AI-Powered Observability Across Mainframes and Cloud-Native Architectures: A Unified Approach to Hybrid IT Operations

Madhu Garimilla

Broadcom
United States

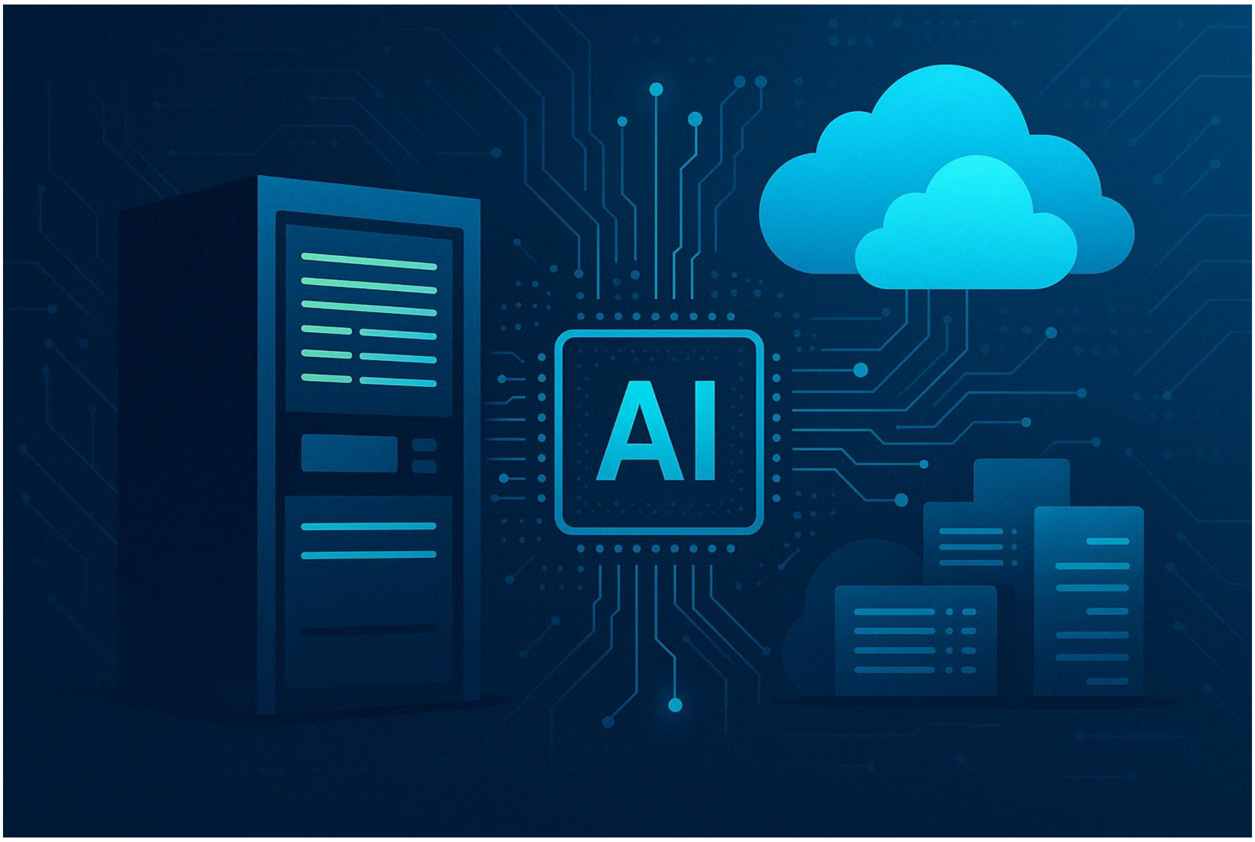

Abstract:
The increasing complexity of hybrid IT ecosystems spanning both legacy mainframes and modern cloud-native systems has heightened the demand for advanced observability. Conventional monitoring techniques fail to offer integrated insights across such heterogeneous platforms, often resulting in operational blind spots. This paper presents a unified observability framework augmented with artificial intelligence (AI), which enables organizations to achieve real-time visibility, intelligent alerting, and proactive issue resolution. By examining architectural design, implementation strategies, and practical examples, this article demonstrates how AI can facilitate seamless operations across hybrid environments while reducing operational costs, improving mean time to resolution (MTTR), and boosting system reliability.

Keywords: AI Powered Observability, Mainframes, Cloud Native Architectures, Hybrid IT Operations, Unified Observability, AIOPs.

1. INTRODUCTION AND MOTIVATION

Enterprise IT environments are undergoing rapid transformation, with organizations striving to modernize infrastructure while maintaining the reliability of legacy systems. Mainframes remain vital for critical applications in sectors such as finance, government, and healthcare. According to IDC, over 71% of Fortune 500 companies still rely on mainframe systems for core transactional processes due to their security, scalability, and uptime [1]. for core transactional processes due to their security, scalability, and uptime. Conversely, 90% of new enterprise applications are cloud-native, driven by demands for agility and distributed development [4].

This duality creates complex operational environments that generate vast volumes of telemetry data in heterogeneous formats. Mainframes emit SMF records, SYSLOG data, and job logs, while cloud-native systems produce distributed traces, time-series metrics, and structured/unstructured logs via platforms like Open Telemetry and Prometheus. The divergence in data structures, ownership, and tooling introduces friction in achieving end-to-end visibility. A 2023 Splunk report found that 68% of organizations experience delayed incident resolution due to siloed observability tools [2]. due to siloed observability tools.

This paper demonstrates that a unified observability approach powered by AI—capable of ingesting, normalizing, correlating, and analyzing data across diverse systems—offers a strategic solution [5].—capable of ingesting, normalizing, correlating, and analyzing data across diverse systems—offers a strategic solution. When implemented effectively, this approach provides contextualized insights, facilitates early anomaly detection, and enables autonomous remediation.

2. UNIFIED OBSERVABILITY FRAMEWORK

The unified observability framework is structured into four key layers:

- **Data Aggregation Layer**: This layer collects telemetry from heterogeneous systems. On the mainframe side, tools like SYSVIEW, OMEGAMON or homegrown log shippers can stream SMF and SYSLOG data. On the cloud-native side, collectors such as Fluentd, Logstash, and OpenTelemetry agents capture real-time metrics and traces. API gateways, edge services, and sidecars enhance observability for ephemeral workloads. Unified data lakes or observability pipelines (e.g., Kafka + Elasticsearch) serve as a backbone for cross-domain telemetry.
- **Normalization and Correlation Layer**: This layer transforms platform-specific data into a standardized schema using techniques like schema-on-write, tagging, and timeline alignment. Metadata such as trace IDs, transaction keys, user sessions, and geographic labels help create a unified context. Event correlation is achieved using rule-based logic, statistical similarity, or AI clustering to detect relationships between events.
- **AI Analytics Layer**: At the core of intelligent observability, this layer uses machine learning algorithms to identify anomalies, forecast failures, and trace root causes. Time-series forecasting with Prophet or ARIMA models, anomaly detection via Isolation Forests or autoencoders, and causal inference with Bayesian networks are commonly applied. Gartner forecasts that by 2026, over 60% of enterprises will integrate AIOps capabilities into their observability stacks [3]. into their observability stacks.
- **Visualization and Control Layer**: Dashboards, alert consoles, and ITSM integrations allow teams to act on insights. Platforms such as Grafana, Kibana, and Datadog provide rich UIs, while incident response tools (e.g., PagerDuty, ServiceNow) enable automated ticketing and escalation. The key is reducing alert fatigue and delivering actionable information.

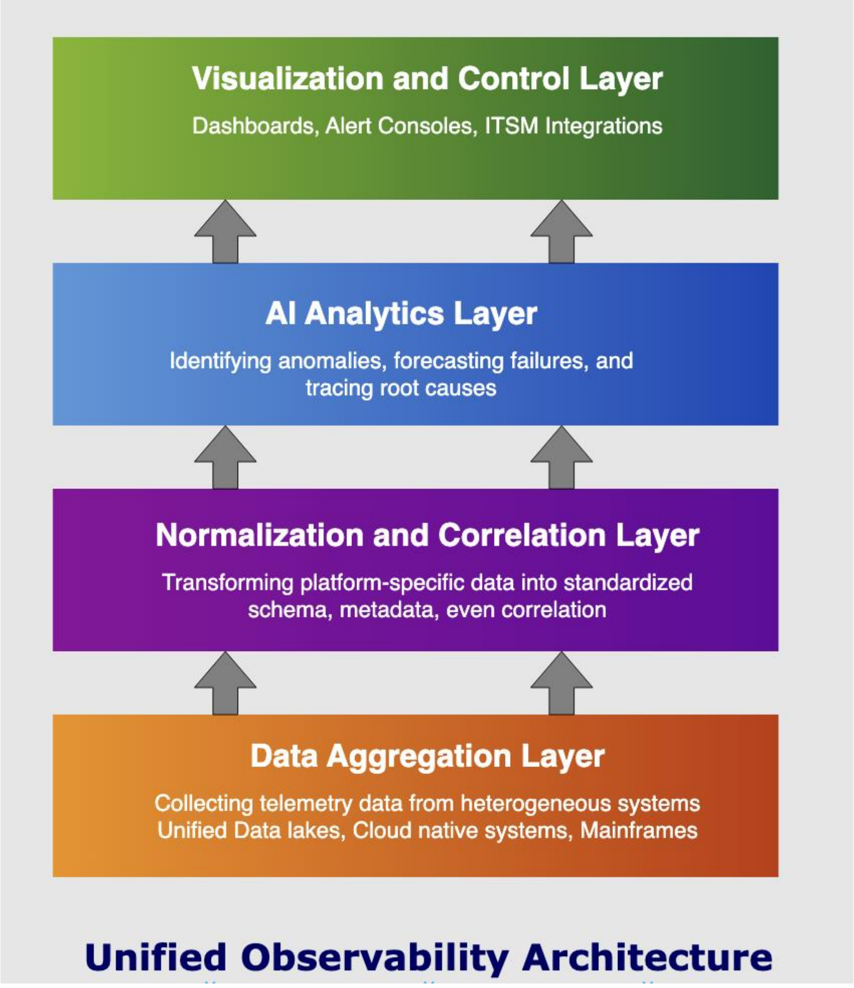
Visualization and Control Layer
Dashboards, Alert Consoles, ITSM Integrations
AI Analytics Layer
Identifying anomalies, forecasting failures, and tracing root causes
Normalization and Correlation Layer
Transforming platform-specific data into standardized schema, metadata, even correlation
Data Aggregation Layer
Collecting telemetry data from heterogeneous systems
Unified Data lakes, Cloud native systems, Mainframes
Unified Observability Architecture

Figure 1: Unified Observability Architecture *Visualizing the interaction between layers in a unified observability framework.*

3. IMPLEMENTATION STRATEGY AND CHALLENGES

Successful adoption of AI-powered observability involves aligning strategy with business goals and operational maturity. Organizations often face a paradox: they need observability to improve stability, yet lack the visibility to make informed decisions on where to begin [6]. To bridge this, a strategic blueprint should consider organizational structure, tool maturity, and data governance policies.

In addition to the core phases of centralization, standardization, intelligent analytics, and automation, enterprises can adopt targeted sub-strategies:

- **Pilot-Based Implementation**: Start with a single business-critical application that spans mainframe and cloud-native services. Use it to demonstrate value, identify integration issues, and refine data schemas.
- **Observability Readiness Assessment**: Conduct an audit of current telemetry sources, data quality, alert rules, and response processes. Identify blind spots and redundant tools. This assessment guides prioritization and architecture choices.
- **Persona-Driven Dashboards**: Tailor visualization layers to roles—SREs, operations staff, business owners—using widgets that map technical metrics to business KPIs (e.g., revenue per transaction latency, SLA breaches per geography).
- **Feedback Loop for Continuous Improvement**: Use AI-driven retrospectives to capture learnings from incidents. Feed insights back into observability models to refine alert thresholds and root cause predictors.

For example, a global logistics provider implemented an observability strategy beginning with their shipment tracking system, which integrated CICS transactions on z/OS and containerized frontend dashboards. By gradually layering AI insights, they reduced false-positive alerts by 60% and improved delivery forecast accuracy by 22%.

These detailed strategies foster adoption that is both systematic and adaptable, ensuring long-term impact and stakeholder alignment.

Implementing AI-powered observability is not a one-size-fits-all process. Organizations should tailor their adoption path based on maturity, existing tooling, and strategic priorities. To operationalize AI-powered observability, organizations can follow a structured approach that gradually matures their capabilities while ensuring alignment with business outcomes. The strategy typically unfolds through four progressive stages, each building on the last to deliver increasing levels of automation and intelligence:

1. **Centralization – Establishing a Unified Telemetry Backbone**: The first step involves consolidating data sources from mainframes, cloud-native platforms, and third-party systems into a central telemetry platform or data lake. This includes deploying log forwarders, setting up metrics collectors (e.g., Prometheus scrapers), and configuring span exporters (e.g., OpenTelemetry). This foundational stage enables visibility but may still rely on manual correlation.
2. **Standardization – Structuring Data for Cross-System Correlation**: Once telemetry is centralized, the focus shifts to creating a consistent schema and enriching data with meaningful tags (e.g., service name, region, customer segment). This phase improves data quality and enables traceability across systems. Organizations should adopt semantic conventions (like OpenTelemetry guidelines) and define SLIs/SLOs to guide observability goals.
3. **Intelligent Analytics – Layering AI for Insights and Forecasting**: In this stage, machine learning models are introduced to enhance anomaly detection, forecast usage patterns, and prioritize incidents. Common examples include applying Isolation Forests to detect memory spikes, using Prophet to forecast storage saturation, or building causal graphs to trace the root cause of latency. Integration with CMDB (Configuration Management Database) can further enhance the model's contextual understanding [7].
4. **Closed-Loop Automation – Driving Proactive Remediation**: The final stage operationalizes observability insights by feeding them into CI/CD pipelines, ITSM systems, and auto-remediation scripts. For example, AI can trigger rollbacks during canary deployment failures or automatically adjust Kubernetes HPA (Horizontal Pod Autoscaler) policies in response to predicted load. This stage reduces human intervention and improves incident response time.

Each stage should be supported by training, tooling assessments, and stakeholder engagement to ensure adoption. Organizations that adopt this phased model are better positioned to scale observability without overwhelming teams or compromising on compliance.

Key Implementation Challenges:

- **Data Volume and Cardinality**: High cardinality tags (e.g., user IDs) can lead to performance issues and inflated storage costs. Teams must balance granularity with cost efficiency using aggregation strategies.
- **Compliance and Privacy**: Sensitive data embedded in logs must be redacted or tokenized. Role-based access control (RBAC) and encryption-in-transit/storage are critical.
- **Organizational Change Management**: Establishing cross-functional observability squads can help unify mainframe operations teams with DevOps/SRE groups, promoting shared responsibility.
- **Toolchain Integration**: Leveraging open APIs and service meshes (e.g., Istio) facilitates consistent observability practices across diverse workloads.

According to a 2024 Forrester study, organizations with mature observability practices witnessed a 43% drop in unplanned outages and saved an average of $2.1 million annually through improved incident prevention [10]. witnessed a 43% drop in unplanned outages and saved an average of $2.1 million annually through improved incident prevention.

4. USE CASES IN HYBRID IT OPERATIONS:

AI-powered observability provides tangible benefits across industries. These extended use cases illustrate deeper integrations:

- **Financial Services**: A multinational investment bank implemented unified observability across its COBOL trading engine and Kubernetes-based reporting service. AI correlated memory spikes on z/OS with concurrent API failures, enabling dynamic traffic rerouting via service mesh and reducing trade delays by 80%.
- **Retail Sector**: A leading online retailer used unsupervised learning to detect early signs of infrastructure strain during seasonal sales. Combined log and metric analysis prevented order drops by alerting operators to memory saturation in backend processes. Integrating these alerts into Grafana dashboards ensured 24/7 visibility and team collaboration.
- **Healthcare Systems**: A public healthcare consortium used a supervised anomaly detection model to monitor patient data ingestion from mainframe sources into cloud-hosted EHR systems. Model retraining based on drift patterns helped maintain over 95% accuracy in alert generation, directly contributing to reduced reporting errors and improved clinical decision support.

Cross-Domain Impact and KPIs:

Table 1: Observability Impact Metrics Across Industries

Industry	Use Case	MTTR Reduction	Alert Volume Reduction	Customer Impact
Banking	DB2 throughput correlation	35%	30%	Improved SLA compliance
Retail	ERP-API flow optimization	50%	45%	Revenue loss prevented
Healthcare	Predictive batch processing alerts	60%	40%	Increased reporting accuracy

- SLA adherence improved by 30% with proactive scaling.
- MTTR dropped from 3 hours to 40 minutes in retail scenarios.
- Alert fatigue reduced by over 50% through deduplication and intelligent alert suppression [8].

5. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

AI-powered observability is transforming IT operations by converging insights across platforms and surfacing root causes with unprecedented speed and precision. Beyond technical gains, this paradigm fosters collaboration, resilience, and data-driven decision-making. As the ecosystem evolves, several future directions emerge:

- **Agentic AI**: Future observability systems will employ AI agents with situational awareness, capable of performing complex, multi-step remediation tasks based on real-time telemetry and historical behavior models.
- **eBPF & WASM Instrumentation**: Innovations like eBPF allow low-overhead monitoring directly from the Linux kernel, while WASM modules can provide portable observability logic that adapts dynamically to runtime environments.
- **Explainable AI (XAI)**: With increasing adoption of AI in operational decision-making, the need for transparency becomes critical. Visual tools that map model decisions to system components will enhance trust and accountability.
- **Industry-Specific Observability Blueprints**: Financial, healthcare, and government sectors will benefit from pre-defined observability templates aligned with compliance standards and operational goals.

Ultimately, unified observability is not merely a technological upgrade—it is an enabler of digital trust, operational agility, and sustained innovation.

REFERENCES:

1. IDC, "Mainframe Modernization Survey", 2023.
2. Splunk, "State of Observability 2023".
3. Gartner, "Emerging Technologies: AIOps", 2023.
4. EMA Research, "The Impact of AIOps on IT Operations", 2023.
5. TechRepublic, "Skills Gaps in AI and Observability", 2022.
6. Dynatrace, "State of Observability 2024".
7. OpenTelemetry Project Documentation.
8. IBM Redbooks, "z/OS Observability Techniques", 2023.
9. Google Research, "Explainable Machine Learning Techniques", 2022.
10. Forrester, "State of IT Operations and Observability", 2024.