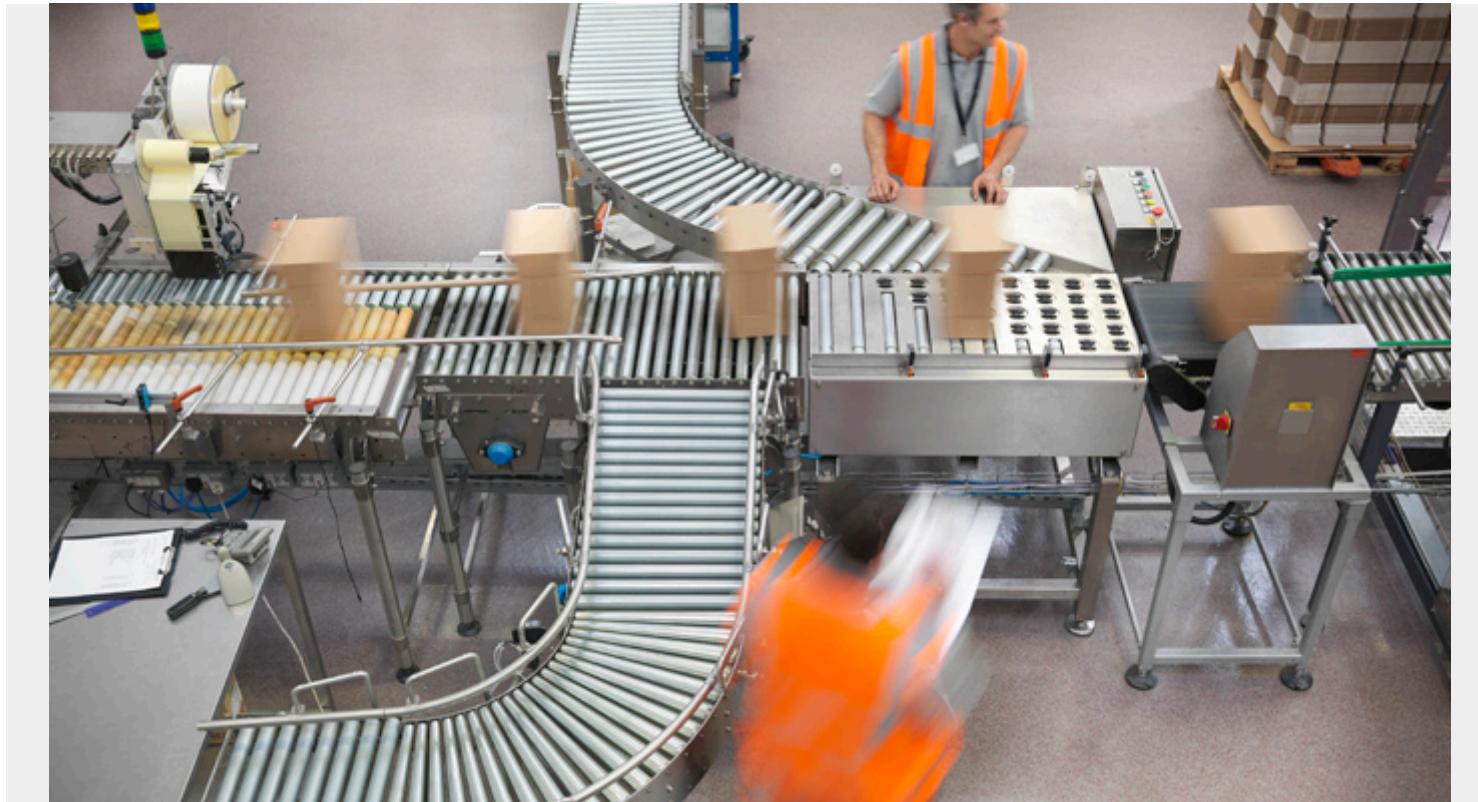


HOW BATCH PROCESSING IS CHANGING IN THE AGE OF AI



Batch processing is [workflow orchestration](#) which executes units of work over a bounded set of inputs. It runs continuously—not just after hours—optimizing for correctness, completeness, and throughput. Today's batch processing coordinates multi-step, dependency-aware work across data, services, and environments, and can be triggered by events just as often as by schedules.

This approach is often called [workload automation](#) (WLA) or [job scheduling](#), but these names don't capture how dramatically it has evolved. For many businesses, batch processing isn't a legacy technology—it's the backbone of their most critical operations, from payroll to AI model training.

In this article, we'll redefine what batch processing means today, examine how it complements stream processing rather than competing with it, and provide guidance on leveraging both approaches for maximum business impact.

The Evolution of Batch Processing: From "After Hours" to 24/7 Orchestration

The outdated definitions of batch processing reflect a business world that no longer exists. For decades, batch was defined by when it ran: after business hours, during overnight windows when systems were otherwise idle. This made sense when businesses closed at 5 PM and computing resources were scarce.

But modern businesses operate 24/7. Global operations, always-on digital services, and real-time

customer expectations mean there are no "idle hours" anymore. The batch jobs that power invoicing, payroll, employee onboarding, and AI model training don't wait for nighttime—they run continuously throughout the day.

What has changed isn't whether batch processing is needed, but how we understand and define it. Modern terminology focuses on "workflow orchestration" rather than "overnight jobs." The distinction between batch and other processing methods is no longer about when you process, but why you process that way.

Today's batch processing has evolved to become:

- **Continuous:** Running 24/7 rather than only during off-peak hours
- **Event-driven:** Triggered by business events (customer orders, file arrivals, API calls) as well as schedules
- **Orchestrated:** Coordinating complex, multi-step dependencies across systems
- **Purpose-built:** Optimized for correlation, completeness, and throughput rather than just deferred execution

Modern Workflow Orchestration: Dependencies and Triggers

Today's batch processing relies on sophisticated orchestration platforms that manage dependencies, triggers, and execution paths without constant human oversight.

Event-Driven and Schedule-Driven Triggers

Modern batch workflows can be initiated in multiple ways:

- **Schedule-based:** Traditional time-based triggers for recurring work (daily reports, weekly reconciliation)
- **Event-based:** Business events that trigger workflow execution (customer places order, file arrives in storage, fraud signal detected)
- **Dependency-based:** Workflows that begin when prerequisite jobs complete or data becomes available
- **Hybrid:** Combinations of schedules, events, and dependencies for complex orchestration

Dependency Management and Execution Paths

A [workflow orchestration platform](#) manages dependencies through flows which define:

- **Sequential steps:** tasks that must run in a specific order
- **Parallel execution:** Independent tasks that can run at the same time to reduce overall duration
- **Conditional branching:** different execution paths based on data, events or business rules
- **Compensating actions:** rollback or corrective steps which run when a task fails

Exception Handling and Observability

Modern systems use exception-based management to notify stakeholders only when intervention is needed:

- **Monitors** track job duration, data volumes, and quality metrics, alerting on anomalies

- **Per-step metrics** provide visibility into performance and bottlenecks
- **Structured logging and tracing** enable rapid troubleshooting and root cause analysis
- **SLA/SLO monitoring** ensures business commitments are met with precise alerting thresholds

The goal is strategic human intervention: people are engaged when their expertise is needed for decisions, not for routine monitoring.

When to Use Batch Processing

Batch processing should be considered when business processes require correlation of events over time, completeness of data, or high-volume throughput. The choice between batch and stream isn't about modernity—it's about matching the processing approach to business requirements.

Use batch processing when:

- **Actions depend on correlating multiple events over time** rather than responding to a single signal
- **Complete, validated datasets are required** for correctness, compliance, or audit purposes
- **Workloads are compute-intensive and throughput-oriented** such as ETL/ELT pipelines, model training, or bulk updates
- **Complex algorithms need access to entire data sets** for analysis or transformation
- **Simpler operations and better cost efficiency** are priorities for large-volume work
- **The work is repetitive** and benefits from standardized orchestration patterns

How Batch and Stream Processing Work Together

The most significant misconception about batch and stream processing is that they compete. In reality, they complement each other, each optimized for different business purposes.

Stream processing excels at acting on signals immediately: detecting fraud as it happens, responding to IoT sensor alerts, or triggering real-time personalization. These workloads process unbounded event flows with millisecond or second-level latency requirements.

But what happens after the immediate action? Consider fraud detection: stream processing catches the suspicious transaction in real-time and flags it instantly. Then batch processing orchestrates everything that follows—notifying stakeholders, blocking accounts, halting payments, opening investigation cases, reconciling related transactions, updating customer communications, and maintaining audit logs. These follow-up steps require correlation across multiple systems and must be completed reliably, but they don't all need to happen in milliseconds.

This pattern appears across industries:

- **E-commerce:** Stream processes the checkout event; batch orchestrates order fulfillment, inventory updates, and invoicing
- **Financial services:** Stream detects anomalies; batch handles account holds, investigation workflows, and regulatory reporting
- **AI and machine learning:** Stream may feed real-time inference features; batch powers data preparation, model training, evaluation, and periodic retraining

The distinction isn't batch versus stream—it's understanding which business processes need immediate response (stream) and which need reliable, complete, correlated processing (batch).

Batch Processing

Processes large volume of data all at once

May require dedicated staff to handle issues

Stream Processing

Analyzes streaming, cross-device data in near real time

The data output rate must be just as fast as the data input rate

Benefits and drawbacks of common data processing types

Comparison: When to Choose Each Approach

Use stream processing when:

- A single event (or very short time window) requires immediate action
- SLAs are measured in milliseconds or seconds
- You can tolerate higher complexity and resource intensity to achieve low latency
- Real-time signals drive time-sensitive decisions (fraud alerts, IoT monitoring, live personalization)

Use batch processing when:

- Actions depend on correlating multiple events over time
- Complete, validated datasets are necessary for correctness and compliance
- Workloads are compute-intensive and throughput-oriented
- Simpler operations and better cost efficiency matter for large volumes
- Business services require reliable, auditable completion (payroll, invoicing, model training)

Use a hybrid approach when:

- Real-time signals trigger multi-step follow-up processes
- You need both fast detection and reliable, auditable completion of downstream work

- Different parts of the workflow have different latency requirements

Stream processing is generally more complex and more compute and resource intensive. If business processes rely on collecting data—even when that collection is relatively small—batch or workflow orchestration may be the better choice. When action is dependent on a series of events correlated over time, batch is the stronger approach.

Why Batch Powers Today's Most Modern Workloads

The perception that batch processing is outdated couldn't be further from reality. The most modern, cutting-edge workloads in technology today depend heavily on batch processing.

AI and Machine Learning Workloads

Consider that OpenAI recently concluded a \$300 billion deal with Oracle to purchase compute resources for model training. Large language models and generative AI represent the frontier of modern computing—and the way effective models are built, taught, and updated is through batch approaches.

There is no AI in use today that doesn't rely on and depend heavily on batch processing to:

- **Prepare and curate training data** from diverse sources
- **Train models** on massive datasets with high computational requirements
- **Evaluate model performance** across test sets and benchmarks
- **Gather and process feedback** for continuous model improvement
- **Retrain and update models** on regular schedules or when performance degrades
- **Backfill features and labels** for historical analysis

Machine learning workloads are batch-intensive by nature: they require complete datasets, complex computations over bounded inputs, and iterative processing cycles. This is batch processing at its most sophisticated—and most essential.

Critical Business Services

Businesses rely on batch processing for their most critical operations, both internal and customer-facing:

Internal services:

- Payroll processing and compensation calculations
- Employee onboarding workflows across HR, IT, and access management systems
- Benefits administration and enrollment
- Compliance reporting and audit log generation

Customer-facing services:

- Invoicing and billing across subscription cycles
- Statement generation and distribution
- Order fulfillment orchestration
- Reconciliation and [financial close processes](#)

The benefit organizations gain from batch processing is straightforward: they can deliver the services they exist to deliver to their customers. The focus is not whether batch processing is needed, but how to enhance the level of automation, make it more efficient, and give it the ability to process ever-growing volumes of data as effectively as possible.

Advantages of Modern Batch Processing

Reliability and Correctness

Batch processing excels at ensuring completeness and correctness because it works with bounded, complete datasets. This allows for:

- Validation of data quality before downstream processing begins
- Detection of anomalies through completeness checks and thresholds
- Reliable retry and reprocessing mechanisms when issues occur
- Full audit trails and lineage tracking for compliance

Cost Efficiency and Throughput

For high-volume workloads, batch processing delivers better economics:

- Right-sized compute resources matched to workload characteristics
- Ability to use spot or preemptible instances for cost savings
- Batching and micro-batching to optimize throughput
- Co-location of compute with data to reduce transfer costs

Without the overhead of maintaining constant low-latency infrastructure, batch systems can process millions of records efficiently.

Operational Simplicity

Compared to real-time or stream processing, batch is significantly less complex:

- Standard orchestration patterns for common workflows
- Simpler infrastructure without the need for specialized streaming platforms
- Easier debugging with clear execution boundaries and complete input/output tracking
- Lower barrier to entry for organizations building data and automation capabilities

Continuous Processing Without Manual Oversight

Modern batch systems run 24/7 with minimal human intervention:

- Exception-based alerting means teams only engage when issues arise
- Platform engineering teams maintain infrastructure while business teams own workflows
- Self-service capabilities let domain experts define and manage their own orchestration
- AI assistance accelerates troubleshooting by surfacing relevant knowledge and past incidents

Considerations for Successful Batch Implementation

While batch processing delivers significant value, successful implementation requires attention to several factors.

Platform Engineering and Self-Service Balance

Many organizations adopt a platform engineering model for their batch processing infrastructure. These teams focus on the technology itself—ensuring it's installed, maintained, up to date, and available. But they rely on users to make optimal use of the tool in business contexts.

This separation of concerns is critical: platform teams handle the care and feeding of the technology, while business, data, and application teams own their specific workflows and outcomes. Whether called self-service or another name, this model prevents bottlenecks and enables domain experts to leverage batch capabilities effectively.

Human Intelligence in Automated Systems

The goal of automation is to minimize unnecessary human intervention, but artificial intelligence hasn't replaced human intelligence. Even in the most modern use cases for automation, there's a concept called "human in the loop" where AI might generate several options, then rely on a human to make the final decision.

The optimization isn't about eliminating humans—it's ensuring that human intervention is utilized and leveraged when needed, while avoiding reliance on it when it's not. This is especially true for:

- High-risk decisions with financial or regulatory implications
- Situations requiring business judgment or context that algorithms can't capture
- Exception handling where past patterns don't provide clear guidance
- Validation of AI-generated insights or recommendations

Debugging and Troubleshooting Capabilities

When issues occur, teams need rapid access to relevant knowledge. One of the major areas of evolution today with generative AI and AI assistance is making humans more effective by giving them quicker and better access to both general public knowledge and institutional knowledge embedded within the organization.

The ideal is enabling AI to guide humans to the information they need as quickly as possible—surfacing relevant runbooks, linking to past similar incidents, and proposing potential root causes based on organizational history.

Training and Organizational Change

As with any platform, there's a learning curve involved in managing modern batch systems. Teams need to understand:

- What triggers workflows and how to configure them
- How to define dependencies and execution paths
- What exception notifications mean and how to respond

- How to implement data quality gates and validation rules
- When to apply retries, compensating actions, or manual intervention

Organizations succeed when they invest in training, create clear runbooks, and build internal expertise rather than depending entirely on external consultants.

Do I Need Modern Batch Processing?

If you're wondering whether batch processing is the right approach for your organization, consider where you might apply workflow orchestration in your business operations. Are there gaps you could fill with better automation?

Common use cases include:

- **Payroll processes** and **employee time sheets** with multi-step approvals and calculations
- **Invoicing and billing cycles** that combine usage data, pricing rules, and customer records
- **Financial reconciliation** and **statement generation** requiring complete transaction sets
- **Data analytics and reporting** that aggregate metrics across time periods
- **AI model training and evaluation** with feature engineering and performance tracking
- **Supply chain and fulfillment orchestration** coordinating inventory, shipping, and notifications
- **Compliance and audit reporting** generating required documentation on schedule
- **ETL/ELT pipelines** moving and transforming data across systems

As a rule of thumb, if you find yourself regularly doing large computing jobs manually, or if you have critical processes that require reliable multi-step orchestration, the right batch processing platform could free up significant time and resources for your organization.

Questions to Consider

When deciding if your organization needs to invest in modern batch processing capabilities, ask:

Reliability and Coordination:

- Do you have jobs that must complete in a specific order? How do you ensure they're submitted and processed correctly?
- Do you have work waiting to start that's contingent on other jobs completing? How do you track each job to completion?
- How do you track dependencies across different systems and teams? How do you know dependent systems will be available when needed?

Operational Efficiency:

- Are you manually checking for new files or triggering processes? How frequently, and does it disrupt other work?
- Could your team focus on higher-value work if routine orchestration was automated?
- Do you have retry mechanisms and failure handling in place, or do failures require manual investigation and restart?

Business Service Delivery:

- Are there critical business services (payroll, invoicing, onboarding) that could be more reliable or efficient with better orchestration?
- Do you need to process growing volumes of data without proportionally growing your operations team?
- Are you building AI or machine learning capabilities that require robust data preparation and model training pipelines?

Platform and Self-Service:

- Do you have a platform engineering function to support orchestration infrastructure, or would teams need to build this from scratch?
- Would domain experts benefit from self-service capabilities to define their own workflows without depending on central IT?
- Do you have observability and monitoring in place to detect issues before they impact business services?

Modernize Your Approach to Batch Processing

Modern batch processing is workflow orchestration that runs continuously, handles your most critical business services, and powers cutting-edge AI workloads. It complements stream processing by handling the reliable, correlated, multi-step work that follows real-time detection.

The distinction between batch and stream is about business purpose, not technology fashion. Choose stream when a single event requires immediate action. Choose batch when actions depend on correlating events over time or when completeness, correctness, and throughput matter most. Use both together when you need fast detection and reliable follow-up.

Organizations that succeed with modern batch processing separate platform engineering from self-service usage, leverage AI assistance for troubleshooting, and ensure strategic human intervention where business judgment matters. They focus not on whether batch processing is needed, but on how to enhance automation, improve efficiency, and scale with growing data volumes.