WHITE PAPER | 2025

From Multi-Million Dollar Tech Debt to Strategic Value:

The Observability Pivot

How a Lean Team Transformed Enterprise Logging Infrastructure
Across 20 Business Units and 100+ Application Teams

80%+	**~91%**	**40%**	**~80%**
TCO Reduction	Query Performance Gain	MTTR Reduction (P1)	Data Volume Reduction

Hemanth Shivanna
AI Solutions Architect | Platform Engineering & SRE Leader
MBA | M.Sc. Information Technology | AWS, Azure, Salesforce, ITIL Certified

February 2025

Executive Summary

This white paper presents a comprehensive case study of an enterprise-scale observability transformation conducted at a leading automotive remarketing B2B service provider. The initiative addressed a critical accumulation of technical debt within the organization's logging and monitoring infrastructure, specifically a Splunk Enterprise deployment processing 3–5 terabytes of log data daily across 20 business units and over 100 application development teams.

Over a 12–18 month period, a lean team of 4 engineers, operating as a strategic task force rather than a large program, re-architected the entire observability pipeline. The transformation introduced Cribl Stream as an intelligent data routing layer and implemented tiered storage with Amazon S3 based on log retrieval frequency, all deployed and managed through Terraform-based Infrastructure as Code (IaC) on AWS.

The results were transformative: the total cost of ownership was reduced by over 80%, query performance improved by approximately 91% (with observed benchmarks showing queries that previously took 14 minutes 47 seconds completing in 1 minute 22 seconds), and mean time to resolution for Priority 1 incidents decreased by 40%. Daily data ingest volume dropped from 3–5 TB to approximately 700 GB through intelligent filtering, noise reduction, and the establishment of actionable alerting gates, all without any loss of operational visibility. Critically, the project modernized enterprise audit logging compliance throughout the transformation, with DevOps best practices and application development team security audits ensuring PII masking and audit trail integrity in a regulated automotive financial services environment.

This paper details the methodology, architecture, governance framework, and cultural transformation required to achieve these outcomes, providing a replicable blueprint for enterprises facing similar observability cost and performance challenges.

Table of Contents

1. The Challenge: Observability Tech Debt at Enterprise Scale

1.1 Industry Context

The observability market has experienced explosive growth, with enterprise spending on logging, monitoring, and APM platforms increasing 25–40% annually.

However, this growth has often been uncontrolled. Organizations adopt powerful platforms like Splunk, Datadog, and Elastic without establishing the governance frameworks necessary to manage data economics at scale. The result is a pattern now recognized across the industry: **observability tech debt**.

1.2 The Organizational Landscape

The subject organization, a leading automotive remarketing B2B service provider, operated a complex IT ecosystem serving 20 distinct business units with over 100 application development teams. The Splunk Enterprise deployment had grown organically over several years, becoming the de facto logging platform across the enterprise.

Key Operational Characteristics

Dimension	State at Project Initiation
Daily Log Ingest Volume	3–5 TB/day
Business Units Supported	20
Application Teams	100+
Annual Platform Run-Rate	Seven-figure ($1M–$2M+)
Compliance Framework	Enterprise Audit / Regulatory Compliance
Industry	Automotive / Fleet / Financial Services
Retention Policies	Mixed / Inconsistent across BUs
Data Pipeline Architecture	Mixed / Inconsistent across BUs

1.3 Symptoms of Critical Tech Debt

The observability platform exhibited three categories of systemic failure that collectively threatened both operational effectiveness and organizational confidence in the tooling:

Performance Degradation

Search queries routinely exceeded five-minute timeout thresholds. Service owners conducting incident triage, where every minute of delay extends mean time to resolution, were forced to wait 10–15 minutes for diagnostic queries that should have completed in seconds. Critical dashboards used for executive reporting and on-call operations became unreliable, often timing out entirely during peak operational hours.

Engineering Trust Erosion

Application teams across the 20 business units had progressively lost confidence in the platform. Development teams began building shadow logging solutions: exporting logs to local files, standing up ad-hoc ELK stacks, and routing critical telemetry to alternative tools. This fragmentation created blind spots in incident response and compliance coverage.

Economic Unsustainability

The annual run-rate had reached seven figures, compounded by significant professional services fees required to maintain operational stability. Leadership questioned the return on investment of the entire observability stack, and budget conversations increasingly centered on whether the platform should be replaced entirely rather than optimized.

"The Moment of Truth"

> *When a service owner told our team that a critical 15-minute diagnostic query was "just how Splunk works," we knew the problem had moved beyond technology. The organization had normalized failure.*

2. Diagnosis: Understanding the Root Causes

Before designing a solution, the team conducted a comprehensive audit of the existing environment. The diagnosis revealed that the tech debt was not a single problem but a convergence of data, architectural, and cultural failures.

2.1 Data Quality Crisis

A systematic audit of all data flowing into the Splunk environment revealed that approximately 30–40% of ingested data provided zero operational value.

This included verbose debug-level logs from applications that had been decommissioned months or years prior but whose forwarders were never removed; duplicated event streams where multiple collection mechanisms captured the same data; overly chatty health checks and heartbeat logs generating millions of events daily with no alerting or dashboard dependencies; and legacy middleware verbose logging left at DEBUG level in production environments.

2.2 Architectural Entropy

The data pipeline architecture had no consistent design pattern. Across the 20 business units, log data entered the Splunk environment through a patchwork of methods: some teams used Universal Forwarders sending directly to indexers, others routed through Heavy Forwarders with inconsistent configurations, and several teams had implemented custom scripts pushing data via HTTP Event Collector (HEC) endpoints. There was no centralized routing, no filtering layer, and no data enrichment pipeline.

2.3 Governance Vacuum

Retention policies were inconsistent and often excessive. Some indexes retained data for 13 months or longer without a documented compliance justification. The enterprise regulatory compliance requirements mandated specific audit trail retention, but the lack of data classification meant all data, including non-compliance-relevant debug logs, was retained at the same costly tier. There was no data ownership model, no onboarding process for new log sources, and no mechanism to decommission obsolete data streams.

Diagnostic Warning Signs (Applicable to Any Organization)

- Retention policies exceeding 13 months without documented compliance justification
- Search queries consistently hitting 5-minute timeout thresholds
- Platform costs rising 20%+ annually while capabilities remain static
- Shadow logging solutions emerging across engineering teams
- No documented data ownership model for log sources

3. Transformation Framework: The Three-Pillar Approach

The transformation was designed around three interdependent pillars, each addressing a distinct dimension of the tech debt challenge. The framework was deliberately sequenced: data intelligence first, architecture second, culture third. The reasoning is straightforward: you cannot optimize what you do not understand, and you cannot sustain what your teams do not embrace.

Pillar 1: Data Intelligence & Governance

The first pillar established visibility into what data was flowing through the system and whether it provided operational, compliance, or analytical value.

Data Classification Framework

Every data source was categorized into one of four tiers based on operational criticality, compliance requirements, and consumption patterns:

Tier	Classification	Routing Destination	Retention
Tier 1	Mission-Critical / Regulatory Audit	Hot Storage (Splunk Indexes)	Per compliance mandate
Tier 2	Operational / Active Monitoring	Hot Storage (Splunk Indexes)	30–90 days
Tier 3	Historical / Low-Frequency	S3 Tiered Storage (by retrieval frequency)	13 months
Tier 4	Zero-Value / Decommissioned	Filtered at Source (Dropped)	N/A

Pillar 2: Intelligent Routing Architecture

The second pillar introduced a strategic data routing layer between log sources and the Splunk indexing tier. Cribl Stream was selected as the routing platform based on its ability to filter, enrich, transform, and route log data in real-time without requiring changes to source applications.

This architectural insertion point was transformative: for the first time, the organization had a centralized control plane for all observability data flowing into the platform. Data could be inspected, classified, and routed based on content, source, and operational tier, all before incurring indexing costs.

Pillar 3: Cultural & Operational Evolution

The third pillar addressed the human dimension. The organization needed to transition from a "log everything, sort later" culture to a disciplined, governed approach to observability data management.

This required building self-service capabilities that empowered application teams to manage their own log streams within guardrails, implementing automated policies that enforced data classification and retention standards, creating feedback loops so teams could see the cost and performance impact of their logging decisions, and establishing a data governance committee with representatives from each business unit to maintain standards over time.

4. Architecture Deep Dive: Before and After

4.1 Before State: Uncontrolled Organic Growth

The pre-transformation architecture reflected years of organic, uncoordinated growth across the enterprise.

Figure 1: Pre-Transformation Architecture

DATA SOURCES 20 BUs, 100+ Teams	**PIPELINE** No Routing Layer	**SPLUNK PLATFORM** Overloaded & Ungoverned	**BUSINESS IMPACT** Accumulated Tech Debt
• Universal Forwarders (direct) • Heavy Forwarders (inconsistent) • HEC Endpoints (unvalidated) • Legacy Scripts • Shadow Logging • Decommissioned Apps • DEBUG Logs in production	✘ No filtering ✘ No classification ✘ No enrichment ✘ No data governance ✘ No cost controls **ALL DATA TREATED EQUALLY** Debug = regulatory audit trails 3–5 TB/day unfiltered	• Indexer Cluster: 3–5 TB/day ALL hot tier • Mixed Retention: no consistent policy • Query Performance: 5–15 min waits • Annual Run-Rate: $1M–$2M+ • Audit Compliance: at risk • Engineering Trust: eroded	• 30–40% zero-value data waste • 5–15 min avg query wait • $1M–$2M+ annual run-rate • 0 data ownership accountability • 0 self-service capabilities • Growing compliance risk

ROOT CAUSE: Democratized data ingestion without governance created unsustainable cost growth and operational degradation.

Key Deficiencies

There was no intermediate processing layer between sources and indexers. All data was treated equally regardless of operational value, meaning debug logs consumed the same expensive indexing resources as compliance-critical audit trails. Retention policies varied by business unit with no centralized governance, and the indexing cluster was oversized to handle volume rather than right-sized for value.

The absence of a routing layer meant the organization had no ability to inspect, classify, or filter data before it incurred indexing costs. Every byte of data, whether a mission-critical regulatory audit event or a verbose health check from a decommissioned application, consumed the same premium storage and processing resources. This architectural gap was the single largest driver of both cost inflation and performance degradation.

The compounding effect was significant: as data volumes grew, query performance degraded, which eroded engineering trust, which drove teams to build shadow solutions, which further fragmented the observability landscape and increased total organizational cost. Breaking this cycle required addressing all three dimensions simultaneously: data quality, architecture, and culture.

4.2 After State: Intelligent, Tiered, Governed

Figure 2: Post-Transformation Architecture

STANDARDIZED SOURCES	CRIBL STREAM	TIERED STORAGE	GOVERNANCE	OUTCOMES
✓ Universal Forwarders (standardized) ✓ HEC Endpoints (validated) ✓ IaC-Managed Configs ✓ Self-Service Onboarding **20 Business Units** 100+ App Teams Governed Data Pipelines Data Classification Framework **Terraform on AWS**	**Filter** → Drop zero-value **Classify** → Tier 1/2/3/4 **Enrich** → Add context **Route** → By tier + content **TIER 4: DROPPED** 30–40% eliminated EC2 Auto-Scaling	**HOT (Splunk)** Tier 1: Regulatory Audit/Critical Tier 2: Operational (30–90d) ~700 GB/day Sub-second queries **S3 TIERED STORAGE** Tier 3: Historical (13 mo) Audit Archives (compliant) Tiered by retrieval frequency CISO/BU retention policies S3 lifecycle automated	✓ Self-Service Onboarding ✓ Automated Guardrails ✓ Cost Attribution ✓ Data Ownership Model ✓ Retention Governance ✓ Volume Monitoring ✓ DevOps Best Practices ✓ PII Masking (Security Audit) ✓ Regulatory Compliance	**80%+** TCO Reduction **~91%** Query Perf. **40%** MTTR Improvement **~80%** Volume Reduction ✓ Regulatory Compliance ✓ Engineering Trust ✓ Self-Service ✓ PII Masking

INFRASTRUCTURE: AWS Cloud | Terraform IaC | EC2 Auto-Scaling | VPC + IAM | CloudWatch Monitoring | S3 Lifecycle Policies

Architectural Highlights

Cribl Stream serves as the intelligent routing layer, inspecting every event in real-time and making routing decisions based on the data classification framework. Tier 4 data (zero-value) is dropped before it ever reaches Splunk, eliminating 30–40% of ingest volume immediately. Tier 1 and Tier 2 data flows to Splunk indexes for real-time querying with optimized retention windows. Tier 3 data is routed to Amazon S3 tiered storage, organized by log retrieval frequency. Infrastructure logs follow CISO-defined retention and purge policies, while application logs adhere to business requirement policies set by each application development team. S3 lifecycle policies automate the transition between storage tiers and eventual purging based on these defined retention windows.

4.3 Infrastructure as Code: Terraform on AWS

The entire architecture was deployed and managed through Terraform, ensuring reproducibility, version control, and self-service provisioning capabilities:

Component	Terraform-Managed Configuration
Cribl Stream Workers	EC2 fleet with auto-scaling groups, deployed via Terraform modules
Splunk Indexer Cluster	Right-sized cluster with S3-tiered storage configuration

Component	Terraform-Managed Configuration
S3 Storage Tiers	Lifecycle policies by retrieval frequency, encryption, access controls, and retention automation per CISO/BU policy
Network & Security	VPC, security groups, IAM roles, and cross-account access policies
Monitoring & Alerting	CloudWatch integration for infrastructure health and capacity planning

The IaC approach was critical to the self-service model: application teams could request new log stream onboarding through standardized Terraform modules, reducing provisioning time from weeks to hours while ensuring compliance with governance standards.

5. Implementation Methodology

The 12–18 month transformation was executed in four phases, each building upon the prior phase's deliverables. The lean team structure (4 engineers) was a deliberate design choice. Smaller teams move faster, communicate more efficiently, and maintain clearer accountability.

Phase 1: Discovery & Assessment (Months 1–3)

The first phase focused on comprehensive environmental assessment and stakeholder alignment. The team conducted a full audit of all data sources, ingest pipelines, and consumption patterns across all 20 business units. This included mapping every Universal Forwarder, Heavy Forwarder, and HEC endpoint; cataloging every index, sourcetype, and retention configuration; interviewing service owners and SRE teams across business units to understand operational dependencies; and documenting enterprise audit and regulatory compliance requirements with the internal audit and legal teams.

Phase 2: Architecture & Proof of Concept (Months 3–6)

Phase 2 introduced the Cribl Stream routing layer in a controlled pilot with two business units representing high-volume, well-understood data streams. The proof of concept validated the data classification framework, demonstrated the filtering and routing capabilities in production conditions, established performance baselines for the new architecture, and confirmed that regulatory audit trail integrity was maintained through the routing layer.

Phase 3: Phased Rollout (Months 6–14)

The production rollout was executed in waves, prioritizing business units by data volume and operational criticality. Each wave followed a standardized playbook: classify data sources against the tiering framework, implement Cribl Stream routes and filters, validate compliance data integrity, cutover to the new pipeline with rollback capability, and decommission legacy direct-to-indexer configurations.

Phase 4: Optimization & Self-Service (Months 14–18)

The final phase focused on building sustainable operational capabilities, including the self-service Terraform modules for application teams, automated guardrails for data classification enforcement, dashboards providing business units visibility into their own data economics, and the governance committee structure for ongoing standards management.

5.1 Implementation Challenges: Resistance, Risk, and Reality

No enterprise transformation of this scale proceeds without friction, and this initiative was no exception. Transparency about the challenges encountered is essential for any organization considering a similar undertaking.

Organizational Resistance to Change

The most significant challenge was not technical; it was human. Change is hard to accept, and across 20 business units with deeply entrenched workflows, the pushback was substantial. Teams had spent years building familiarity with existing pipelines, and the proposed transformation was met with a recurring and pointed concern: *"Are we just walking from one tech debt to another with a different name?"*

This was a legitimate question, and dismissing it would have been a strategic mistake. Instead, the team addressed it head-on by establishing transparent metrics and demonstrating incremental value through the pilot phase. Each wave of the rollout included before-and-after performance comparisons shared with the affected business units, allowing teams to see tangible proof that the new architecture was delivering real improvements rather than merely shifting complexity. The data governance framework, with clear ownership models, cost attribution, and self-service capabilities, was deliberately designed to answer this concern structurally: unlike the previous state, every decision in the new architecture was documented, measurable, and reversible.

Migration Risk and Rollback Planning

The risk of data loss or compliance gaps during migration was a constant concern, particularly in a regulated environment where audit trail integrity and defensible regulatory reporting are non-negotiable. The team mitigated this through well-documented self-service rollout procedures that enabled each business unit to migrate from the legacy Splunk direct-ingest model to the Cribl Stream routing layer at their own pace. Critically, every migration wave included full rollback capability. If any data integrity issue surfaced, the team could revert to the previous pipeline configuration within hours.

Support engagement was automated through a tiered support model: Tier 1 issues (configuration questions, onboarding assistance) were handled through self-service documentation and automated Terraform modules; Tier 2 issues (pipeline anomalies, routing misconfigurations) were escalated to the core engineering team; and Tier 3 issues (compliance concerns, architectural decisions) were escalated to the project lead and governance committee. This tiered approach ensured that the lean team of 4 engineers was not overwhelmed by support requests while maintaining high responsiveness for critical issues.

6. Business Impact and Results

The transformation delivered measurable, auditable business impact across four dimensions:

80%+	~91%	40%	~80%
Total Cost of Ownership Reduction	Query Performance Improvement	MTTR Reduction (P1 Incidents)	Daily Ingest Volume Reduction

6.1 Financial Impact

The annual platform run-rate was reduced by over 80%. This was achieved through the elimination of 30–40% of zero-value data at the Cribl Stream routing layer (reducing Splunk license consumption), right-sizing the indexer cluster to match actual operational data volumes, moving long-term retention to cost-effective S3 tiered storage organized by log retrieval frequency with automated lifecycle policies, and significantly reducing professional services dependency through self-service capabilities and IaC automation.

6.2 Operational Performance

Query performance improved by approximately 91%. In a representative benchmark, a critical diagnostic query that previously took 14 minutes and 47 seconds completed in 1 minute and 22 seconds post-transformation. It is important to note that the pre-transformation baseline reflected a genuinely degraded system. Years of accumulated tech debt had pushed query performance well beyond normal operating parameters. The improvement therefore represents both the recovery from a pathological state and the architectural gains from the new tiered approach.

This improvement was driven by dramatically reduced index sizes (less data to search after eliminating zero-value ingest), optimized data models and search-time field extractions, S3 tiered storage offloading historical data from expensive hot indexes, and the elimination of resource contention from zero-value data processing.

6.3 Data Volume Reduction

Total daily ingestion was reduced from 3–5 TB per day to approximately 700 GB, a reduction of roughly 80%. This was not achieved by indiscriminately dropping data, but through a disciplined approach rooted in a single operating principle: **cut down noise, retain signal, and establish actionable alerts**.

The Cribl Stream routing layer served as the enforcement point for this principle. Every data stream was evaluated against the tiering framework: did this log source contribute to incident detection, compliance requirements, or operational decision-making? If not, it was either filtered, sampled, or dropped entirely.

Critically, the team established an ongoing monitoring gate, a recurring review process to ensure that data hygiene standards were maintained over time and that new application onboarding adhered to the classification framework. This gate included automated volume anomaly detection, quarterly data classification reviews with business unit representatives, and self-service dashboards showing each team’s ingest patterns against their allocated thresholds. The gate transformed noise reduction from a one-time project into a sustainable operational discipline.

6.4 Incident Response

Mean Time to Resolution for Priority 1 incidents decreased by 40%. This compound improvement resulted from faster query performance (engineers could iterate on diagnostic hypotheses rapidly), higher data quality (relevant signals were no longer buried in noise), restored trust in the platform (teams stopped using shadow logging solutions and returned to the centralized platform), and improved alert fidelity (reduced false positives from noisy, low-value data sources).

"The Moment That Defined Success"

> *The true metric of success was not the spreadsheet. It was the relief on a service owner’s face when a critical diagnostic query that used to take nearly 15 minutes finished in under 90 seconds. That moment, when an engineer realizes their tooling works for them rather than against them, is when technical leadership compounds into organizational trust.*

7. Governance and Self-Service Model

7.1 The Self-Service Framework

A core design principle was that the transformation had to be sustainable without the original project team. The self-service model empowered application teams across all 20 business units to manage their own observability data within established guardrails.

Self-Service Capabilities

Capability	Mechanism	Governance Guardrail
New Log Source Onboarding	Terraform Module Request	Auto-classified per tiering framework
Retention Policy Changes	Governed Configuration	Must cite compliance justification
Data Volume Monitoring	Self-Service Dashboard	Alerts at 80% of allocated budget
Pipeline Modifications	Cribl Stream Routes	Peer-reviewed, IaC version-controlled
Index Creation	Standardized Request Process	Naming conventions, lifecycle policies enforced

7.2 Automated Guardrails

Governance was enforced through automation, not manual review. Automated policies included data classification validation at the Cribl Stream routing layer (unclassified data was quarantined, not dropped, to prevent data loss), retention policy enforcement through Terraform-managed index lifecycle configurations, volume anomaly detection alerting data owners when ingest patterns deviated significantly from baselines, and cost attribution dashboards providing each business unit transparency into their platform consumption.

7.3 Compliance and Security Continuity

Throughout the transformation, enterprise audit compliance was maintained as a non-negotiable constraint, with the overarching goal of strengthening internal control reliability, reducing compliance risk, and enabling defensible regulatory reporting. Rather than relying solely on cryptographic mechanisms, the team adopted a comprehensive DevOps best practices approach to compliance assurance, integrated with the application development team’s existing security audit requirements.

Specific measures included PII masking enforced at the Cribl Stream routing layer, ensuring that personally identifiable information was redacted before data reached any storage tier, a critical requirement validated through each application development team’s security audit process. Documented chain-of-custody for all data classification and retention decisions ensured auditability. Automated compliance reporting demonstrated data lifecycle adherence across all tiers. Quarterly audit reviews with internal audit and legal teams validated the new architecture’s compliance posture. The security audit framework required each application team to certify that their data pipelines met PII masking standards before production deployment, creating a distributed accountability model that scaled across all 20 business units.

8. Lessons Learned and Recommendations

8.1 Strategic Lessons

Start with Data Intelligence, Not Architecture

The temptation in observability optimization is to lead with tooling decisions. Our experience demonstrated that understanding your data (classifying it, quantifying its value, and mapping its consumption) must precede any architectural changes. The data audit in Phase 1 revealed that 30–40% of our ingest was zero-value, a finding that fundamentally shaped every subsequent architectural decision.

Lean Teams Outperform Large Programs

A team of 4 engineers transformed observability infrastructure serving 20 business units and 100+ application teams. The lean structure enabled rapid decision-making, clear accountability, and direct communication with stakeholders. Large transformation programs often spend more time coordinating than executing.

Culture Change Requires Self-Service, Not Mandates

Top-down mandates to "log less" are ineffective and breed resentment. The self-service model with transparent cost attribution created natural incentives for data hygiene. When teams can see the cost of their logging decisions and have the tools to optimize independently, behavior changes organically.

Acknowledge and Address Resistance Directly

The most valuable lesson from this transformation was that resistance is not an obstacle, it is signal. When teams pushed back with concerns about trading one form of tech debt for another, that pushback forced the team to build stronger governance, clearer metrics, and more transparent decision-making processes. The resulting architecture was better because of the scrutiny, not despite it.

8.2 Tactical Recommendations

For Organizations Using Splunk

Begin with the Monitoring Console App to audit instance health, indexing rates, and search performance baselines. Identify your top 10 sourcetypes by volume and assess their operational value. Evaluate S3 tiered storage for any retention requirements exceeding 30 days. Consider a data routing layer (Cribl Stream or equivalent) as the single highest-ROI architectural investment.

For Organizations Using Datadog, Elastic, or Other Platforms

The principles in this paper are platform-agnostic. Every observability platform faces the same fundamental challenge: democratized data ingestion without governance leads to unsustainable cost growth. The three-pillar framework (Data Intelligence, Intelligent Routing, Cultural Evolution) applies regardless of vendor.

Red Flags: Is Your Organization Accumulating Observability Debt?

- Retention policies exceeding 13 months without a documented compliance "why"
- Search/query performance consistently degrading quarter over quarter
- Platform costs rising 20%+ annually while operational capabilities remain flat
- Engineering teams building shadow logging or monitoring solutions
- No data ownership model or cost attribution by business unit
- Professional services fees exceeding 15% of total platform costs

9. Industry Applicability and Future Outlook

9.1 Cross-Industry Relevance

While this case study originated in the automotive remarketing and fleet services sector, the patterns, challenges, and solutions described are universally applicable across regulated industries. Financial services organizations facing similar regulatory audit, PCI-DSS, and retention requirements will find the governance framework directly transferable. Healthcare organizations managing HIPAA-compliant observability data face analogous data classification challenges. Technology companies experiencing hypergrowth in log volumes share the same economic pressures regardless of their observability platform choice.

9.2 The 2025–2026 Observability Landscape

The observability industry is converging on several trends that reinforce the approach described in this paper:

AI/ML-Driven Data Intelligence

Machine learning models are increasingly being applied to automated data classification, anomaly detection in ingest patterns, and predictive cost optimization. The data governance framework established in this transformation provides the labeled, classified dataset required to train these models effectively.

OpenTelemetry and Vendor Neutrality

The rise of OpenTelemetry as a vendor-neutral observability standard reinforces the value of an intelligent routing layer. Organizations that decouple data collection from data storage (as this architecture does with Cribl Stream) are positioned to adopt multi-destination strategies without re-instrumenting applications.

FinOps for Observability

The emergence of FinOps practices specifically for observability spending mirrors this paper’s emphasis on cost attribution, data economics, and business unit accountability. The self-service governance model described here is an early implementation of what is becoming an industry-standard approach.

10. About the Author

Hemanth Shivanna

AI Solutions Architect | Platform Engineering & SRE Leader

24+ Years of IT Strategy, Infrastructure, and Operational Excellence

Hemanth Shivanna is an AI Solutions Architect at Elite Technology Solutions with over 24 years of experience spanning IT strategy, platform engineering, Site Reliability Engineering (SRE), and managed services operations. He previously served as Senior Manager of IT & Managed Services Operations at Openlane Inc. (formerly KAR Global), where he led the observability transformation described in this white paper.

Hemanth’s career has been defined by a "player-coach" approach to technical leadership, combining hands-on architectural expertise with the strategic vision to align infrastructure investments with business outcomes. His track record includes $2.2M+ in documented annual cost savings, 90% reduction in production outages, 67% faster project delivery cycles, and leadership of globally distributed teams of 35+ engineers across regulated environments.

Education

MBA (Master of Business Administration)
M.Sc., Information Technology / Information Systems
B.S., Information Technology

Professional Certifications

Certification Domain	Active Credentials
Amazon Web Services (AWS)	Solutions Architect, Cloud Practitioner, AI Practitioner
Microsoft Azure	Azure Fundamentals, Azure AI Fundamentals
Salesforce	Agentic AI Specialist, AI Associate, AI Specialist
ITIL	ITIL v4 Foundation
Microsoft	MS-900, Azure Administration
Cisco	CCNA (Routing & Switching)

Areas of Expertise

Cloud Infrastructure & Architecture (AWS, Azure), Observability & Monitoring (Splunk, Cribl, LogicMonitor, Datadog), Site Reliability Engineering (SRE), Infrastructure as Code (Terraform, Ansible), AI/ML Integration & Strategy, IT Service Management (ITSM/ITIL), Enterprise Security & Compliance (Enterprise Audit, PCI-DSS), Team Leadership & Organizational Development.

Connect

LinkedIn: linkedin.com/in/hemanthshivanna
Website: hemanthshivanna.com

Disclaimer: This white paper reflects the author's professional experience. Company-specific details have been anonymized. Methodologies and recommendations are the author's own. Quantitative results are based on documented outcomes and generalized where necessary.