

From Ad-Hoc Scripts to Orchestrated Pipelines: Architecting a Resilient ELT Framework for Developer Productivity Metrics

Yuvraj Agrawal, Pallav Jain

Adobe Inc.

India

ABSTRACT

Developer Productivity Dashboards are essential for visualizing DevOps performance metrics such as Deployment Frequency and Change Failure Rate (DORA). However, the utility of these dashboards is frequently undermined by data reliability issues. In early iterations of our platform, ad-hoc ingestion scripts (Cron jobs) led to “silent failures,” where data gaps went undetected for days, eroding organizational trust. This paper reports on our experience migrating from legacy scheduling to a robust Extract-Load-Transform (ELT) pipeline using Directed Acyclic Graph (DAG) orchestration and Medallion Architecture. We detail the operational benefits of decoupling data extraction from transformation, the necessity of immutable raw history for metric redefinition, and the implementation of state-based dependency management. Our experience suggests that treating the metrics pipeline as a production-grade distributed system is a prerequisite for sustainable engineering analytics.

Keywords:—Data Engineering, ELT, Apache Airflow, DORA Metrics, Medallion Architecture, DevOps.

I. INTRODUCTION

The measurement of software delivery performance using the DORA (DevOps Research and Assessment) framework has become a standard practice for high-performing technology organizations [1]. Metrics such as *Deployment Frequency* and *Lead Time for Changes* provide objective signals on engineering velocity and stability. However, while the definitions of these metrics are well-documented, the architectural patterns required to compute them reliably at scale are less explored.

In our organization’s initial implementation, metrics were aggregated using isolated Python scripts scheduled via standard Cron utilities. This approach appeared sufficient for a single team but collapsed as we scaled to ingest data from Jira, GitHub, and Jenkins across multiple repositories and business units.

We encountered a critical data quality phenomenon which we formally define as the **Phantom Zero**:

Phantom Zero: An failure mode in data pipelines where an upstream extraction error (e.g., a silent API pagination failure or timeout) is interpreted by downstream transformation logic as a valid empty set. This results in the calculation and reporting of “zero activity” (e.g., 0 Deployments) rather than a system exception, making the failure indistinguishable from a day of low engineering activity.

In our legacy system, Phantom Zeros often persisted for days before being detected manually, significantly damaging stakeholder trust in the dashboard. A dashboard that reports “0 incidents” is useless if the user cannot distinguish between a stable system and a broken sensor.

This paper presents an architectural experience

report on solving this “Last Mile” problem of data trust. We describe our migration to an orchestrated ELT pipeline that eliminates Phantom Zeros through:

- 1) **State-Based Execution:** Transformations run only when dependencies are strictly satisfied, utilizing “Hard Gates” to prevent data pollution.
- 2) **Time-Travel Capability:** The ability to recompute historical metrics when business definitions change without re-fetching raw data from rate-limited APIs.
- 3) **Auditability:** A clear lineage from raw JSON payloads to final aggregated metrics.

The scope of this work is limited to batch-oriented organizational health metrics; real-time operational telemetry (e.g., server CPU usage) is outside the scope of this architecture.

II. THE ARCHITECTURAL PARADIGM

Our system ingests data from three primary disparate sources: structured issue tracking data (Jira), semi-structured version control metadata (GitHub), and unstructured build logs (Jenkins). To manage this complexity, we adopted an ELT (Extract-Load-Transform) paradigm supported by a Medallion Architecture.

A. The Case for ELT over ETL

In our early ETL (Extract-Transform-Load) attempts, we transformed raw API responses into metric values immediately upon ingestion. This proved fragile due to the evolving nature of metric definitions. For instance, when the organization redefined “Change Failure” to exclude “P4 Incidents,” the ETL model required re-fetching nine months of historical data from

the Jira API. This operation triggered strict API rate limits and resulted in partial data gaps.

By shifting to **ELT**, we prioritize loading raw JSON payloads into storage immediately. Transformations occur downstream within the database engine. This decoupling allows us to iterate on metric definitions (the “Transform” step) using stored history, independent of the source system’s availability or rate limits.

B. Medallion Architecture Implementation

We organize data into three distinct layers of refinement to ensure data quality and traceability.

1) *Bronze Layer (Raw Audit)*: The Bronze layer acts as an immutable, append-only store for raw data.

- **Design:** We store the exact JSON payload returned by the source API, enriched only with ingestion metadata (‘fetched_at’, ‘source_system’, ‘execution_id’). No validation is performed at this stage to maximize ingestion throughput.
- **Retention Strategy:** We initially retained Bronze data indefinitely. However, we observed that debugging utility diminishes rapidly over time. We now enforce a 30-day retention policy, which covers approximately 99% of debugging scenarios while controlling storage costs.
- **Operational Value:** The Bronze layer serves as the reliable fallback mechanism. In three separate instances of downstream logic errors, we were able to replay the Bronze data to correct the Silver layer without making a single external API call.

2) *Silver Layer (Standardization)*: The Silver layer is responsible for normalization and cleaning.

- **Schema Enforcement:** Timestamps from diverse sources (e.g., Jenkins’ Unix epoch vs. GitHub’s ISO 8601) are converted to a uniform UTC ‘datetime’ object.
- **Identity Mapping:** A significant challenge in DORA metrics is identity resolution. A user may exist as ‘jdoe’ in GitHub but ‘jane.doe@company.com’ in Jira. The Silver layer applies a mapping logic to resolve these disparate identities into a single ‘contributor_id’.
- **Filtering:** Invalid records, such as aborted Jenkins builds or “bot” commits, are filtered out at this stage.

3) *Gold Layer (Business Aggregates)*: The Gold layer is optimized for read performance and consumption by the dashboard visualization tool (e.g., Metabase or Tableau).

- **Pre-Aggregation:** Instead of storing individual events, this layer stores aggregated metrics (e.g., “Daily Deployment Count,” “Median Lead Time”).
- **Query Latency:** Dashboard queries against the Gold layer scan hundreds of documents, whereas queries against the Silver layer would require scanning tens of thousands of records. This decoupling ensures dashboard responsiveness is not linear to the volume of raw data.

C. Storage Strategy: Time-Series Optimization

Given the temporal nature of DORA metrics, we utilize MongoDB Time-Series collections (version 5.0+). We observed that the choice of the ‘metaField’

(partitioning key) is the single most critical design decision.

- **Initial Failure:** We initially used ‘repository_name’ as the metadata field. However, in a microservices environment with hundreds of small repositories, this led to high cardinality and bucket fragmentation, degrading compression ratios.
- **Optimized Approach:** We refactored to use ‘team_id’ or ‘product_area’. This aligned storage locality with query patterns, as dashboards are typically filtered by team, resulting in a 40% improvement in query performance.

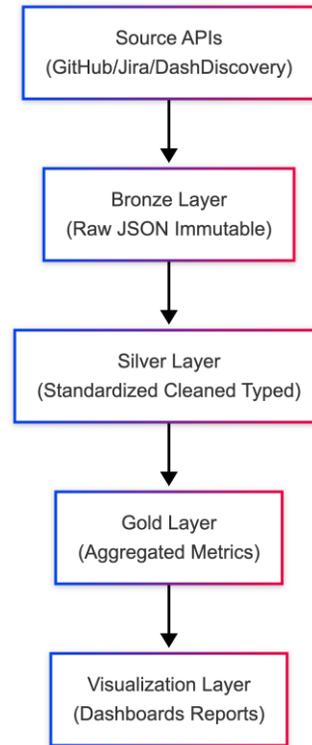


Fig. 1. The Medallion Architecture applied to Developer Productivity Metrics.

III. ORCHESTRATION STRATEGY

The structural layers described above are managed by an Orchestration Engine (Apache Airflow), modeling the pipeline as a Directed Acyclic Graph (DAG). This shift from time-based to state-based execution was a major contributing factor to improved reliability.

A. State-Based Dependency Management

In our legacy Cron approach, scripts were loosely coupled by time (e.g., Script A runs at 1:00 AM, Script B at 1:30 AM). If Script A failed or ran long, Script B would execute on stale data. In the DAG approach, the ‘Calculate_Gold_Metrics’ task has a strict, blocking dependency on the success of ‘Normalize_Silver_Data’. If the GitHub API is unavailable and the Bronze ingestion fails, the pipeline halts immediately. This “Hard Gate” mechanism ensures that we never publish partial or corrupted data to the Gold layer, prioritizing data correctness over freshness.

B. Idempotency and the “Backfill” Capability

A requirement for a robust ELT pipeline is idempotency: the ability to run the pipeline multiple times for the same execution date without creating duplicate records. This is achieved via “Upsert” (Update or Insert) operations.

- **Composite Keys:** We define unique keys at the database level for the Silver and Gold layers. For example, a Gold record is uniquely identified by ‘date, team_id, metric_type’.
- **Atomic Writes:** Writes to the Gold layer are atomic. If a re-run occurs, the existing record for that day is replaced.

This idempotency enables “Time Travel” or Backfilling. When we introduced a new metric (“Mean Time To Recovery”), we defined the logic in the DAG and triggered a backfill for the past 365 days. The scheduler spawned worker processes to recompute the Gold layer from the existing Silver data in parallel, completing a year’s worth of analytics in under 20 minutes.

C. Resiliency Patterns

To handle the inherent instability of external APIs, we implemented several resiliency patterns within the DAGs:

- 1) **Exponential Backoff:** Tasks utilize a retry policy with exponential backoff (e.g., 3 retries, doubling delay from 5 to 20 minutes). We capped the maximum delay to 45 minutes to prevent “zombie tasks” from blocking the queue indefinitely.
- 2) **Dead Letter Queues (DLQ):** If a task fails after all retries, the payload metadata is sent to a DLQ for manual inspection, preventing the failure from crashing the entire orchestration server.
- 3) **Sensor Tasks:** We utilize “Sensors” to verify data volume before processing. If the number of fetched Jira tickets is 90% lower than the 30-day moving average, the pipeline pauses and alerts an engineer, preventing the “Phantom Zero” issue where a silent API failure mimics a day of zero activity.

IV. REAL-TIME ALERTING: THE “PUSH” MODEL

While the DORA metrics themselves are calculated in batches (e.g., daily or every 12 hours), the *notification* of critical regressions must be timely. A dashboard is a passive tool; relying on engineers to proactively check it introduces significant latency between a metric regression and remediation. To bridge this gap, we implemented an active alerting layer.

A. The Latency Gap in “Pull” Models

Our legacy alerting system relied on a separate Cron job that polled the database every 60 minutes to check for threshold breaches (e.g., Change Failure Rate > 15%).

$$Latency_{alert} \approx T_{poll} + T_{processing}. \quad (1)$$

This “Pull” model was inefficient. It wasted computational resources querying the database when no data had changed, and it introduced an unavoidable latency window. In one incident, a surge in deployment failures at 2:05 PM was not detected until the 3:00 PM poll, delaying the incident response by 55 minutes.

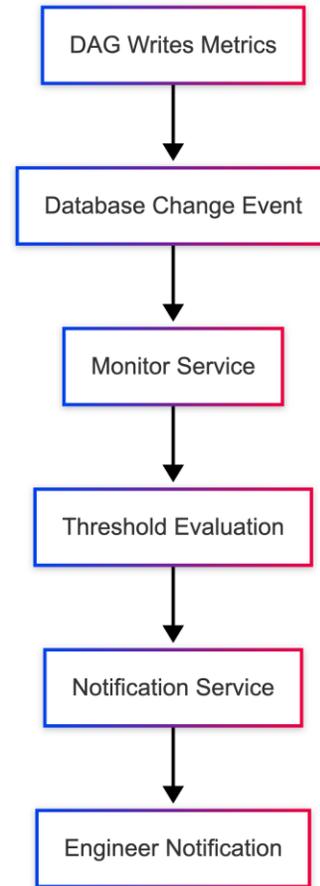


Fig. 2. Real-Time Alerting Architecture via Change Streams.

B. Event-Driven Architecture via Change Streams

We shifted to an event-driven “Push” model using MongoDB Change Streams. This mechanism allows applications to subscribe to real-time data changes at the database level without the complexity of tailing the operation log (oplog).

1) **Architecture Components:** The alerting pipeline consists of three decoupled components:

- 1) **The Producer (DAG):** The Airflow DAG writes a calculated metric to the Gold layer. This is an atomic write operation.
- 2) **The Event Bus (Database):** The database emits a change event containing the document delta.
- 3) **The Consumer (Monitor Service):** A lightweight, always-on microservice listens for ‘insert’ or ‘replace’ events on the Gold collection.

2) **Resiliency via Resume Tokens:** A critical challenge in event-driven systems is reliability. If the Monitor Service crashes or is redeployed, events occurring during downtime

could be lost. We leveraged the ‘resume_token’ feature of Change Streams. The Monitor Service persists the token of the

last successfully processed event to persistent storage (Redis). Upon restart, it presents this token to the database, which replays all events that occurred during the downtime.

- **Observation:** This distinct delivery semantic ensures that no alerts are missed, even during system maintenance windows.

V. OPERATIONAL IMPACT AND OBSERVED OUTCOMES

The migration from ad-hoc scripting to orchestrated ELT yielded measurable improvements in the operability and reliability of our metrics platform. Table I summarizes our observations comparing the legacy system to the new architecture over a 6-month operation period.

TABLE I
OPERATIONAL COMPARISON: LEGACY VS. ORCHESTRATED

Metric	Legacy (Cron)	Orchestrated (DAG)
Failure Detection	2-5 Days (User reported)	<1 Hour (Automated)
Data Lineage	None (Black Box)	Full Task History
Backfill Effort	Manual Scripting (Days)	UI Trigger (Minutes)
Recovery Time	High (Manual Cleanup)	Low (Idempotent Re-run)

A. Scenario Analysis: The Schema Change

During a routine update, the Jira API schema changed, renaming a critical field used for “Lead Time” calculation.

- **Legacy Outcome (Projected):** The script would have either crashed or, worse, inserted null values into the database. Recovery would have required writing a custom one-off script to delete the corrupted rows and re-fetch data.
- **Observed DAG Outcome:** The *Extract* task failed validation immediately. The pipeline halted, preventing any corruption of the Silver or Gold layers. The Data Engineering team was alerted via Slack. Once the parser was patched, we simply clicked “Clear State” in the Airflow UI. The pipeline resumed from the point of failure, automatically catching up on the missed data without duplicate ingestion.

VI. LESSONS LEARNED

Throughout the design and implementation of this architecture, we identified several non-obvious principles that are critical for engineering analytics pipelines.

A. Bronze is for Audit, Not Querying

In our initial design, we attempted to build “live” dashboards directly on top of the Bronze (Raw) layer to reduce latency. This was a mistake. The schema inconsistency of raw API data led to complex, slow, and fragile queries that broke

whenever a vendor changed their API response format. **Lesson:** Strictly enforce the Medallion separation. The Gold layer creates a necessary contract between data engineering and data visualization.

B. Alert on Absence of Data

The most dangerous failure mode in analytics is not a crash, but a successful run that processes zero records (the “Phantom Zero”). Standard error monitoring does not catch this. **Lesson:** We added “Data Volume Checks” as a final step in our DAGs. If the number of processed rows deviates by more than 2 standard deviations from the 30-day moving average, the pipeline alerts the team, even if the code executed without error.

C. Token Rotation Management

Hardcoded API tokens in environment variables proved to be a single point of failure and a security risk. **Lesson:** We integrated our orchestration engine with a Secrets Manager. Tasks fetch fresh credentials dynamically at runtime. This significantly reduced authentication failures caused by expired tokens and improved our security posture.

VII. LIMITATIONS

While superior in reliability, the orchestrated ELT architecture introduces trade-offs that must be acknowledged.

A. Infrastructure Complexity

A Cron job requires a single server and a text file. An Airflow-based architecture requires a Scheduler, a Web Server, a Metadata Database (PostgreSQL), and a queuing mechanism (Redis). For small teams with low data volume and simple metrics, the operational overhead of maintaining this distributed system may outweigh the benefits of observability.

B. Batch Latency

Although the alerting layer is real-time relative to the write event, the metrics themselves are typically batch-processed (e.g., every 6 or 12 hours). The system reflects the state of the world as of the last DAG run. While sufficient for DORA metrics (which track long-term trends), this architecture is not suitable for real-time operational monitoring (e.g., server CPU spikes), which requires a streaming architecture (e.g., Apache Kafka or Flink).

VIII. FUTURE WORK

While the current orchestrated ELT architecture has stabilized our reporting of DORA metrics, we have identified two primary areas for further evolution: the transition to streaming architectures for operational metrics and the adoption of machine learning for anomaly detection.

A. Beyond Batch: Streaming Ingestion

Our current Airflow implementation operates on a micro-batch schedule (e.g., every 4–6 hours). While sufficient for strategic metrics like *Deployment Frequency*, this latency is

suboptimal for operational metrics such as *Time to Restore Service* during active incidents. Future work involves implementing a hybrid architecture. We aim to introduce a streaming ingestion layer (using Apache Kafka or Confluent) alongside the batch layer. This would allow us to compute “Time to Recovery” in near-real-time using windowed aggregations.

(e.g., via Apache Flink), providing immediate feedback to incident commanders while retaining the batch layer for accurate historical trend analysis.

C. AI-Driven Anomaly Detection

Currently, our pipeline relies on static thresholds for data quality checks (e.g., the 90% volume drop rule discussed in Section III). While effective for catastrophic failures, static thresholds struggle with seasonality (e.g., holidays, weekends) and subtle degradations. We are prototyping an unsupervised machine learning approach to learn the “normal” heartbeat of our engineering data. By replacing static rules with dynamic, seasonality-aware baselines, we aim to detect subtle data quality issues—such as a gradual drift in API latency or a slow decline in ingestion volume—before they trigger hard alerts, further reducing the cognitive load on the platform engineering team.

IX. CONCLUSION

The transition from ad-hoc scripting to a disciplined ELT pipeline is not merely a technical upgrade; it is a prerequisite for building trust in Developer Productivity Dashboards. By adopting a Medallion Architecture, we ensure data auditability and consistency. By utilizing DAG-based orchestration, we eliminate silent failures and enable robust failure recovery. Finally, by layering event-driven Change Streams on top of the storage layer, we bridge the gap between historical analysis and real-time operational alerting.

Our experience suggests that the primary challenge in engineering analytics is not the definition of the metrics themselves, but the maintenance of confidence in the data supply chain that drives them. A dashboard that users do not trust is a dashboard that users do not use.

REFERENCES

- [1] N. Forsgren, J. Humble, and G. Kim, *Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations*, IT Revolution Press, 2018.
- [2] Google Cloud, “DevOps Research and Assessment (DORA),” [Online]. Available: <https://cloud.google.com/blog/products/devops-sre>.
- [3] Apache Software Foundation, “Apache Airflow Documentation,” 2024. [Online]. Available:

<https://airflow.apache.org/docs/>

- [4] MongoDB Inc., “Time Series Collections,” *MongoDB Manual*, 2024. [Online]. Available: <https://www.mongodb.com/docs/manual/core/timeseries-collections/>
- [5] M. Armbrust et al., “Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores,” *Proc. VLDB Endow.*, vol. 13, no. 12, pp. 3411–3424, 2020.

APPENDIX A

AIRFLOW DASHBOARD SCREENSHOTS

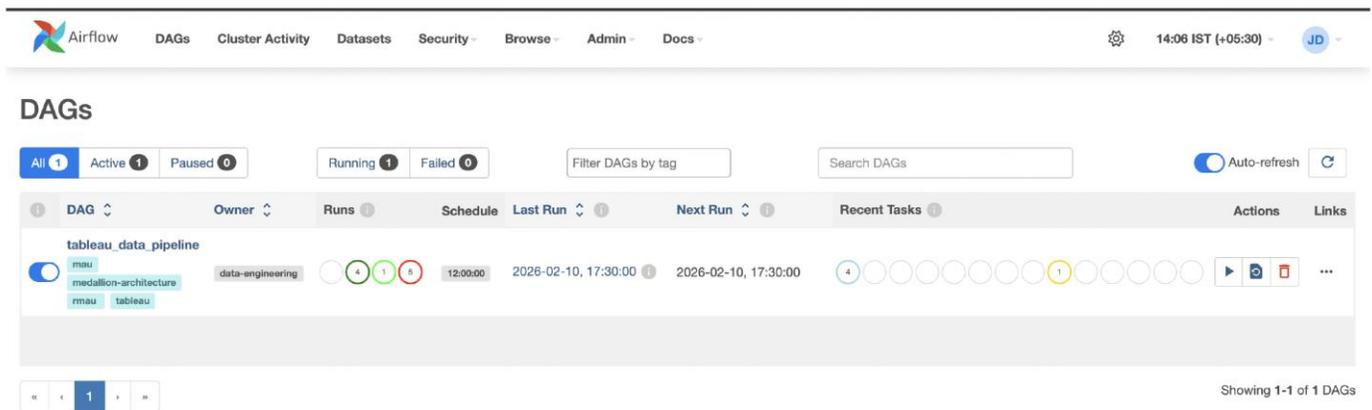


Fig. 3. Airflow DAG Overview showing tasks and execution state.

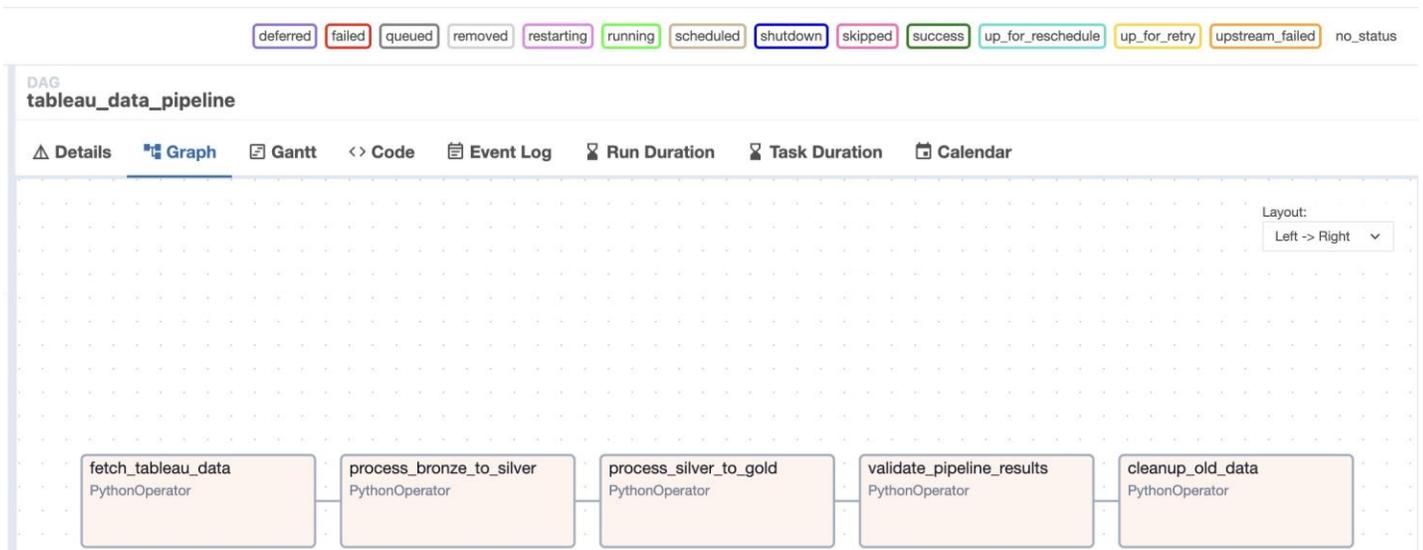


Fig. 4. Airflow Graph View highlighting pipeline stages.