ARCHITECTURE DESIGN GUIDE

A guide to practical investigations

Modern investigations require AI-ready architecture

Cribl®

AI can transform investigations, allowing rapid analysis of vast datasets

Yes, AI can dramatically accelerate investigative speed and scale but also overwhelm fragmented, cost-constrained telemetry architectures — requiring unified, accessible, high-performance data pipelines and data stores to fully realize AI-enabled operations.

AI is real and has a lot of promise, however it cannot fix a fragmented data architecture, and in many cases, it can magnify it. AI agents require broad, high-speed, and cost efficient access to complete telemetry across the enterprise. Without architectural modernization, AI workloads place additional strain on already brittle systems, overwhelming ingestion pipelines, search infrastructure, and licensing models.

To enable AI-driven investigations, organizations must rethink their data management strategy. The foundation is a simplified, unified ingest and search architecture that makes all relevant telemetry accessible without license bottlenecks, while optimizing for performance, scalability, and cost control. Data should be collected once, enriched at ingestion, routed intelligently, and stored in a format that supports elastic, high-concurrency search for both human analysts and AI agents. By decoupling data ingestion, storage, and analytics, composable architecture provides flexible pipelines, expanded data access, and scalable search across massive telemetry volumes.

The result is an AI-ready environment where telemetry becomes a shared strategic asset that powers faster investigations, reduces operational friction, lowers tool costs, and enables security and AI teams to operate with greater speed, scale, and confidence in an increasingly complex digital landscape.

AI can't fix what telemetry breaks

AI is often positioned as the cure for slow, manual security investigations, yet many organizations are discovering that progress remains frustratingly incremental. The issue is not the intelligence of the algorithms, it is the condition of the data beneath them. When telemetry is fragmented across disconnected tools, inconsistently structured, or enriched only after ingestion, AI cannot meaningfully accelerate outcomes. It simply moves faster through the same bottlenecks. As cloud and hybrid environments scale, telemetry volumes surge and overwhelm legacy architectures. To control costs, teams sample logs, delay enrichment, or isolate pipelines by platform. The result is greater noise, reduced context, and increased operational friction. AI systems query and correlate data at a scale humans never could, but they excel with structured, contextualized telemetry. Organizations seeing real gains are redesigning their data foundations, building flexible, enriched pipelines that support multiple tools simultaneously. In the AI era investigative speed depends not on smarter models, but on AI-ready data.

Investigative architectures: Today

Investigators are fighting their architecture: Security and IT investigators are pressured to move faster while environments become more complex.

Telemetry is expanding rapidly across cloud, applications, networks, and identities. AI is shifting from passive insight to autonomous action.

Agentic systems don't just summarize dashboards, they generate hypotheses, execute parallel queries, and move across systems at machine speed. Most architectures, including observability tools and SIEMs, were never designed for that scale or concurrency.

Investigation workflows remain fragmented — security data lives in one platform, observability in another, and business analytics elsewhere. Teams optimize telemetry for their own tools, creating inconsistent schemas, enrichment methods, and retention policies. Investigators must pivot between systems, manually stitch together context, and rebuild timelines from disconnected datasets. Collaboration slows, and root cause analysis becomes inefficient and uncertain.

AI inherits these weaknesses. Even advanced models struggle when data is incomplete or inconsistent. Many AI initiatives stall because the underlying data foundation is fragile. Without structured, normalized telemetry, AI requires more queries, more compute, and more cost to produce reliable answers. Agentic investigation also drives query explosion. Humans investigate sequentially and prune quickly. AI evaluates multiple hypotheses in parallel across systems, sharply increasing concurrency and infrastructure strain. Architectures built for human-scale workloads hit performance limits, and licensing costs escalate.

The answer isn't another tool, it's an investigation-optimized architecture: comprehensive telemetry collection, pipeline based normalization, flexible routing, and scalable search across real-time and historical data. The future of investigations depends not just on smarter models, but on modernizing data architecture to support AI at scale.

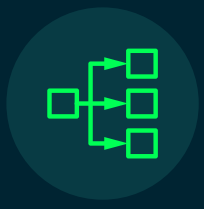

Connect Your Resources

Network connectivity architecture defines how devices, systems, and applications connect and communicate across an organization. By establishing traffic flow paths, segmentation boundaries, and security enforcement points across core, distribution, access, edge, and cloud layers, it creates a structured and resilient digital ecosystem. This design ensures scalability, performance optimization, operational continuity, and predictable visibility into system interactions.

That structured visibility becomes a powerful foundation for modern investigations. Defined traffic paths and logical segmentation produce consistent, high-value telemetry at known control points, making analysis faster and more accurate.

Cribl Search-in-Place capability uses a network connectivity architecture by enabling teams to query telemetry anywhere they have access to (connected), such as object storage, data lakes or Analytics services and platforms without rehydration or re-indexing. By using the organization's structured connectivity architecture, Search-in Place delivers federated, high-speed access to distributed data, reducing cost and enabling scalable, AI-ready investigations across hybrid environments. .

Investigative architectures: Tomorrow

Strategic advantage of AI-driven investigations: Transform investigations by accelerating detection, improving accuracy, and scaling intelligence.

Artificial intelligence, particularly agentic AI, marks a fundamental shift in how investigations are conducted across security, IT, and platform engineering. Like the PC and cloud before it, AI goes beyond efficiency gains; it reshapes workflows, operational scale, and decision making. The impact on investigative operations is profound.

Traditional investigations are human-driven. Analysts move between dashboards, manually query logs, and piece together telemetry with human-generated context such as tickets, chat threads, and change records. As telemetry volumes and system complexity increase, this model becomes slow and difficult to sustain. AI-driven investigations transform this approach. Agentic systems continuously interrogate telemetry, test multiple hypotheses simultaneously, correlate signals across domains, and enrich findings in real time. Rather than executing a handful of manual queries, AI can run dozens or hundreds instantly, surfacing anomalies, mapping causality, and prioritizing meaningful risk with far greater speed and precision.

However, realizing these benefits requires architectural readiness. AI cannot operate effectively on fragmented, siloed, or unstructured telemetry, it depends on data that is normalized at ingestion, enriched with identity and business context, and accessible at scale without performance or cost constraints. When machine telemetry and human-generated data are fused into a unified, queryable layer, AI gains the context needed to reason effectively

The implication is clear: AI does not replace investigators; it elevates them. Organizations that modernize their telemetry architecture to support agentic workloads gain a decisive operational advantage. They move from reactive analysis to proactive, scalable intelligence, enabling faster, more confident decisions in an increasingly complex threat landscape. This delivers clear advantages:

- Faster detection and resolution: AI compresses mean time to detect (MTTD) and mean time to resolve (MTTR) by automating correlation, enrichment, and hypothesis testing.
- Improved accuracy: Structured data and AI reasoning reduce false positives and surface high-confidence insights quicker.
- Scalable operations: AI agents handle exponential query growth without requiring proportional headcount expansion.
- Human amplification: Analysts shift from manual log hunting to higher-value oversight and decision-making with access to human-based context; from on-call schedules to runbooks and post-incident reviews that can ensure faster response and better investigations.

In the AI era investigative success is no longer defined by how much data you collect. **It's defined by how intelligently you activate it.**

AI will fundamentally reshape investigations

Success depends on modern, scalable data architectures that ensure broad access, high performance search, cost efficiency, and human-guided oversight to fully realize AI's potential.

The Positives

Artificial intelligence accelerates and enhances security and IT investigations. AI agents can run 10–100x more queries than humans, rapidly analyzing logs, metrics, and traces to detect anomalies and correlate signals across systems. With scalable telemetry architecture, AI automates data retrieval, surfaces context, and guides analysts to root causes, reducing detection and resolution time. Rather than replacing investigators, AI amplifies them by enabling smaller teams to move faster with greater analytical depth while preserving human judgment.

The Challenges

Artificial intelligence can introduce new risks in security and IT investigations. AI agents generating 10–100x more queries may overwhelm legacy systems, increase costs, and expose architectural weaknesses. Poor data quality or incomplete telemetry can lead to false positives, missed threats, and misplaced confidence in AI-driven conclusions. Over-reliance on automation may reduce analyst expertise and raise governance concerns. Without strong data foundations and oversight. AI can amplify complexity and operational risk instead of improving outcomes.

The bottom line

AI will transform investigations, but its effectiveness depends on the underlying data architecture. Organizations that modernize their telemetry pipelines to ensure broad data accessibility, elastic scalability, cost efficiency, and high-performance search will enable faster, more accurate, and more collaborative investigations. Those that do not may find that AI accelerates cost, infrastructure strain, and confusion rather than clarity. Ultimately, the determining factor will not be the sophistication of the AI itself, but whether the organization's data foundation is built to support it.

Telemetry pipelines

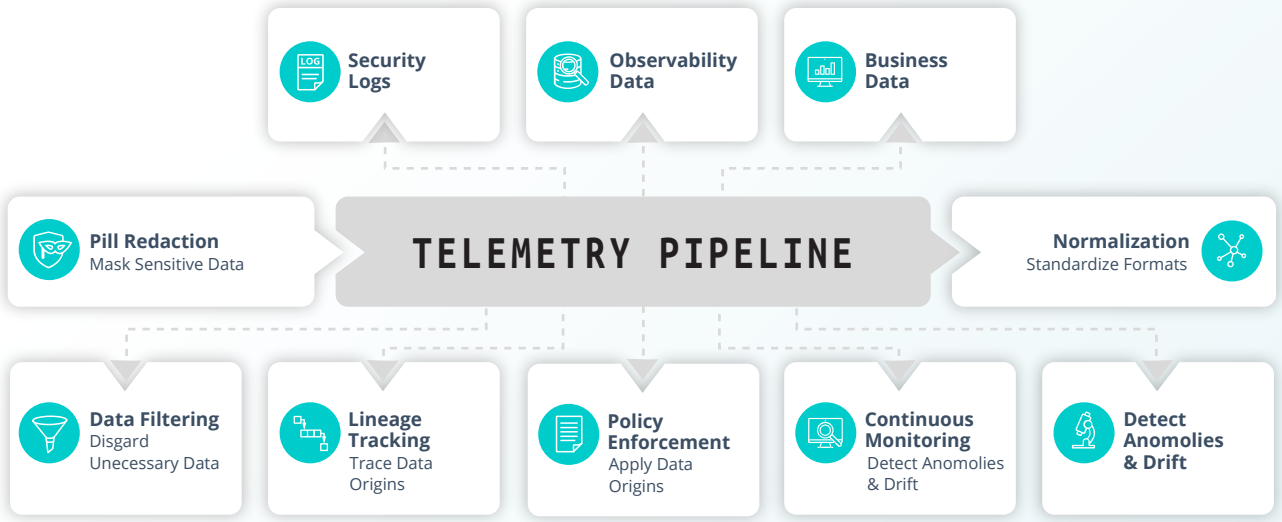

Telemetry pipelines have emerged as a foundational architectural layer in contemporary data ecosystems, serving as the primary front end for ingestion, shaping, and delivery of observability and security telemetry (logs, metrics, traces) to downstream analysis platforms.

In modern distributed environments, this layer quietly controls the essentials: investigative speed, data quality and relevance, ingestion and storage costs, and overall downstream performance. Through normalization, enrichment, filtering, routing, and lifecycle management, pipelines decide what data matters, how it's structured, and where it flows. Many organizations have adopted independent, vendor-agnostic pipelines for their core strengths: transparent and modifiable routing, true multi-destination support, and the freedom to change tools without re-instrumentation or lock-in. That openness was the original value proposition.

Yet as vendors embed pipelines deeper into their own ecosystems, priorities often shift toward proprietary integrations. Routing options narrow, portability declines, and ironically the vendor lock-in pipelines were designed to eliminate begins to return.

In 2026, with telemetry volumes surging, multi-tool stacks becoming standard, and budgets under pressure, the pipeline is no longer just infrastructure. It has become a strategic control point.

Teams that keep this layer open and interoperable, especially via standards, retain real advantages: precise cost control at ingestion, optimal routing of high-value signals, rapid adaptation to new tools, and full governance over sensitive data flows.

The bottom line

Telemetry pipelines are now the center of gravity for observability, security operations, and data driven agility. How organizations manage this layer will increasingly determine how quickly they investigate, how affordably they operate, and how readily they evolve in a constantly shifting landscape.

Engineering a composable analysis platform

The days of data-to-SIEM-to-insights are over

Modern detection and response require defenders to rethink traditional models. Designing a composable analysis architecture is not about abandoning centralized analysis platforms, it's about evolving them into interoperable, automated, and adaptable frameworks capable of keeping pace with cloud scale, AI proliferation, and increasingly sophisticated threats. The shift is moving from monolithic platforms toward modular ecosystems that deliver agility, precision, and automation.

A composable analysis stack deconstructs detection and response into functional layers: data collection, normalization, enrichment, storage, detection engineering, automation, and investigation. Rather than forcing one platform to serve as the entire brain of the operation, each layer is optimized with the right tool and appropriate level of automation. This flexibility allows organizations to automate confidently where maturity is high, yet still maintain oversight where needed.

The trade-off is familiar: platformization offers simplicity but often limits agility and best-of-breed provides specialization but demands integration discipline. A composable approach balances unifying tools through APIs, data pipelines, and orchestration layers. With clear security strategy and governance guiding the design, organizations can achieve scalable observability, detection-as-code, AI-driven automation and federated search. The result is a future-ready analysis ecosystem built for autonomy, resilience, and continuous evolution.

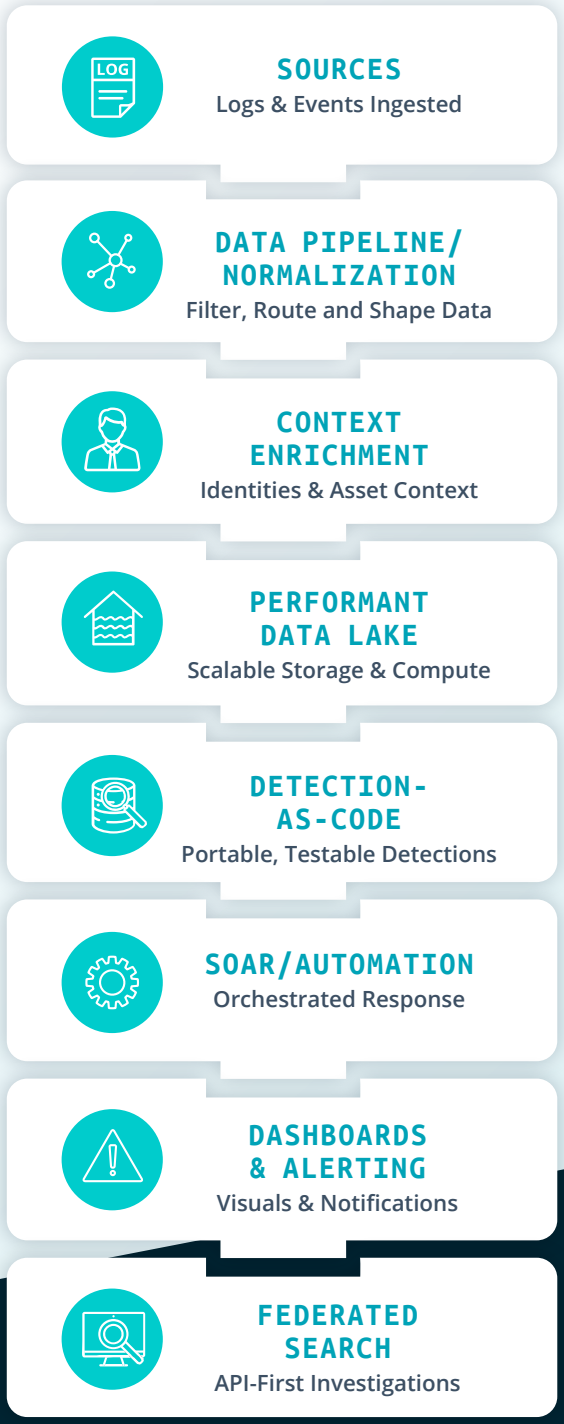

Cribl architecture and capabilities

A modern investigation-ready architecture aligns directly with the Cribl product portfolio, where each foundational layer, collection, pipeline, storage, and analysis is powered by a purpose-built component of the Cribl platform. Rather than stitching together disconnected tools, Cribl delivers an integrated yet modular approach.

Cribl Edge collects and processes telemetry at the source across endpoints, servers, and cloud workloads. By filtering, transforming, and routing data before it leaves the environment, it reduces noise and downstream costs. Built-in buffering ensures resilient delivery during outages, while its lightweight design enables rapid deployment across distributed systems.

Cribl Stream centralizes telemetry processing, normalizing formats, enriching events, and routing data to the right destinations. By decoupling sources from analytics platforms, it provides flexibility and cost control. Replay capabilities allow historical data to be reprocessed without recollection, keeping investigations agile and efficient.

Cribl Lake combines scalable object storage with high-performance analytics in a unified lakehouse model. It supports single ingestion with tiered retention and direct, in-place search. Automated lifecycle policies and governance controls ensure cost optimization, compliance, and long-term, searchable data retention.

Cribl Search delivers high-speed analytics through a unified ingest-to-query architecture with built-in Notebooks for collaborative, interactive investigations. It reduces friction by combining rapid on-boarding, automatic schematization, AI-powered parsing, and distributed search, making telemetry analysis-ready in minutes. Search-in-Place enables teams to query data directly where it resides, including object storage, eliminating rehydration delays and unnecessary movement so both recent and historical data remain accessible. Embedded Notebooks provide a collaborative workspace where queries, results, context, and AI guidance coexist. Investigators can use natural language to explore data, pivot across datasets, document findings, and create repeatable workflows. Agentic AI assists with query generation, summarization, and next-step recommendations, reducing manual effort while keeping analysts in control.

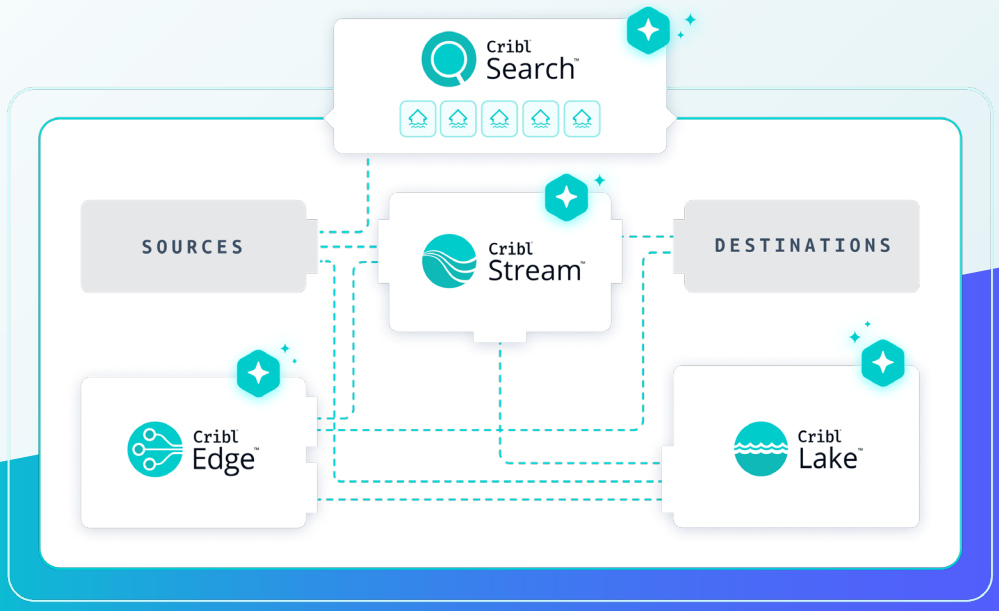

Together, these capabilities unify ingestion, storage, search, and collaboration, accelerating root cause analysis and turning raw telemetry into actionable insight with less complexity.

CONCLUSION

Accelerate investigations with Cribl

Across security, IT, and platform engineering, investigations have become more complex. Telemetry volumes are rising nearly 30% annually as cloud and hybrid environments expand, while AI adoption accelerates. By 2026, agentic SRE and SecOps systems may generate 10–100x more queries than human analysts, overwhelming legacy SIEM and observability platforms not designed for AI-scale workloads. Costs rise, performance limits surface, and teams are expected to operate as "10x investigators" without the infrastructure to support them. A central challenge is fragmented access to investigation-ready data. Telemetry is often restricted by licensing, siloed tools, role-based controls, or ingestion gaps due to cost constraints resulting in incomplete visibility. Analysts pivot between systems, chase logs, and lack of context, extending downtime and operational strain.
The answer is architectural. A modern network architecture structures how systems connect and can fully communicate across the enterprise — creating an observable, resilient ecosystem accessible from a single dashboard. When combined with resilient telemetry collection, intelligent enrichment, scalable storage, and collaborative search, it's where human context meets AI-ready telemetry data and becomes a strategic asset.

Investigating with Cribl

Cribl enables organizations to modernize telemetry architecture for AI-enabled investigations. With Cribl Stream, teams can collect data once, enrich and filter it intelligently, and route it flexibly to reduce unnecessary ingestion costs and remove downstream bottlenecks. Cribl Search delivers high-speed, elastic search across massive telemetry volumes, ensuring both humans and AI agents can access the data they need without license constraints. By unifying pipelines and expanding governed access to telemetry, Cribl transforms data from a fragmented cost center into a scalable foundation for faster, more conclusive investigations. By unifying and optimizing telemetry, organizations enable governed access and high-speed analysis for both humans and AI, driving faster, more decisive investigations at enterprise scale.

Cribl

Cribl, the AI Platform for Telemetry, empowers enterprises to manage and analyze telemetry for both humans and agents. Trusted by organizations worldwide, including half of the Fortune 100, Cribl bridges the gap between AI ambition and infrastructure reality. No lock-in. No data loss. No compromises. Cribl's vendor-agnostic platform ensures data remains portable and interoperable. By cost-effectively handling increasing data volume and variety without delay, Cribl gives enterprises the choice, control, and flexibility to build what's next.

Learn more: cribl.io | Try now: Cribl sandboxes | Join us: Slack community | Follow us: LinkedIn and X (Twitter)

SB-0070-EN-1-0326