



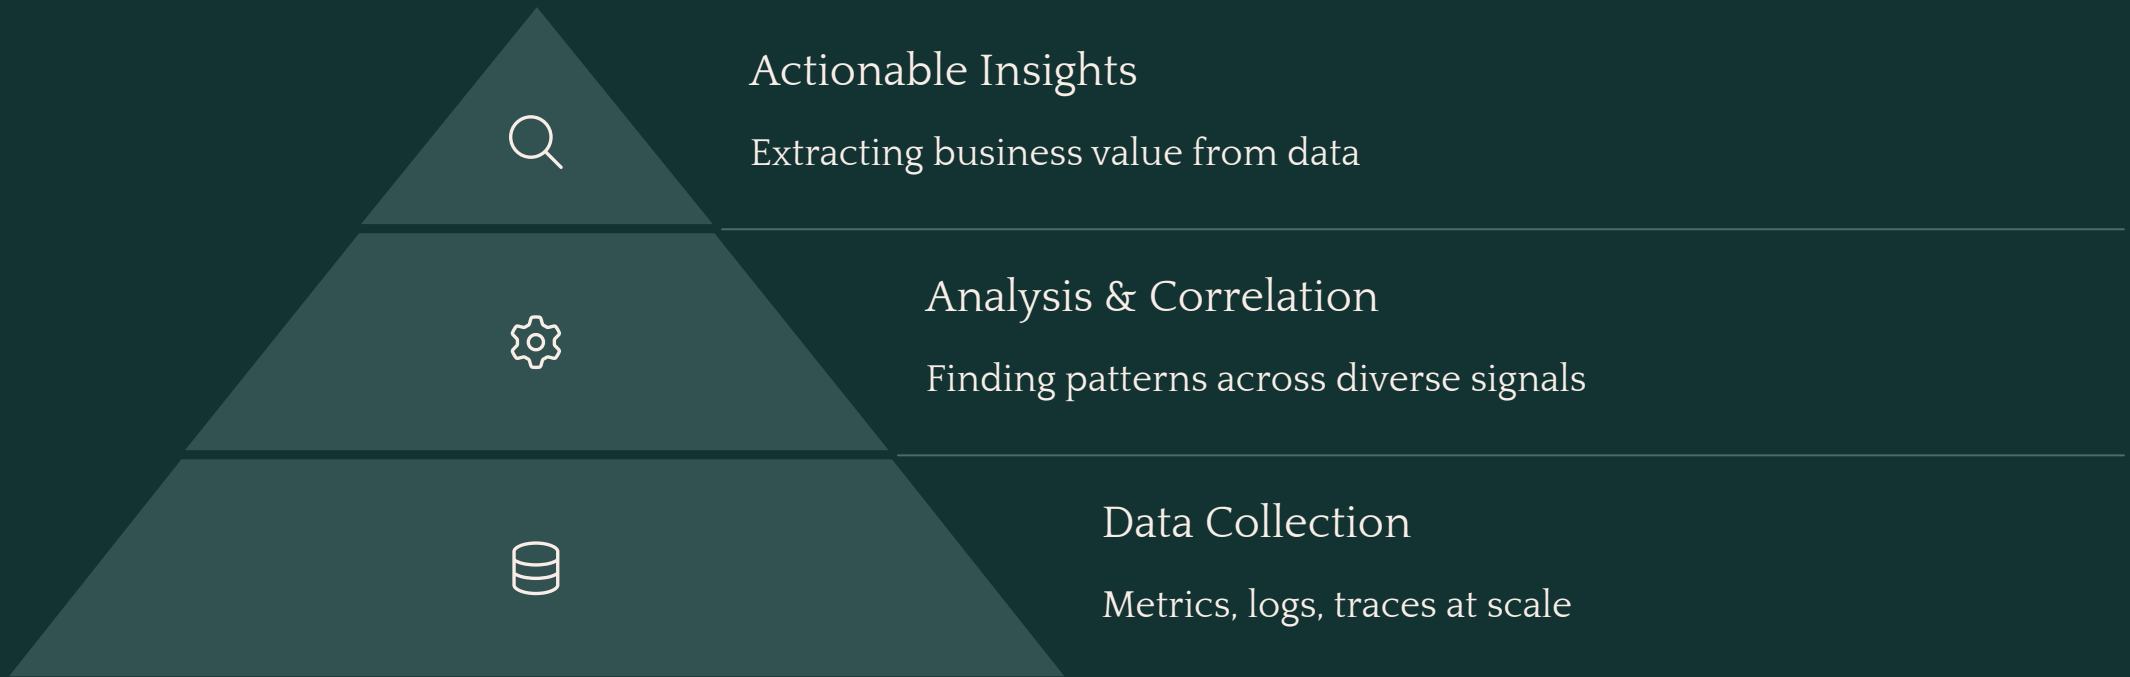
# Beyond Reactive Monitoring: Implementing AI-Powered Observability with Splunk for Sub-Minute Incident Resolution

Modern distributed systems generate overwhelming volumes of metrics, traces, and logs across microservice architectures, making traditional monitoring approaches inadequate. This presentation explores how Splunk's AI capabilities can transform observability practices from reactive firefighting to predictive system management.

Drawing from the latest Observability Forecast findings, we'll examine how leading organizations position observability as a mission-critical capability, allowing teams to maintain high reliability even during peak traffic periods.

By: Prabhu Govindasamy Varadaraj

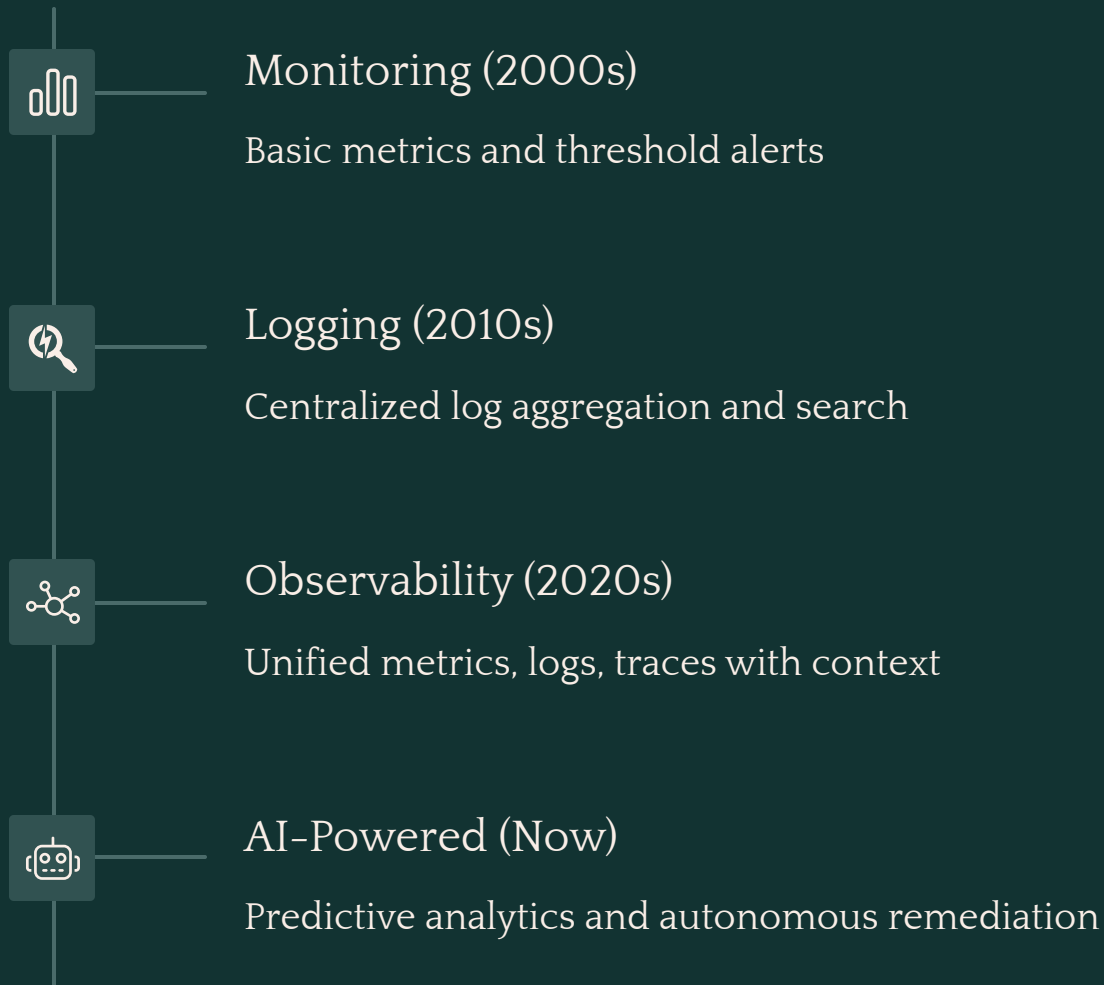
# The Observability Challenge



The complexity of modern distributed systems has created an observability crisis. IT teams struggle with siloed monitoring tools that generate thousands of alerts without context, leading to alert fatigue and delayed responses.

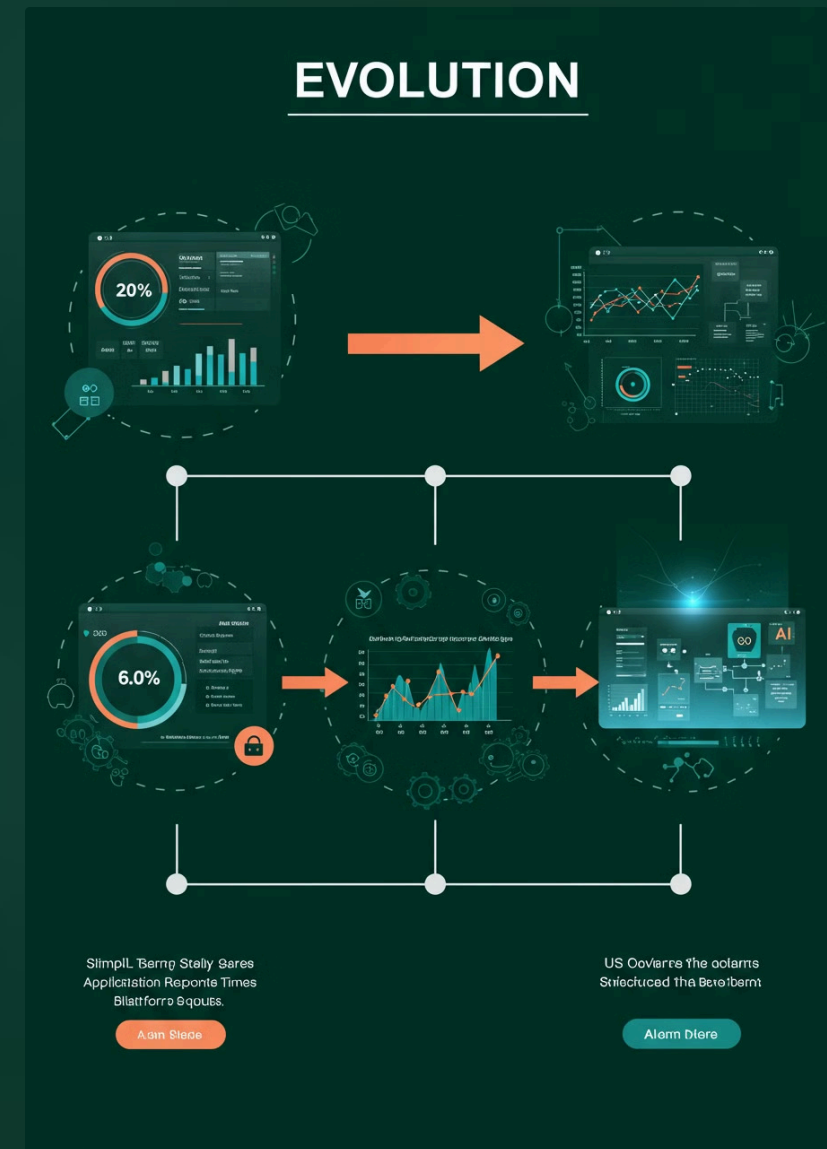
Organizations face challenges at each level of the observability pyramid - from collecting the right data to deriving actionable insights that drive business decisions. Without a unified approach, teams spend more time troubleshooting than innovating.

# The Evolution of Observability



Observability has evolved dramatically over the past two decades. What began as simple uptime monitoring has transformed into sophisticated systems that provide deep visibility into complex distributed architectures.

The latest evolution integrates artificial intelligence, enabling systems to not only report on what's happening but predict future states and automate responses. This shift represents a fundamental change in how organizations approach system reliability and performance management.





# Observability as a Strategic Imperative

94%

Digital Importance

Percentage of enterprises that consider digital services critical to business

54%

Observability Gap

Organizations lacking full-stack observability

45%

MTTR Reduction

Average improvement with AI-powered observability

3.5x

ROI

Return on investment from mature observability practices

The latest Observability Forecast reveals that while digital services have become the primary business driver for most enterprises, many still lack the visibility needed to ensure reliability. This gap creates significant business risk as downtime directly impacts revenue and customer trust.

Leading organizations have recognized observability as a strategic capability that drives competitive advantage. Those with mature practices report significantly faster incident resolution times and measurable improvements in customer experience metrics.



# Splunk's AI-Powered Observability Platform



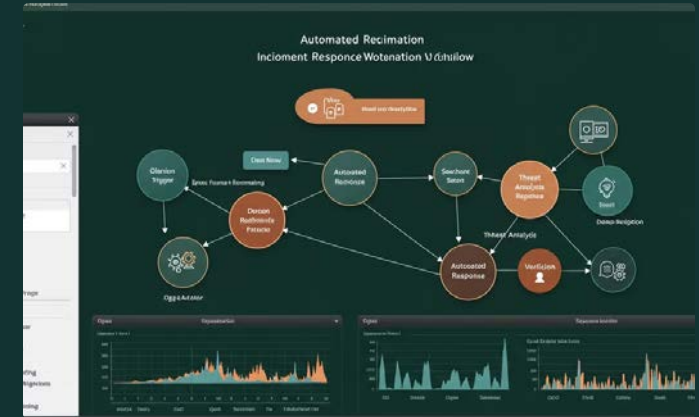
# Anomaly Detection

Machine learning algorithms automatically identify unusual patterns across metrics, logs, and traces, detecting potential issues before they impact users.



## Unified Analysis

Correlation engine connects seemingly unrelated events across your infrastructure, providing context that speeds troubleshooting and root cause analysis.

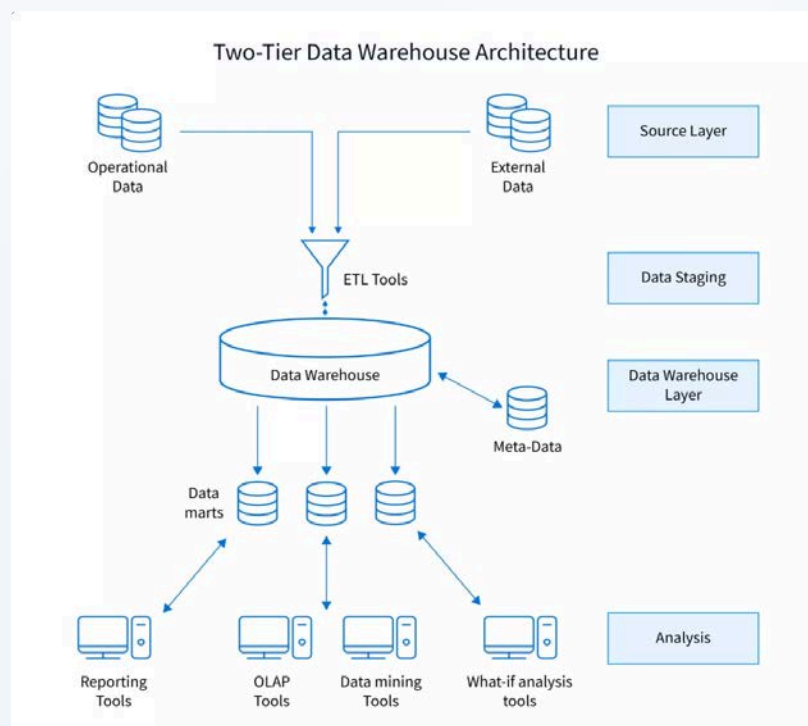


## Automated Remediation

Predefined playbooks automatically execute remediation actions when specific conditions are detected, reducing mean time to resolution.

Splunk's platform goes beyond traditional monitoring by applying machine learning at every layer of the observability stack. The system continuously learns from normal operating patterns to establish dynamic baselines that adapt to your environment's unique characteristics.

# Technical Deep Dive: Architecture



## Data Collection



- Universal Forwarders for logs
- OpenTelemetry for metrics and traces
- Cloud service integrations

## Processing Pipeline



- Field extraction and enrichment
- Data normalization
- Correlation ID injection

## AI/ML Layer



- Feature extraction
- Pattern recognition models
- Anomaly detection engines

## Visualization & Response



- Real-time dashboards
- Automated alerting
- Remediation workflows

Splunk's architecture is designed for massive scalability, processing millions of events per second while maintaining query performance. The pipeline begins with distributed data collection across your entire technology stack, feeding into a processing layer that normalizes and enriches data for analysis.

# Feature Selection & Model Tuning




Effective AI-powered observability depends heavily on selecting the right features for model training. Our approach combines traditional infrastructure metrics with application context, user experience indicators, and business KPIs to create a holistic view of system health.

Model tuning requires balancing sensitivity and specificity - you want to catch real issues without generating false positives. We use historical incident data to continuously refine models, gradually reducing false alarms while maintaining detection capability.



# Case Study: E-Commerce Peak Performance



A major retailer implemented Splunk's AI-powered observability before their annual sale event, which typically generates 10x normal traffic. Historical data was used to train models on expected patterns and potential failure modes.

During the event, the system detected an emerging bottleneck in the payment processing service as traffic increased. The anomaly score rose above threshold at 11:45, triggering automated scaling policies 15 minutes before traditional alerts would have fired. This proactive action prevented a checkout slowdown that would have cost an estimated \$2.3M in lost sales.



# Building Autonomous Remediation Workflows

## Anomaly Detection

ML algorithms identify deviations from normal patterns across infrastructure and application metrics

## Root Cause Analysis

Correlation engine identifies likely sources of the anomaly using causal inference techniques

## Response Selection

Decision engine chooses appropriate remediation action based on context and historical success

## Automated Execution

System triggers predefined workflows through integration with orchestration platforms

## Verification & Learning

Results are monitored and fed back into the system to improve future responses

Autonomous remediation represents the highest maturity level in observability practices. By connecting AI-powered detection with automated response mechanisms, organizations can resolve issues before they impact users.

Building these workflows requires careful planning and a phased approach. We recommend starting with low-risk, high-frequency incidents like auto-scaling or restart operations before progressing to more complex remediation scenarios.

# Implementation Framework

## Data Strategy

- Identify critical services and components
- Define instrumentation standards
- Establish data retention policies
- Create sampling strategies for high-volume telemetry

## Technical Implementation

- Deploy collection agents and integrations
- Configure data pipelines and enrichment
- Build service maps and dependency tracking
- Implement baseline ML models

## Operational Integration

- Align with incident management processes
- Establish SLOs and measurement frameworks
- Create role-specific dashboards and alerts
- Develop runbooks for common scenarios

Successful implementation requires alignment across data strategy, technical architecture, and operational processes. Our framework addresses each dimension with practical guidance based on hundreds of enterprise deployments.

Organizations typically complete initial implementation in 8-12 weeks, with an additional 4-6 months to fully mature their practices. The framework emphasizes incremental value delivery, with defined milestones providing measurable benefits throughout the journey.

# Next Steps on Your Observability Journey



## Assess Your Current Maturity

Complete our online assessment to identify gaps and opportunities in your current observability practices



## Develop Your Strategy

Work with our consultants to create a tailored roadmap addressing your specific technical environment and business goals



## Start Your Pilot

Implement AI-powered observability in a targeted area to demonstrate value and build organizational expertise



## Build Internal Capabilities

Train your teams on observability best practices and Splunk's AI capabilities through our certification programs

Wherever you are in your observability journey, practical next steps can accelerate your progress toward predictive operations and autonomous remediation. Our team of specialists can help you assess your current state and develop a customized implementation plan.

The technology landscape continues to evolve, but organizations that build strong observability foundations today will be positioned to adapt and thrive. We invite you to partner with us in transforming your operations through AI-powered observability.

Thank you