SMARTER ORCHESTRATION, REAL RESULTS: HOW CONTROL-M CONNECTS YOUR DATA ECOSYSTEM



Disconnected tools. Missed SLAs. Broken pipelines. For many data teams, the challenge isn't building—it's making everything work together reliably.

Control-M bridges that gap. It connects the technologies you already use—like Snowflake, SageMaker, and Tableau—into unified and resilient workflows.

Here's how Control-M helps you orchestrate across your entire data stack, delivering operational efficiency without adding complexity.

Orchestrate full ML pipelines with Control-M and Amazon SageMaker

Machine learning delivers real value only when it reaches production. Control-M helps you get there faster by transforming SageMaker workflows into orchestrated, resilient pipelines.

Instead of hand-cranked scripts and brittle logic, Control-M offers:

- Direct integration to launch and monitor SageMaker training and inference jobs
- Built-in SLA tracking, parallel job execution (up to 50 per agent), and predictive alerts
- Seamless coordination with upstream data prep and downstream deployment steps

Real-world example: A data science team used Control-M to automate everything from data prep to monitoring. The result: a repeatable ML lifecycle that integrates directly with CI/CD pipelines.

Run smarter ELT workflows with Snowflake

Snowflake is built for scale—but it takes orchestration to turn SQL jobs, UDFs, and transformations into stable, production-ready workflows.

Control-M integrates directly with Snowflake so you can:

- Automate, orchestrate, and monitor SQL and UDF executions with clear visibility
- Automatically trigger downstream steps based on completion status
- Ensure consistency and data integrity across platforms and tools

Use case: Orchestrate an end-to-end pipeline from Kafka ingestion to Snowflake transformation to Tableau visualization—all within a single, governed Control-M workflow.

Scale Airflow with centralized orchestration

Apache Airflow is great for DAG orchestration—but complexity can grow quickly. That's where Control-M comes in.

Control-M works with Airflow to:

- Centralize DAG execution and integrate with other toolchains
- Track job status visually with SLA and error-handling logic
- Trigger Airflow jobs from upstream events or use them within larger workflows

Best of both worlds: Use Airflow where it excels but orchestrate across your stack with Control-M to ensure continuity and control.

Deliver always-fresh analytics with Power BI, Looker, and Tableau

Dashboards only add value when they reflect current, accurate data. Control-M connects BI platforms to the rest of your pipeline, ensuring updates happen on time, every time.

Through native and API-based integrations, you can:

- Trigger dashboard refreshes after upstream ELT or ML jobs complete
- Link model outputs directly to reporting tools and decision-making
- Monitor and resolve issues proactively before end users are affected

Example: Push SageMaker inference results into Power BI or Tableau dashboards immediately upon model run completion—no manual intervention required.

Why Control-M?

If you're running modern data platforms across multi-cloud environments, Control-M helps everything run better—without reinventing your stack. It delivers:

- **Unified governance**: Centralized visibility, management, monitoring, and alerting
- **Resilience**: Auto-recovery, retries, and SLA intelligence built-in
- **Seamless integration:** Up to three integrations added monthly, including AWS, Azure, GCP, and many more

Control-M is the orchestration engine that makes your stack smarter—by making it work together.

See it in action

Whether you're deploying models, syncing dashboards, or scheduling ELT jobs, Control-M helps you build smarter workflows that don't break under pressure.

Explore real-world demos and integration guides—no forms, no sales pitch.
Take the Control-M product tour to see how it fits your stack.