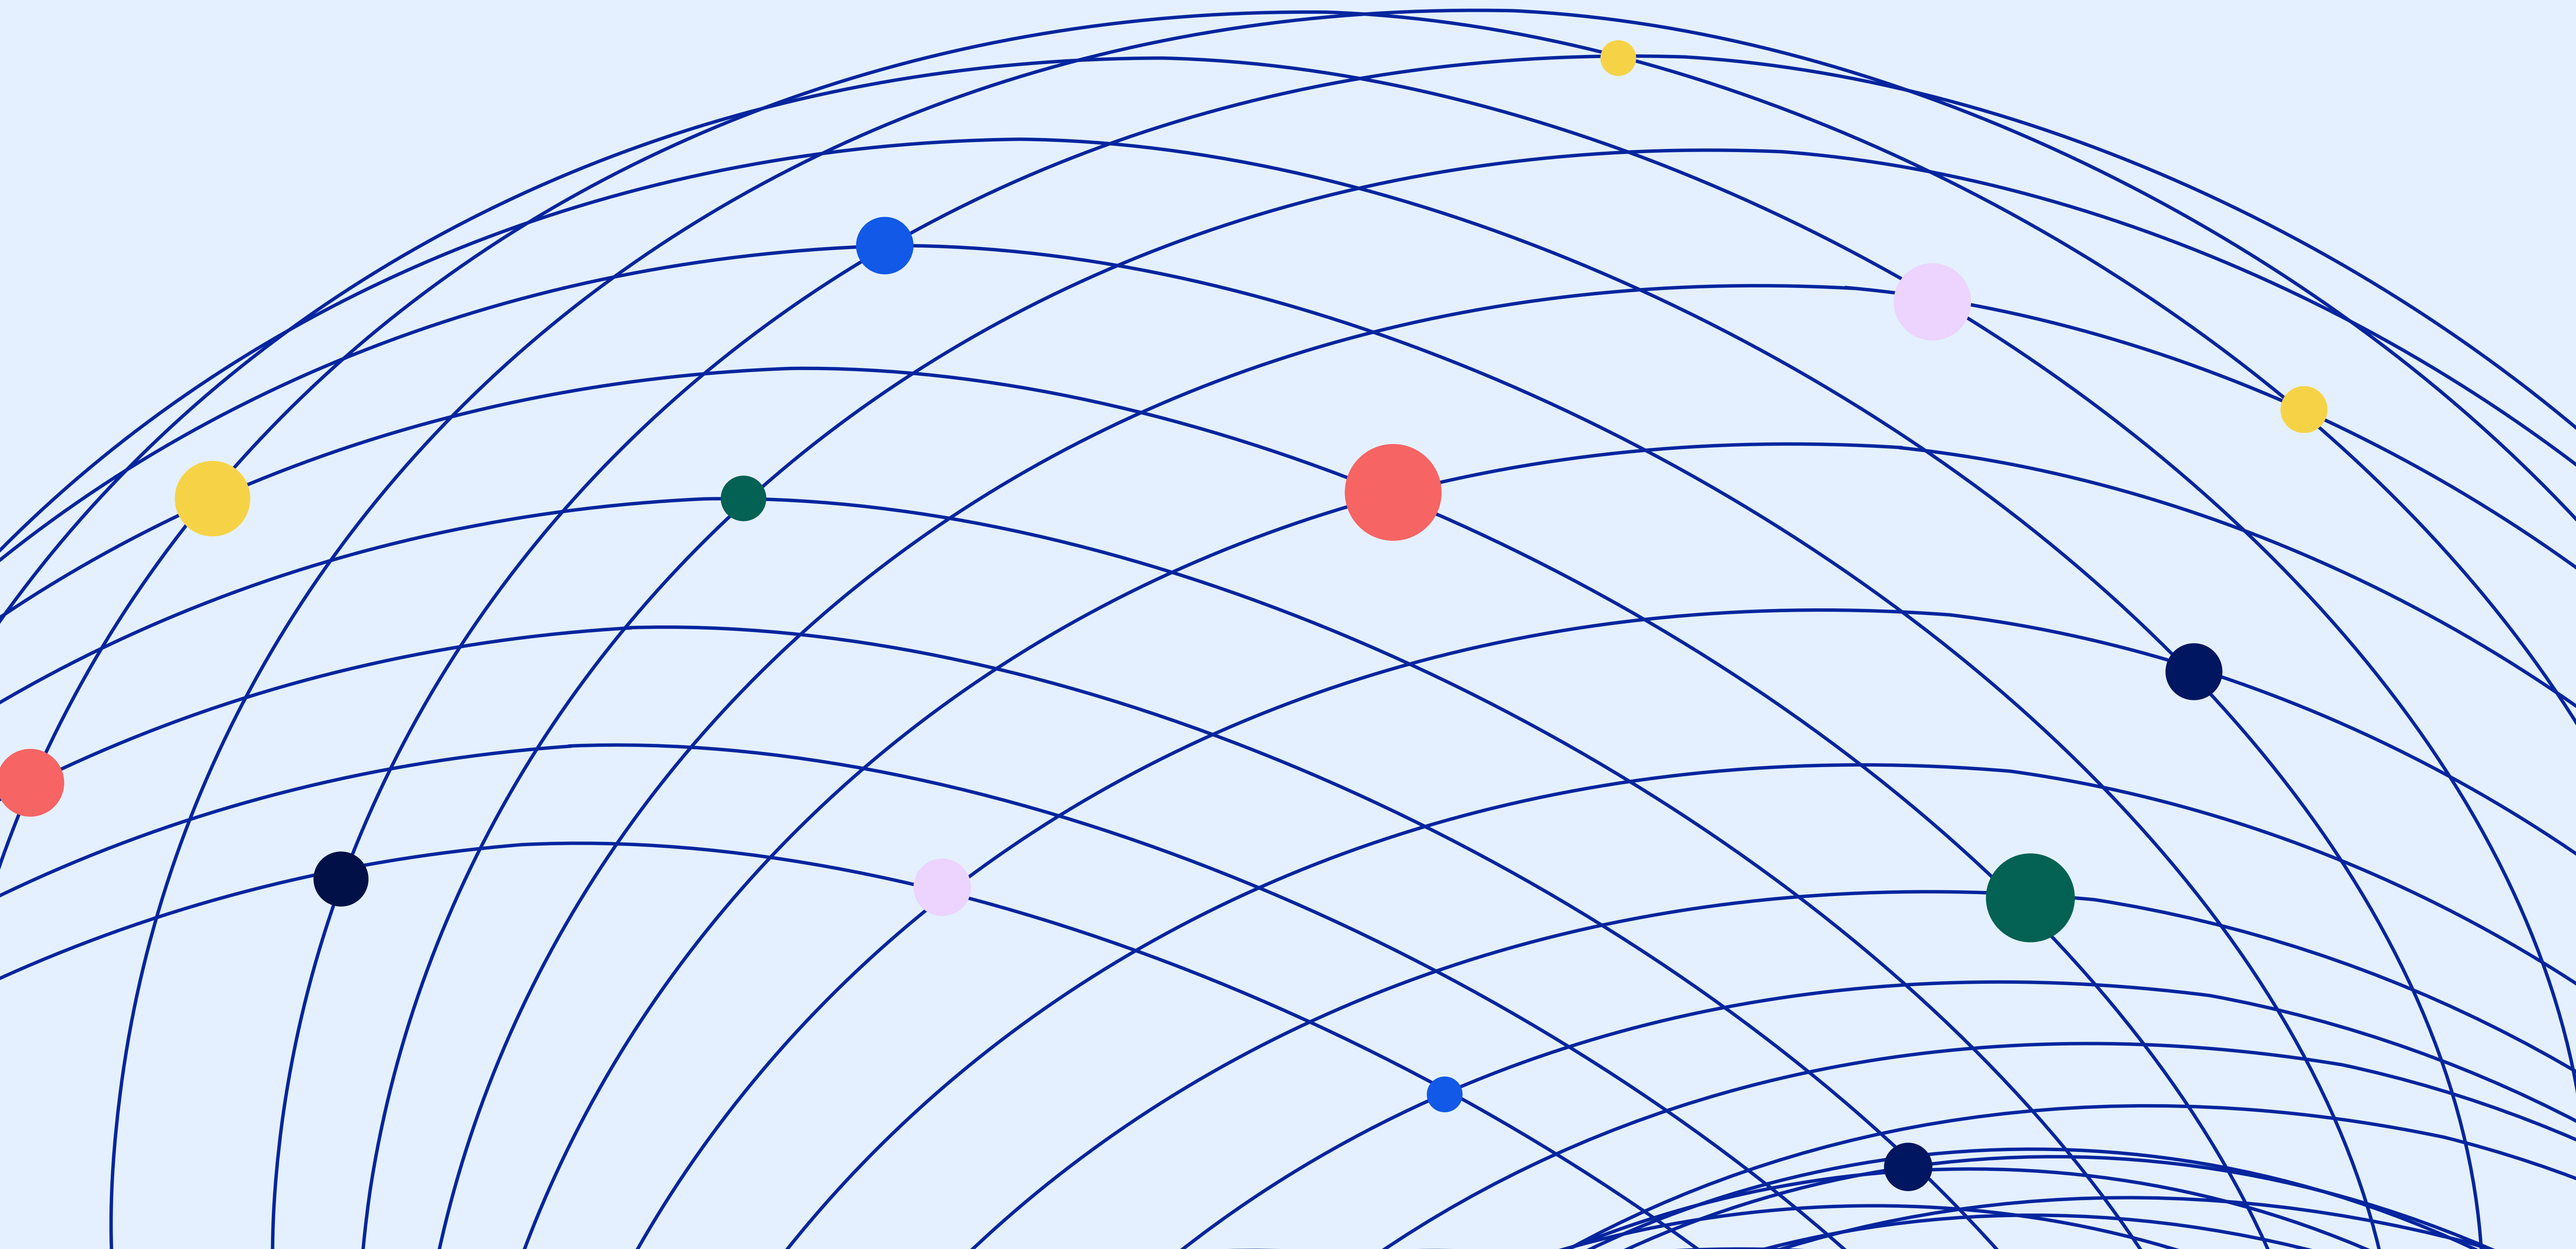

Security Has a Context Problem

Modern security organizations are not failing because they lack data. They are failing because they lack **context**.

Enterprise environments today generate enormous volumes of signals across security and IT systems: vulnerability scanners, cloud platforms, identity providers, endpoint agents, ticketing systems, asset inventories, threat intelligence feeds, and more.

Each of these systems contains a fragment of the truth about the environment. But none of them maintain a **complete or continuously updated understanding of the organization's exposure landscape.**

Critical information about risk is often buried in places security tools cannot easily interpret:

- Asset ownership hidden in service catalogs
- Compensating controls described in tickets
- Business context documented in wikis
- Operational exceptions discussed in chat threads

As environments evolve, this fragmented data quickly becomes **stale, inconsistent, and incomplete**. The result is familiar to every security team:

- Thousands of "critical" vulnerabilities with no clear prioritization
- Manual investigation across disconnected systems
- Remediation efforts that fail to reduce actual risk

Security programs today operate with **visibility but not actionable understanding.** To solve this problem, security platforms must move beyond collecting data toward **maintaining living contextual knowledge of the enterprise environment.**

This is the role of the **Agentic Security Data Fabric.**

The Emergence of the Agentic Security Data Fabric

A traditional data fabric focuses on integrating data from multiple systems.

An **Agentic Security Data Fabric** goes further. It continuously **discovers, reconciles, and contextualizes security and IT data across the enterprise**, maintaining a living model of the organization's digital environment. Rather than acting as a passive data pipeline, the fabric becomes an **active context layer** that represents how assets, exposures, identities, and business operations relate to one another.

AI agents continuously operate within the fabric to:

- Discover new signals and assets
- Normalize and reconcile conflicting data
- Extract context from structured and unstructured sources
- Identify relationships between entities
- Detect inconsistencies and missing information
- Update the organization's exposure model as environments change

The result is a continuously maintained **Security Knowledge Graph** that models the enterprise environment and its evolving exposure landscape. This graph becomes the foundation for **risk reasoning, prioritization, and remediation coordination.**

From Data Aggregation to Context Synthesis

Traditional security platforms aggregate data. Agentic data fabrics synthesize context.

In most organizations, security teams must manually correlate signals from multiple systems in order to understand the true impact of a vulnerability or misconfiguration. Answering even simple questions often requires significant manual effort:

- Which business services depend on this asset?
- Is this vulnerability externally reachable?
- Are there compensating controls already mitigating the risk?
- Who owns remediation?
- Does fixing this issue actually reduce exposure?

Because these answers exist across multiple disconnected systems, security teams spend large amounts of time performing **manual context reconstruction.**

The Agentic Security Data Fabric eliminates this burden by maintaining a **continuously updated contextual model** of the enterprise environment. This allows security teams to move from reactive analysis toward **continuous exposure reasoning.**

The Role of Agentic AI

Artificial intelligence transforms the data fabric from a static integration layer into a s**elf-maintaining context engine.**

Enterprise environments change continuously:

- New assets appear
- Ownership changes
- Controls evolve
- Vulnerabilities emerge and disappear

Maintaining an accurate exposure model therefore requires **continuous reconciliation of signals across systems.** Agentic AI performs this work automatically.

AI agents operating within the fabric continuously:

- Detect new entities and relationships
- Reconcile conflicting asset ownership data
- Extract context from documents, tickets, and collaboration platforms
- Identify stale or missing information
- Maintain the security knowledge graph as environments evolve

This creates a **closed-loop context system**, ensuring that the organization's exposure model remains accurate over time. Rather than requiring constant manual data maintenance, the fabric **maintains itself**.

Relationship to Data Mesh and Cybersecurity Mesh Architecture

The emergence of data fabrics sometimes creates confusion with related architectural concepts such as **Data Mesh** and **Cybersecurity Mesh Architecture (CSMA)**. These concepts address different layers of the modern enterprise data and security stack.

Data Mesh

Data Mesh is an **operating model** for managing data across organizations. It promotes domain ownership of data, treating data as a product and decentralizing responsibility across business units. While data mesh focuses on **organizational structure and governance**, the data fabric focuses on the **technology layer that enables seamless data integration and contextualization.** In practice, data mesh implementations can be built on top of a data fabric.

Cybersecurity Mesh Architecture

Cybersecurity Mesh Architecture focuses on **security control interoperability.** Rather than isolated tools operating independently, CSMA enables security systems to share signals, policies, and intelligence across distributed environments. The Agentic Security Data Fabric complements this model by providing the **context layer that enables exposure reasoning and coordinated remediation across the cybersecurity mesh**. If the cybersecurity mesh is the **nervous system of security operations**, the data fabric acts as the **brain that maintains understanding of the environment.**

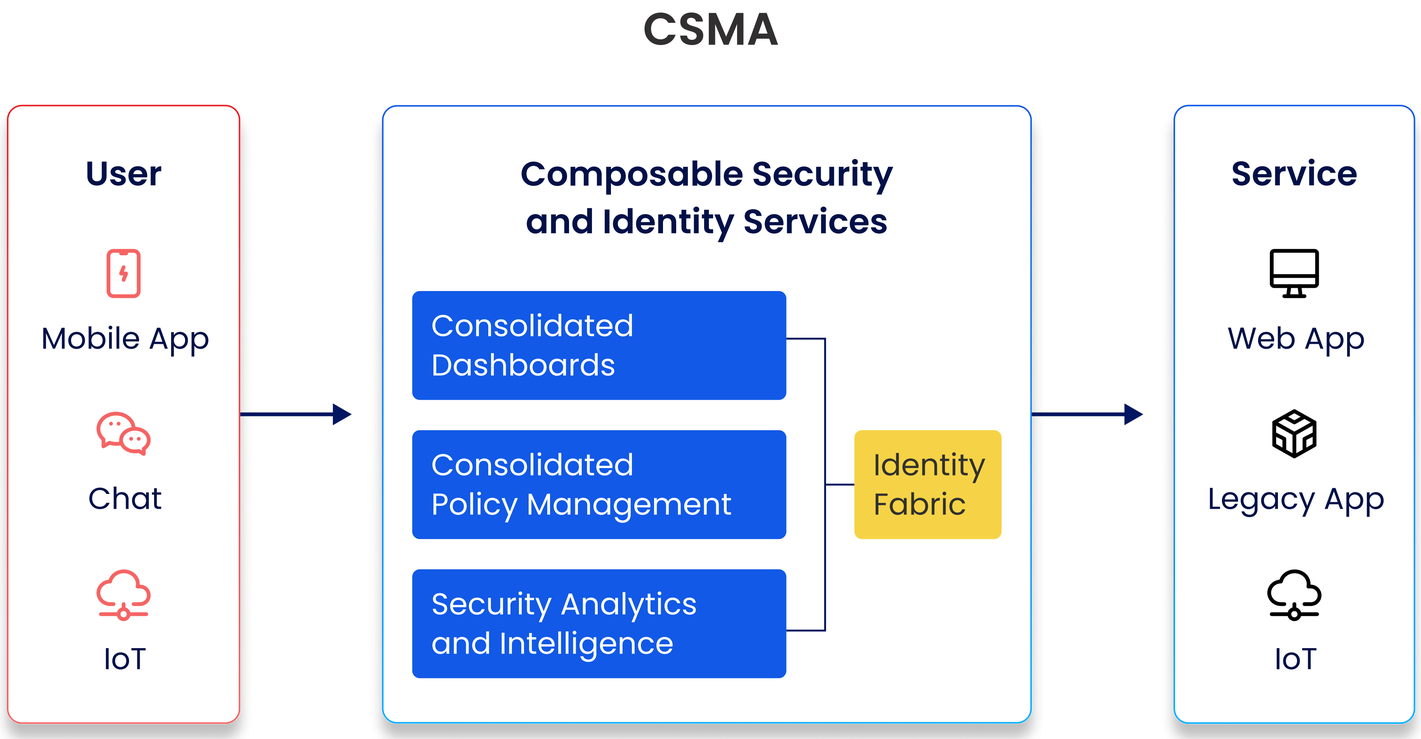

Enabling Continuous Exposure Management

Exposure management requires more than identifying and scoring vulnerabilities. It requires understanding the **context in which exposures exist**.

Security teams must determine:

- Which assets are affected
- How those assets support business operations
- What controls mitigate the exposure
- Whether attackers could realistically exploit it
- Who owns remediation
- Whether remediation actually reduced risk

These answers rarely exist within a single system. The Agentic Security Data Fabric continuously synthesizes signals from security, IT, and collaboration platforms to maintain an accurate exposure model. This enables several critical capabilities:

Contextual Risk Prioritization

Security teams can focus on the small subset of exposures that truly threaten business operations.

Attack Path Analysis

Understanding relationships between assets enables identification of realistic attacker paths.

Blast Radius Estimation

Security teams can evaluate how exposures impact critical systems and services.

Remediation Coordination

The platform identifies responsible teams and orchestrates remediation workflows.

Continuous Validation

The system verifies whether remediation efforts actually reduced exposure.

Instead of producing static reports, the fabric powers **continuous exposure reduction loops.**

Architecture of the Agentic Security Data Fabric

The Agentic Security Data Fabric operates across several functional layers that transform fragmented signals into actionable security context:

Signal Ingestion

Data is continuously collected from across the enterprise environment, including:

- Security platforms
- IT and asset management systems
- Identity and access management systems
- Cloud infrastructure platforms
- Collaboration and knowledge systems
- Custom enterprise applications

These signals form the raw inputs used to model the organization's exposure landscape.

Data Normalization

Signals from different systems use different schemas and formats. Normalization standardizes this data so that entities such as assets, vulnerabilities, identities, and controls can be consistently interpreted across sources. This ensures that data from disparate platforms can be accurately compared and analyzed.

Context Extraction

Many critical insights about risk exist within **unstructured data**.

AI models extract contextual information from sources such as:

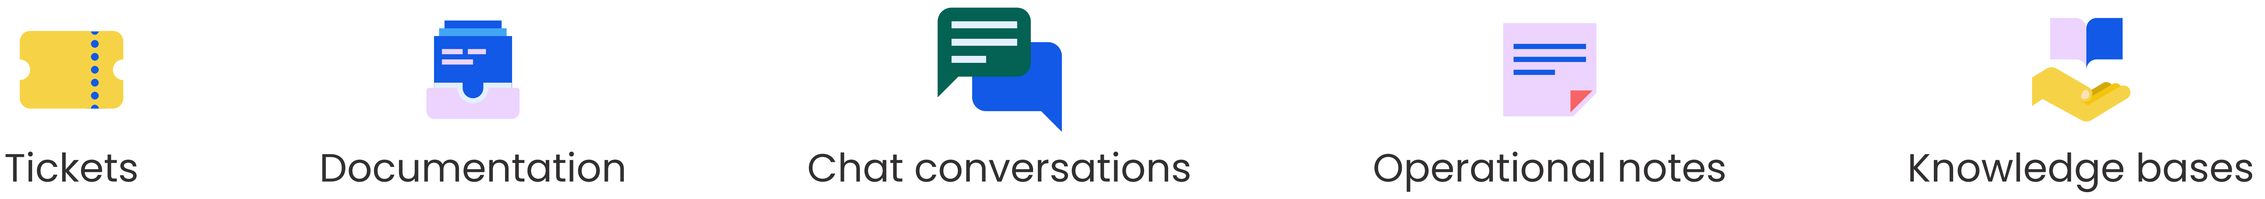

This context often includes:

- Business criticality
- Asset ownership
- Operational dependencies
- Compensating controls
- Risk exceptions/appetite

Data Harmonization

Enterprise data frequently contains inconsistencies and conflicts. The fabric continuously reconciles these discrepancies through:

- Entity correlation
- Deduplication
- Conflict resolution
- Consensus analysis

This process ensures that the platform maintains a **consistent and trustworthy representation of the environment.**

Security Knowledge Graph

Once harmonized, entities and relationships are organized into a continuously evolving **Security Knowledge Graph**. This graph represents:

- Assets and services
- Vulnerabilities and exposures
- Identities and permissions
- Controls and mitigations
- Organizational relationships
- Business dependencies

The graph acts as a **living model of the organization's exposure landscape**, enabling reasoning across complex relationships.

Decision and Execution Layer

Applications and AI agents operate on top of the knowledge graph to perform:

- Exposure prioritization
- Attack path reasoning
- Remediation coordination
- Continuous validation of risk reduction

This layer enables security teams to move beyond analysis toward **operational risk reduction.**

The Foundation of Agentic Exposure Management

Security programs today are overwhelmed by fragmented data and disconnected tools. Visibility or scoring alone are no longer sufficient. Organizations need systems that can **continuously maintain contextual understanding of their environment and coordinate action across teams.**

The Agentic Security Data Fabric provides this foundation. By synthesizing fragmented security and IT signals into a continuously maintained exposure model, the fabric enables organizations to:

- Prioritize risk with business context
- Coordinate remediation across teams
- Reduce exposure faster and more reliably

This architecture represents the foundation of **Agentic Exposure Management**, a new approach to security operations where intelligent systems continuously reason about risk and drive remediation outcomes.

Organizations that adopt this model move beyond managing findings toward **continuously reducing real-world risk**.

Tonic Agentic CTEM Platform

An exposure intelligence architecture powering continuous, governed risk reduction.

Decision & Execution Layer

Tonic Application & MCP

Exposure intelligence and agentic workflows delivered wherever security teams work — inside the **Tonic app**, through **MCP-powered integrations**, or embedded in the tools and workflows they already rely on.

Tonic Agents | 3rd party Agents

Domain-specific agents that continuously prioritize, coordinate, and validate exposure reduction at scale.

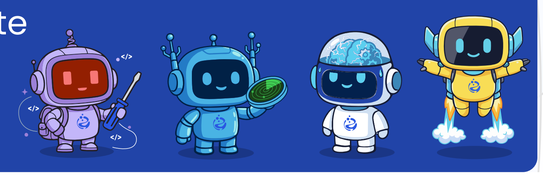

Agentic Control Plane

The trust layer that **controls what agents can see, do, and change** — enforcing permissions, requiring human approval for high-risk actions, and maintaining a complete audit trail.

Security Data Fabric

Exposure Intelligence

A multi-dimensional **exposure graph** that reasons across business, operational, and adversarial context to drive **precise risk decisions and actions.**

Context Engine

Continuously infers business criticality, ownership, exploitability, and blast radius — transforming **structured and unstructured data** into **decision-ready context**.

Data Foundation

A continuously self-improving fabric that **ingests, reconciles, and normalizes all security and IT-related signals** across the enterprise.