

Metrics, Logs, and Traces: A Unified Approach to Observability in Microservices

Pradeep Bhosale*

Citation: Bhosale P. Metrics, Logs, and Traces: A Unified Approach to Observability in Microservices. *J Artif Intell Mach Learn & Data Sci* 2022, 1(1), 2084-2088. DOI: doi.org/10.51219/JAIMLD/Pradeep-bhosale/458

Received: 03 June, 2022; Accepted: 28 June, 2022; Published: 30 June, 2022

*Corresponding author: Pradeep Bhosale, Senior Software Engineer (Independent Researcher), USA, E-mail: bhosale.pradeep1987@gmail.com

Copyright: © 2022 Bhosale P., This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ABSTRACT

Microservices architectures bring flexibility and modularity at scale, yet they also introduce operational complexity; services are scattered, with ephemeral pods and dynamic routing. Understanding system behavior under these conditions demands robust observability. Traditionally, organizations collect metrics (quantitative measures), logs (time-stamped event records), and traces (end-to-end request flows) in disparate silos. However, the synergy of these three pillars when unified yields deeper insights into root causes of performance bottlenecks or errors.

This paper explores a unified approach to observability in microservices, focusing on metrics, logs, and traces as complementary data sources. We describe the architectural components needed to ingest, store, and correlate these signals effectively; highlight anti-patterns (like ignoring distributed traces or over-collecting logs without index strategies); and provide best practices for bridging these signals via consistent instrumentation and tagging. Through visual diagrams, code examples, and real-world case studies, we illustrate how to debug cross-service latencies, identify resource constraints, and pinpoint failing dependencies in microservices-based systems. We also discuss how advanced solutions like open standards (Open Telemetry), centralized logging platforms, and distributed tracing frameworks enable more holistic DevOps workflows. Ultimately, this paper offers a roadmap for organizations aiming to build or evolve a comprehensive microservices observability strategy, bridging everyday debugging tasks with advanced, data-driven insights for system resilience.

Keywords: Microservices, Observability, Metrics, Logs, Traces, Distributed Tracing, DevOps, OpenTelemetry, Performance, Debugging, Classes, DevOps

1. Introduction

1.1. The observability challenge in microservices

Modern software architectures often rely on microservices, each running in containers orchestrated by platforms like Kubernetes. These microservices typically communicate via REST, gRPC, or messaging. This distributed model increases complexity around debugging performance or correctness issues: partial failures, unexpected latencies, or ephemeral container restarts can hamper root cause analysis^{1,2}.

Observability; the capability to infer the internal state of a system from external outputs becomes pivotal. Historically, teams used logs or metrics in isolation; over time, distributed tracing emerged for cross-service request tracking. A unified approach merging metrics, logs, and traces reveals more comprehensive system behavior, enabling faster triage and deeper analysis.

1.2. Scope and structure

This paper addresses:

- **Defining observability:** Contrasting monitoring with observability and outlining the “three pillars.”
- **Metrics:** Collection, storage, and usage patterns in microservices.
- **Logs:** Best practices for structured logging, indexing, correlation.
- **Traces:** Distributed tracing to see how requests propagate across services, diagnosing bottlenecks.
- Unifying these signals under a consistent tagging or instrumentation approach.
- Anti-Patterns that degrade system clarity or hamper root cause investigations.
- Real-World Scenarios and code/diagram-based examples for adopting a holistic approach.

2. Observability Fundamentals: Metrics, Logs, and Traces

2.1. Observability vs. monitoring

Monitoring typically focuses on known failure modes or metrics thresholds. Observability aims at ensuring that if an unknown or complex issue arises, the system’s emitted data (metrics, logs, traces) can help explain it. In a microservices environment, these signals are more critical due to ephemeral pod lifecycles and numerous service interactions³.

2.2. The three pillars

- **Metrics:** Quantifiable measures (CPU usage, request rates, latencies). Often aggregated over time and used for trend analysis or alerting.
- **Logs:** Event records with timestamps that detail errors, warnings, or debug info. Typically used to pinpoint the precise cause or sequence of events.
- **Traces:** Show an end-to-end path of a request across multiple services, each sub-operation’s timing, enabling quick identification of slow or error-prone segments⁴.

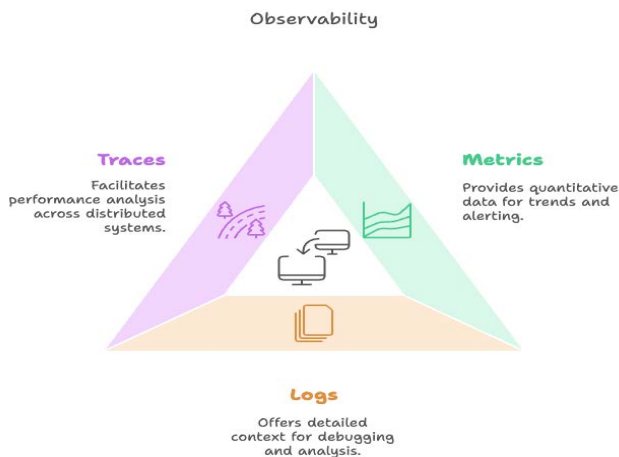


Figure 1: Three pillars of observability.

3. Metrics in Microservices

3.1. Collecting metrics

Microservices typically expose performance counters (HTTP request rates, latencies, memory usage) via instrumentation libraries. Tools like Prometheus scrape these endpoints, storing time-series data for queries or alerts⁵. For container environments,

cAdvisor or node exporters provide CPU, memory, and network usage.

Snippet (Prometheus scraping a microservice):

```
scrape_configs:
- job_name: 'myservice'
  kubernetes_sd_configs:
    - role: endpoints
  relabel_configs:
    - source_labels: [__meta_kubernetes_service_label_app]
      regex: myservice
      action: keep
```

3.2. Aggregation and dashboards

Grafana or alternative visualization platforms let teams see live or historical metrics in line charts, histograms, or heatmaps. Some common metrics:

- **HTTP requests:** requests_total, error_rate, p95 latency.
- **System resources:** CPU load, memory usage, disk I/O.
- **Custom domain metrics:** e.g., “cart_size,” “orders_placed,” or ML inference requests.

3.3. Anti-pattern: Over-collecting metrics without strategy

Issue: Capturing every possible metric at high cardinalities (like user_id dimension).

Effect: High storage costs, complicated queries.

Remedy: Focus on key operational metrics, use robust dimensional strategies (like only collecting user_id for selective debug sessions).

4. Logs for Detailed Context

4.1. Structured logging

Logs give time-stamped event data with human-readable or structured (JSON) fields. In microservices, logs frequently contain:

- Timestamp
- Service Name
- Pod ID or container ID
- Log Level (DEBUG, INFO, WARN, ERROR)
- Message including user or request IDs

Anti-Pattern: Using random or unstructured logs that hamper log correlation across services. A recommended approach is structured JSON logs with consistent fields (like trace_id, user_id)⁶.

4.2. Centralized log platforms

A typical pipeline sees logs from each container captured by a sidecar (Fluentd, Logstash) or node agent, then shipped to a central store (ElasticSearch, Splunk). This allows full-text indexing, letting devs quickly locate error messages or suspicious events across multiple microservices.



Figure 2: Centralized logging in containers.

4.3. Searching and alerting on logs

When a service errors out or an unexpected stack trace appears, ops teams can query the logs. They can also set up alerts for certain patterns: e.g., “If more than 50 ERROR logs from ‘checkout-service’ appear in 1 minute, raise an alert.”

5. Traces: The Missing Link

5.1. Distributed tracing basics

In microservices, a single user request might traverse multiple services. Traces record each service call as a “span,” detailing the operation name, start time, and duration. Tools like Jaeger or Zipkin visualize how these spans form a request tree, letting devs see which service contributed the largest latency chunk⁷.

Snippet (pseudo-code for instrumentation):

```
Span span = tracer.buildSpan("checkout").start();
try {
  // call user-service
  // call payment-service
} finally {
  span.finish();
}
```

5.2. Trace visualization

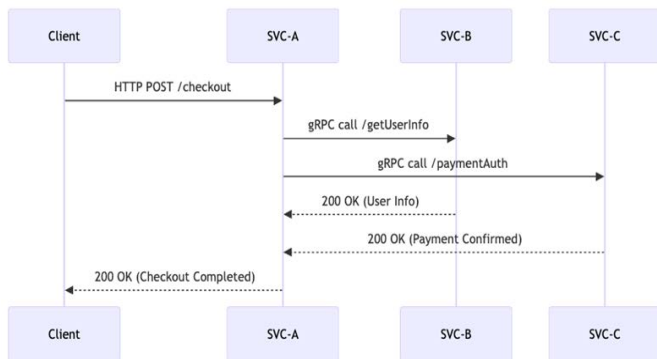


Figure 3: Multi-service request flow.

The trace aggregator (e.g., Jaeger) collects these spans, allowing an operator to see the timeline for each interaction.

5.3. Anti-pattern: Failing to propagate trace context

- **Issue:** If SVC-A doesn’t forward the trace ID or parent span ID to SVC-B or SVC-C, the aggregator sees disjoint traces, losing the end-to-end view.
- **Remedy:** Standardize an approach for injecting/extracting trace headers (like x-b3-traceid or W3C Trace Context) in each service.

6. Unifying Metrics, Logs, and Traces

6.1. Tagging and correlation

One powerful approach: each service logs with a trace_id or correlation_id. Metrics also might embed that ID for certain counters. The distributed tracer ensures each sub-span has the same ID. This synergy allows an operator to jump from a metric anomaly to the relevant logs, or from logs to the aggregated trace⁸.

6.2. Observability platforms

Vendors or open-source solutions unify these pillars:

- **Elastic stack:** Kibana for logs, metric beat for metrics, APM for distributed tracing.
- **Datadog:** Merges logs, metrics, and traces in one interface.
- **Open telemetry:** Emerging standard for unified instrumentation, offering SDKs to produce logs, metrics, traces in a consistent format.

7. Anti-Patterns in Observability

- **Siloed tools:** Using separate platforms for metrics/logs/traces with zero cross-linking. Operators must manually cross-reference timestamps or grep logs.
- **No standard tagging:** Each service uses different fields for “trace_id,” “request_id,” making correlation painful.
- **Excessive log verbosity:** Dumping all internal debug logs into production, overwhelming indexes.
- **Ignoring traces:** Relying solely on metrics/logs, leading to blind spots for cross-service latencies.

8. Implementation Approaches in Microservices

8.1. Sidecar or library

Sidecar: Some deploy specialized containers that intercept traffic, injecting or extracting trace headers. Tools like Envoy in a service mesh scenario (Istio) can produce distributed tracing data. Alternatively, each microservice can use an instrumentation library (like Open Telemetry, Brave for Java) that automatically wraps HTTP or gRPC calls⁹.

8.2. CI/CD integration

Pipelines may:

- Validate instrumentation presence (like scanning code for the base tracing library).
- Deploy to staging clusters, run synthetic transactions verifying trace spans are collected.
- Possibly set up ephemeral environment watchers for performance metrics.

9. Real-world case study #1: E-commerce observability

9.1. Context

A mid-sized e-commerce platform runs 15 microservices handling user, product catalog, checkout, payment, shipping, etc. They integrated a “three pillars” approach:

- **Metrics:** Each service emits Prometheus metrics for request throughput, error rates.
- **Logs:** JSON logs shipped to an ELK stack for query and correlation.
- **Traces:** They use Jaeger instrumentation in each service. The x-b3-traceid header is propagated via gRPC calls.

9.2. Key insights

- They discovered an unexpected 200 ms overhead in the shipping service by analyzing spans in Jaeger.
- During a CPU usage spike in the checkout service, logs revealed frequent “could not connect to DB” errors, correlating with a metric spike in error rate.
- The synergy saved hours of guesswork and manual cross-referencing. Operators improved auto-scaling thresholds after analyzing metrics and logs together.

10. Real-world case study #2: AdTech DSP

10.1. Scenario

An AdTech demand-side platform receives thousands of QPS in real-time bidding. Latency under 100 ms is paramount. They implement:

- Low-level metrics for each microservice, tracking p95 latencies.
- Structured logs containing “trace_id” in each request log line.
- Distributed tracing via OpenTelemetry integrated into each microservice pipeline.

10.2. Observed gains

- They quickly debug slow auctions in the aggregator service, identifying that user-profiling calls took 60+ ms.
- By pivoting from metrics (where aggregator’s latency soared) to traces, they pinpointed the “profile-service” call was stalling due to an unexpected DB lock.
- Logs confirmed the DB was missing an index, leading to query lock contentions. The fix was swift once identified.

11. Organizational and DevOps culture

11.1. Autonomy with governance

Each microservice team might choose their instrumentation approach but unify on a standard schema for logs and trace headers. A central SRE or DevOps group ensures consistent deployment of monitoring sidecars or libraries, standard label usage, and uniform metrics naming¹⁰.

11.2. Ongoing training

Even the best instrumentation can fail if developers don’t interpret or leverage these data sets effectively. Workshops on reading logs, building Grafana dashboards, or analysing distributed traces fosters a culture of continuous improvement in debugging and performance tuning.

12. Advanced Techniques

12.1. Service mesh observability

Service meshes like Istio or Linkerd can automatically intercept traffic, injecting trace IDs or collecting stats. This approach can remove burdens from application code, though advanced correlation might still require some in-service instrumentation for method-level spans or contextual logs.

12.2. AI/ML for anomaly detection

As data volumes grow, some organizations layer machine learning on top of logs, metrics, or trace data to detect anomalies (like unusual latency patterns or error spikes). While still an emerging approach, it can highlight hidden correlations or drift over time.

13. Building a Unified Observability Stack

13.1. Component Roles

- **Data collection:** Libraries (OpenTelemetry, Prometheus exporters), sidecars, or service mesh injection.
- **Transport:** Aggregators for logs (Fluentd) or traces (Jaeger agent).

- **Storage:** Time-series DB for metrics (Prometheus TSDB, InfluxDB), log DB (ElasticSearch), and trace DB (Jaeger or Zipkin).
- **Dashboards and alerting:** Tools like Grafana, Kibana, or custom UIs for end-user interactions.
- **Correlation:** Shared “trace_id” or “correlation_id” across logs, metrics, and traces. This might mean ingest transformations or consistent naming in code¹¹.

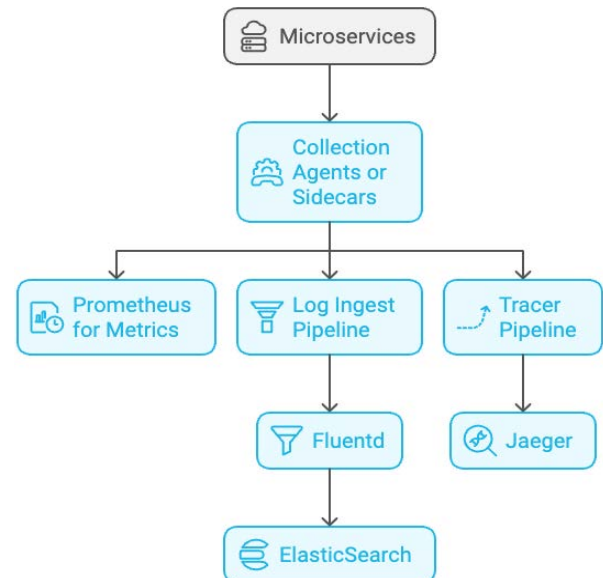


Figure 4: Observability stack.

Then combined queries, dashboards, or specialized correlation layers unify these signals.

14. Anti-Pattern Consolidation

- Ignoring or partial instrumentation in microservices (some produce traces, others don’t) → incomplete distributed view.
- Siloed Tools for metrics/logs/traces with zero cross-links → manual correlation is time-consuming.
- Excessive cardinalities in metrics → ballooning storage costs, hamper query performance.
- Low resolution or ephemeral logs → discarding logs too quickly or collecting insufficient detail hampers debugging.
- No standard naming for services or resources → confusing to follow ephemeral container names or inconsistent label usage.

15. Best Practices Summary

- **Define common instrumentation:** Standardize a single approach or library (OpenTelemetry) for metrics, logs, and traces.
- **Correlate data:** Insert trace_id or request_id in logs, connect metrics to those IDs, letting you pivot from a metric spike to the relevant logs/traces quickly.
- **Keep data balanced:** Avoid storing everything at infinite retention. Summaries or roll-ups can keep costs in check.
- Use Dashboards that unify all signals. E.g., a single place to see a service’s CPU usage, error logs, and slow trace spans in the same time window.
- **Train teams:** Observability is only as good as the culture that invests in reading, analyzing, and refining the signals.

16. Conclusion

In a microservices architecture, diagnosing performance or reliability issues demands a robust observability framework that ties together metrics, logs, and traces. By collecting quantitative time-series data (metrics), capturing event details (logs), and mapping end-to-end request flows (traces), teams gain a far more comprehensive vantage point for debugging. The synergy is especially crucial under ephemeral container lifecycles and dynamic routing where partial failures or cross-service latencies might otherwise remain opaque.

A unified approach means standard instrumentation libraries or sidecars producing consistent signals, stable correlation via `trace_id` or `span_id`, and integrated storage plus dashboards that let operators pivot from high-level anomalies to specific logs or trace timelines. This approach transforms “reactive monitoring” into “proactive observability,” accelerating root cause analysis, facilitating data-driven operational decisions, and building organizational confidence in rapid iteration cycles.

As microservices continue to proliferate, ongoing innovations like consolidated open-source solutions, further standardization in Open Telemetry, or AI-driven anomaly detection will refine how metrics, logs, and traces are leveraged. However, the foundational best practices introduced here remain vital for bridging ephemeral container sprawl and ensuring a stable, high-performance system. By combining metrics, logs, and traces in a single perspective, DevOps and SRE teams can unify their approach to diagnosing and optimizing microservices in dynamic, cloud-native environments.

17. References

1. Fowler M and Lewis J. “Microservices Resource Guide,” 2016.
2. Newman S. Building Microservices, O’Reilly Media, 2015.

3. Gilt Tech Blog, “Challenges in Observing Container-based Microservices,” 2018.
4. Krishnan S. “Distributed Tracing in a DevOps Culture,” ACM DevOps Conf, 2019.
5. <https://prometheus.io/>
6. <https://www.elastic.co/>
7. <https://www.jaegertracing.io/>
8. Netflix Tech Blog, “Correlating Service Metrics with Traces,” 2017.
9. CNCF Webinar, “Service Mesh Observability with Envoy,” 2019.
10. Molesky J and Sato T. “DevOps in Distributed Systems,” IEEE Software, 2013;30.
11. <https://opentelemetry.io/>
12. Brandolini A. Introducing EventStorming, Leanpub, 2013.
13. Blum A and Mansfield G. “Multi-Dimensional Logging Strategies,” ACMQueue, 2018;14.
14. Cockcroft G. “Multi-Service Debugging with Distributed Traces,” ACM SoCC Workshops, 2019.
15. Microservices Observability Whitepaper, “Advanced Tagging Approaches,” 2020.
16. <https://argo-cd.readthedocs.io/>
17. Datadog Blog, “Metrics, Logs, and Traces for Container Monitoring,” 2019.
18. AWS Observability Solutions, “Correlating Traces with Cloud Metrics,” 2021.