Enterprise Strategy Group

SHOWCASE

MAY 2025

High Cardinality at Scale: Rethinking Observability for Cloud-native Environments

Torsten Volk, Principal Analyst

Abstract: The continuously increasing complexity of modern IT environments has led to the fragmentation of observability efforts. Today, 69% of organizations rely on six or more observability platforms, with nearly one-third managing over ten.[1] As a result, 27% of IT leaders acknowledged that their telemetry data is siloed across disparate tools and teams, making it difficult for their organizations to achieve full-stack system visibility.[2]

Monitoring application health and performance across hybrid, multi-cloud environments has become increasingly resource-intensive and prone to blind spots. This paper examines the operational challenges created by fragmented observability and, using Apica's unified telemetry pipeline and object store architecture as a case study, outlines the requirements for a cohesive approach that enables organizations to bridge silos, enrich data with business context, and accelerate the delivery of actionable insights.

Overview – The Problem

As enterprises modernize a growing share of their application portfolio, most find themselves managing a heterogeneous mix of traditional monolithic applications, including mainframe and modern distributed microservices, typically spanning on-premises, multi-cloud, and edge environments.[3] Each application stack generates distinct streams of telemetry data, including logs, traces, metrics, and events, across diverse layers of infrastructure, applications, networks, security systems, business processes, AI workloads, DevOps pipelines, and user activities. Successful observability platforms need to continuously collect, interpret, and store telemetry data within their application and business context to optimize operational efficiency, ensure security, and enable proactive data-driven decision-making.

[1] Source: Enterprise Strategy Group Research Report, *Transforming Observability and Monitoring Through AI*, April 2025.
[2] Ibid.
[3] Source: Enterprise Strategy Group Complete Research Report, *Application Modernization and the Role of Platform Engineering*, October 2024.

Figure 1. Enterprise Telemetry Sources and Categories

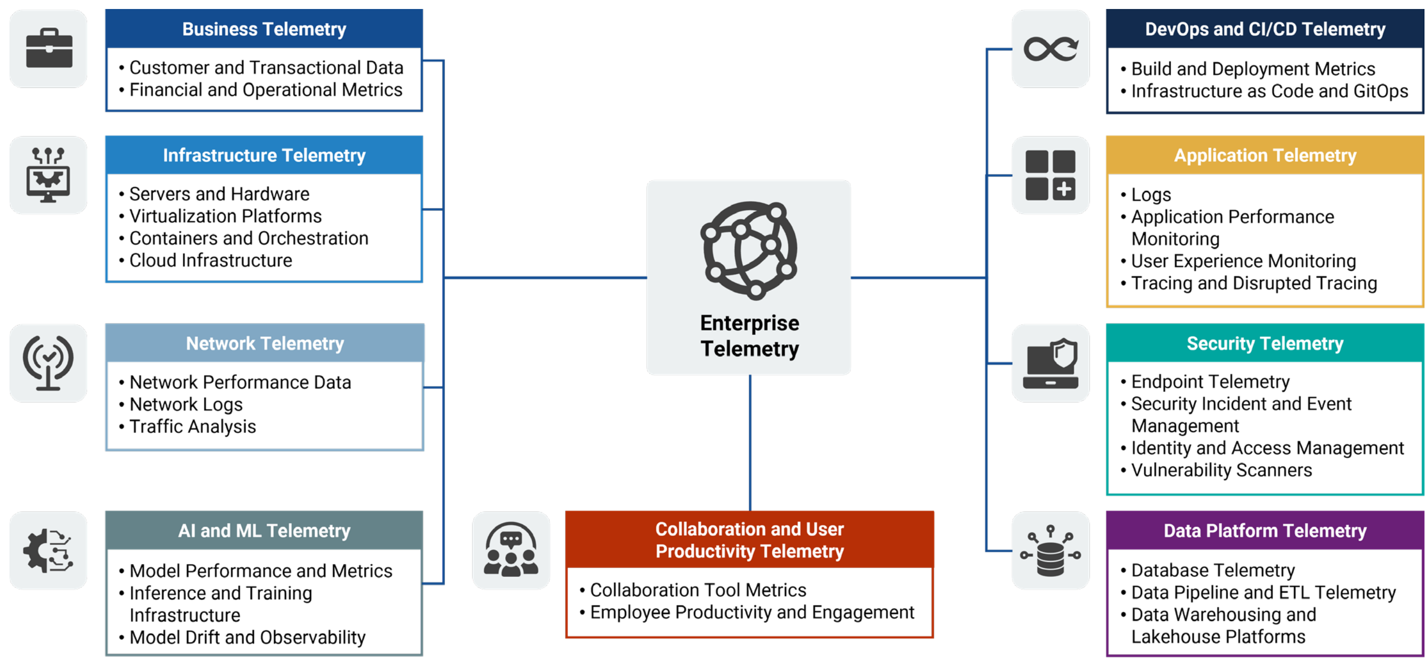

Source: Enterprise Strategy Group, now part of Omdia

The OpenTelemetry Revolution: More Context and Exponential Data Growth

OpenTelemetry has quietly become the de facto standard for capturing, collecting, and exporting telemetry data across today's cloud-native and distributed architectures. By providing a broad set of instrumentation libraries and auto-instrumentation capabilities, OpenTelemetry removes much of the pain traditionally associated with gathering complete telemetry across every layer of the application stack—regardless of whether the data resides in hybrid environments, multiple clouds, or the rapidly evolving edge.

As organizations dismantle their monolithic apps and move to microservices, telemetry volumes grow exponentially. OpenTelemetry not only makes comprehensive data capture possible across wildly diverse application architectures and infrastructure platforms, it also contextualizes this data: Kubernetes metadata, cloud provider details, device IDs, version numbers, geolocation tags, user activity counts, and more. Automatic context propagation across services dramatically increases attribute cardinality by linking telemetry to application, infrastructure, and cloud environments in real time.

This results in massively richer observability data—a clear advantage—but it also accelerates the growth of telemetry data volumes beyond what many traditional observability platforms were ever designed to handle. The industry now faces the double-edged sword of needing to manage exponentially larger data sets while still delivering cost-effective, actionable insights at scale.

The Observability Crisis: Four Systemic Failures Enterprises Must Overcome

As organizations scale their application environments across hybrid and multi-cloud infrastructures, traditional observability approaches reveal critical systemic weaknesses. The following four interconnected pain points prevent enterprises from achieving effective, sustainable, and secure observability at scale:

1. **Operational inefficiency from tool sprawl.** The proliferation of observability tools creates substantial operational overhead. Teams must maintain expertise across multiple platforms, each with its own query

language, visualization capabilities, and management requirements. This fragmentation leads to increased training costs, reduced productivity, and ultimately slower mean time to resolution when issues arise.

2. **Financial unsustainability with traditional solutions.** As data volumes grow exponentially through increased OpenTelemetry adoption, traditional solutions become financially untenable. Their licensing models, which scale with data ingestion, force difficult decisions about what telemetry data to keep versus what to discard—often sacrificing observability for cost control.
3. **High-cardinality data management challenges.** Modern cloud-native environments generate unprecedented numbers of dimensions and labels (service names, container IDs, pod names, regions). Traditional observability architectures were not designed for this reality and often struggle to take on these high-cardinality loads, resulting in:
    - Degraded query performance.
    - Unpredictable storage requirements, causing cost spikes.
    - Inability to apply granular retention policies to high-dimension data.
4. **Siloed security and operations teams.** The separation between IT operations and security monitoring creates dangerous visibility gaps. When security and operations use different tools with different data stores, correlating security events with operational telemetry becomes very difficult, extending detection and response times for potential threats.

These structural weaknesses are no longer edge cases; they are the new normal in large-scale IT environments. Addressing them requires a fundamental shift, moving from fragmented, cost-prohibitive observability strategies toward an integrated, scalable approach that prioritizes full-stack visibility, operational efficiency, and proactive security. The following sections explore the key requirements and design principles for achieving this transformation.

How Apica Transforms Fragmented Telemetry into Unified Observability

Apica's approach tackles these four weaknesses via a comprehensive architectural transformation (see Figure 2).

Figure 2. Apica's Telemetry Data Processing Sequence

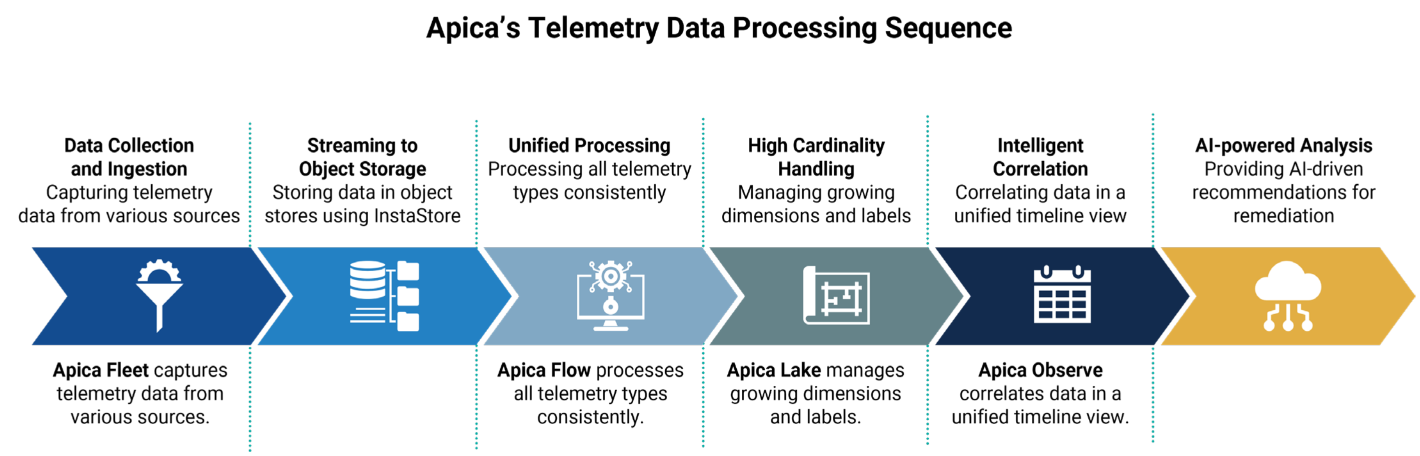

Source: Enterprise Strategy Group, now part of Omdia

Data Collection and Ingestion

Apica's architecture tackles one of the core challenges of modern observability: bringing together logs, metrics, traces, and events into a unified, actionable view without losing their individual analytical value. Specialized adaptors for OpenTelemetry, Prometheus, Fluentbit, Kubernetes, cloud providers, and network devices first normalize incoming data into a standardized internal structure. Along the way, Apica enriches metadata, harmonizes schemas, and injects correlation IDs to maintain context across diverse telemetry streams.

At the core of this architecture is a multi-tier data model that preserves source-specific attributes while adding a common metadata framework (i.e., unified timestamps, consistent entity identifiers, and environmental tags). This creates a hierarchical structure where source-specific details live within dedicated namespaces, while still allowing the data to be analyzed within its relevant context.

The critical enabler here is robust correlation across telemetry types. Apica's architecture establishes the relationships needed to accelerate root cause analysis, offering a unified timeline view without sacrificing the depth of individual telemetry types. This standardization lays the foundation for consistent, high-performance processing at scale while keeping the unique strengths of logs, metrics, traces, and events intact.

Streaming to Object Storage

At the core of Apica's architecture is InstaStore™, a patented technology that streams telemetry data directly to cost-effective object storage (e.g. AWS S3 or Azure Blog Storage) as the primary storage tier, rather than relying on expensive memory-intensive pipelines or proprietary storage layers.

Because data is streamed immediately and persistently into object storage, telemetry is durably recorded from the moment of ingestion. Unlike in-memory-first pipelines, which risk data loss during node failures, network interruptions, or processing bottlenecks, Apica's architecture ensures that once data hits the pipeline, it is safely stored without relying on volatile memory buffers.

Unified Processing Framework

Apica's unified processing framework bridges the traditional divide between different telemetry types by applying consistent transformation, enrichment, and analytics operations across all data. The architecture implements specialized indexing strategies optimized for each telemetry type while maintaining a common query interface, enabling users to seamlessly work with logs, metrics, traces, and events using familiar query patterns.

Unlike traditional systems that require different processing techniques for each data type, Apica's framework handles high-throughput numerical processing for metrics while simultaneously supporting full-text search for logs, hierarchical relationship analysis for traces, and state-transition tracking for events. This architectural consistency eliminates the need to switch between specialized tools while maintaining the unique analytical capabilities required for each telemetry type. The processing framework adapts dynamically to workload characteristics, automatically scaling resources based on data volume, cardinality, and query patterns. This adaptability ensures consistent performance regardless of which telemetry type dominates the workload, enabling reliable analysis even as the observability data mix evolves over time.

High-cardinality Native Design

Apica takes a new approach to managing high-cardinality telemetry data by separating metadata from the core payload at the architecture level. Traditional time-series databases often crumble under the exponential explosion of dimension combinations. In contrast, Apica uses a multi-tier, distributed indexing strategy where high-cardinality fields are dynamically stored and managed independently from the telemetry content itself.

This design unlocks a series of important advantages: Advanced metadata compression, smarter storage tiering based on dimension complexity, and the ability to scale predictably without the cost and performance penalties that typically come with high-dimension environments. As a result, Apica delivers consistently fast query performance, even as cardinality expands into the millions of unique time series combinations, while also enabling precise retention policies at a granular level—capabilities that traditional observability architectures simply were never built to handle.

Intelligent Correlation

Apica's intelligent timeline view automatically correlates metrics, logs, traces, and events in a single interface. This eliminates the manual effort of piecing together information from different tools, dramatically accelerating troubleshooting:

- **Unified investigation:** Correlates all telemetry types in a single view.
- **Cross-stack visibility:** Identifies relationships between infrastructure, applications, and business metrics.
- **Complete context:** Ensures full investigation context with 100% data reliability.
- **AI-powered recommendations:** Suggests precise remediation actions based on historical patterns.

AI-powered Root Cause Analysis and Remediation

Apica's architecture goes beyond traditional observability by integrating machine learning directly into the telemetry pipeline. With a unified, contextualized data model spanning logs, metrics, traces, and events, Apica creates the foundation for AI to surface meaningful insights. The system connects the dots across infrastructure, applications, and business metrics to uncover root causes that would otherwise remain hidden.

The AI engine continuously learns from historical incidents, building a living knowledge base of what worked to fix similar issues in the past. When it detects an emerging problem, Apica does not just alert the team; it proactively recommends targeted remediation steps based on proven outcomes. This approach lightens the cognitive load on operations teams and enables organizations to address problems before users ever notice an impact. It moves observability from being a reactive monitoring exercise to becoming a true performance optimization engine—one that drives measurable business outcomes at scale.

Conclusion: The Future of Enterprise Observability

Management of telemetry data pipelines and storage has become the limiting factor in today's world, where one traditional enterprise application is often modernized into 20 or 30 individual microservices. With infrastructure and application layers now auto-instrumented by default, organizations are collecting massive volumes of logs, metrics, traces, and events. But volume alone does not drive value. What matters is the ability to correlate this telemetry across different formats and sources to generate real-time, actionable insights.

For example, a latency spike in a user-facing service might first show up as a metric in a dashboard. To understand the impact, we need to trace the request through the microservices chain and discover that the bottleneck is in the authentication service. We then dive into logs from that service and see a flood of failed token validation attempts tied to a specific API client. None of these signals—metrics, traces, or logs—tells the full story on its own. But when combined in context, they paint a complete picture: What happened, where it happened, why it happened, and what to do about it.

This convergence is what modern observability demands: A telemetry pipeline that not only handles scale but also preserves the relationships across data types, systems, and services, enabling teams to move from isolated symptoms to true root cause understanding. Apica focuses on exactly these capabilities through its unified architectural approach.

The future belongs to unified telemetry pipelines that can process all data types equally well while maintaining their unique analytical value. By eliminating data silos, implementing cost-effective storage strategies, and providing intelligent correlation across telemetry types, enterprises can transform observability from a costly operational burden into a strategic advantage that drives business value.

Organizations that successfully implement this unified approach are likely to benefit from dramatically faster mean time to resolution, proactive issue prevention through AI-driven insights, and significantly reduced infrastructure costs. Perhaps most importantly, they will gain the operational agility needed to adapt to rapidly evolving technology landscapes without sacrificing visibility or control. In the end, unified observability is not just about better monitoring; it is about enabling faster innovation, more reliable services, and superior customer experiences in an increasingly complex digital world.

About Enterprise Strategy Group
Enterprise Strategy Group, now part of Omdia, provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com