

ORCHESTRATING TRUSTWORTHY AI AGENTS: A PRODUCTION GOVERNANCE PIPELINE WITH CONTROL-M, AMAZON BEDROCK, AND SNOWFLAKE



Learn how Control-M orchestrates a production-grade AI governance pipeline with Amazon Bedrock, Snowflake, DataBrew, SES, and QuickSight — reliably.



Getting an AI agent to produce a compelling response in a playground is straightforward. Getting that same agent to produce reliable, trustworthy output every quarter – on fresh data, in the right sequence, with results delivered to the right people and systems – is an entirely different problem.

The gap between a demo and a production AI workflow usually is not the agent itself. It is everything around it, such as data validation, orchestration, synchronization, and operational visibility.

Agents require validated upstream data before they reason over it. A hallucination caused by a stale governance snapshot or a missing risk indicator is not simply a model problem — it is a pipeline reliability problem. At the same time, downstream systems must act on the agent's output correctly: reports need to be generated, dashboards refreshed, recommendations archived, and governance decisions persisted for auditability.

All of this must happen in the correct sequence, with operational visibility, dependency

management, and failure handling across the entire workflow.

This is the orchestration problem. It is also one of the biggest challenges organizations face when moving AI agents from experimentation into production.

In this post, we show you how to build an end-to-end, production-grade AI governance pipeline using Control-M from BMC for workflow orchestration, Amazon Bedrock Agents for AI-powered portfolio reasoning, and AWS analytics services including AWS DataBrew, Amazon Simple Email Service (Amazon SES), Amazon QuickSight, and Amazon Simple Storage Service (Amazon S3). The pipeline ingests S&P 500 governance data from Snowflake, invokes an Amazon Bedrock Agent to generate strategic portfolio rebalancing recommendations, and orchestrates the complete downstream lifecycle including PDF report generation, email distribution, decision persistence, and dashboard refresh.

What we built and why

This is Phase 3 of a three-part series on operationalizing data and AI workflows with Control-M. [Phase 1](#) covered building a reliable ETL and descriptive analytics pipeline for S&P 500 market data, orchestrated by Control-M and visualized through Amazon QuickSight. [Phase 2](#) extended that foundation into predictive analytics, orchestrated by Control-M to generate return forecasts, volatility estimates, confidence scores, and buy/sell signals.

Phase 3 takes the next step: instead of generating only predictive scores, an AI agent now reasons over current market conditions, evaluates governance signals, and produces portfolio rebalance recommendations that drive downstream operational actions.

The business outcome is a quarterly governance workflow for strategic portfolio rebalancing.

An Amazon Bedrock Agent – the Sentinel-1 Strategic Asset Allocator / Portfolio Rebalancing Agent – ingests governance-ready market intelligence from Snowflake, compares it against the current portfolio allocation, and determines whether market fragility signals such as declining breadth, volatility expansion, and return dispersion justify a more defensive allocation strategy.

When elevated risk conditions are detected, the workflow recommends trimming Technology exposure and redistributing allocations toward Staples, Utilities, and Cash.


Control-M orchestrates the downstream operational lifecycle generating executive PDF governance reports, archiving reports to Amazon S3, distributing reports through Amazon SES, persisting governance decisions back into Snowflake, and refreshing Amazon QuickSight dashboards for operational visibility.

This is what transforms an AI workflow from an isolated model invocation into a production-grade governance platform.

Architecture

AI-Governed Market Intelligence Orchestration Architecture

Control-M Orchestrated Workflow for AI-Powered Portfolio Rebalancing and Governance



The Phase 3 pipeline builds on the Snowflake and S3 data foundation established in the earlier phases. Control-M orchestrates the complete workflow lifecycle across Snowflake, AWS DataBrew, Amazon Bedrock, AWS Lambda, Amazon SES, Amazon S3, and Amazon QuickSight.

At a high level, the orchestration workflow follows this dependency chain:




Figure 1: Phase 3 end-to-end AI governance pipeline orchestrated by Control-M

Each job starts only after its upstream dependency completes successfully. If any stage fails, Control-M halts the workflow before incomplete or unreliable governance outputs reach downstream users.

Two architectural decisions became especially important during implementation.

The first was separating Bedrock reasoning from PDF generation. Earlier versions attempted to pass PDF content directly through the Bedrock response. This caused corruption issues because binary data was being handled as text. The final architecture keeps the Bedrock Agent focused on reasoning and structured JSON generation while Lambda handles PDF generation, chart rendering, S3 archival, and Snowflake persistence.

The second was introducing an S3 watcher as a synchronization gate. Even after Bedrock Flow completed, the generated PDF could still be uploading through the downstream Lambda process.


The Control-M Managed File Transfer (MFT) watcher solved this by waiting until the report physically existed in S3 before releasing the email and dashboard refresh steps.

This transformed the workflow into a deterministic orchestration pipeline instead of a collection of loosely connected cloud services.

SM_BEDROCK

Sub Application: (no value)

Bedrock-portfolio-rebalance



Prerequisites

Before deploying this solution, ensure you have the following in place:

- An active AWS account with administrative access
- Access to Amazon Bedrock with Anthropic Claude 3.5 Sonnet enabled in your AWS Region
- A Snowflake account with the S&P 500 curated governance datasets from Phases 1 and 2
- Control-M environment (v9.0.21+) with the Control-M for AWS integration plug-in installed
- An Amazon S3 bucket for staging governance snapshots and archiving PDF reports
- IAM roles configured for Bedrock, Lambda, S3, SES, and QuickSight access (see Security Considerations)
- Amazon SES configured with verified sender and recipient identities for email distribution
- AWS Lambda functions deployed for PDF generation, S3 archival, and Snowflake persistence (with required Lambda layers for ReportLab and Matplotlib)
- An Amazon QuickSight account with a Snowflake connection configured for governance dashboard visualization

Step 1: Export the governance snapshot from Snowflake

The S&P 500 historical and curated governance datasets reside in Snowflake. Before the Bedrock Agent can reason over the data, the workflow exports governance indicators into a structured JSON snapshot staged in Amazon S3.

The governance context comes from VW_MARKET_TREND_CONTEXT, a curated Snowflake view that provides market breadth indicators, volatility metrics, Bollinger Band width, return dispersion, moving-average strength signals, and governance risk scores.

The workflow exports this data into market_snapshot.json, which becomes the governance context consumed by the Bedrock Agent.

```
COPY INTO @SP500_CURATED.GOVERNANCE.SP500_S3_STAGE/market_snapshot.json
FROM (
  SELECT ARRAY_AGG(
    OBJECT_CONSTRUCT(
      'date', DATE,
      'risk_score', RISK_SCORE,
      'vol_change', VOL_CHANGE,
      'breadth_change', BREADTH_CHANGE,
      'avg_volatility', AVG_VOLATILITY,
      'pct_overbought', PCT_OVERBOUGHT,
      'return_dispersion', RETURN_DISPERSION,
      'dispersion_change', DISPERSION_CHANGE,
      'pct_above_ma30', PCT_ABOVE_MA30,
      'avg_bb_width', AVG_BB_WIDTH
    )
  ) AS market_data
FROM SP500_CURATED.MARKET.VW_MARKET_TREND_CONTEXT
)
FILE_FORMAT = (TYPE = 'JSON' COMPRESSION = NONE)
```

```
OVERWRITE = TRUE  
SINGLE = TRUE;
```

Exporting governance snapshots into Amazon S3 before invoking the Bedrock Agent creates a reproducible AI input artifact. Every workflow run becomes traceable to the exact governance snapshot used during reasoning. It also decouples Bedrock execution from direct database dependencies, improving operational reliability and auditability.

The workflow later writes governance outputs back into Snowflake tables including:

- BF_REPORT_RUNS
- BF_REPORT_ALLOCATIONS
- BEDROCK_DECISION_RESULTS
- DECISION_LOG

These tables provide governance history, workflow metadata, recommendation tracking, and operational auditability for downstream reporting and dashboarding.

Step 2: Validate the snapshot for AI readiness

Before the Bedrock Agent processes the governance snapshot, AWS DataBrew validates the exported data. This becomes the AI-ready data gate within the orchestration pipeline.

The validation layer checks for conditions that could cause the Bedrock Agent to produce misleading governance recommendations:

- Empty governance snapshots
- Missing governance indicators such as risk_score or breadth_change
- Duplicate records
- Structurally invalid numerical values
- Schema inconsistencies

This step becomes one of the most important operational safeguards in the architecture.

AI systems reasoning over incomplete governance data often do not fail obviously. Instead, they produce confident-looking recommendations grounded in unreliable inputs. By validating the snapshot before Bedrock execution begins, Control-M ensures that only trusted governance data reaches the reasoning layer.

If validation fails, the Bedrock workflow never starts.

This is one of the clearest differences between a demo AI workflow and a production AI workflow.

Step 3: Invoke the Bedrock Portfolio Rebalancing Agent

With validated governance data staged in Amazon S3, Control-M invokes the Bedrock workflow.

The workflow uses Amazon Bedrock Flow and the Sentinel-1 Portfolio Rebalancing Agent running on Anthropic Claude 3.5 Sonnet. The agent benchmarks current market conditions against the 2018 stress period and evaluates whether defensive portfolio positioning is warranted.

Model selection

The agent uses Anthropic Claude 3.5 Sonnet on Amazon Bedrock. This model was selected for its:

- Strong performance on structured reasoning and financial analysis tasks
- 200K context window — sufficient for ingesting full governance snapshots without truncation
- Consistent structured output generation with strong JSON schema compliance
- Cost-effective balance between reasoning quality and invocation latency for quarterly governance workloads

Agent configuration

The Sentinel-1 Portfolio Rebalancing Agent is configured with:

- Data context: S3 bucket containing Current_Portfolio.csv and the historical 2018 stress period benchmark snapshot (market_snapshot.json)
- Action Group: ReportArchiver tool receives the structured JSON analysis, generates a PDF governance report with allocation charts, archives it to S3, and returns a secure pre-signed download link
- Instructions: Evaluate volatility expansion, market breadth deterioration, return dispersion, and sector concentration against historical stress periods. Classify market state, produce rebalancing recommendations, and invoke ReportArchiver with the complete structured JSON analysis.

The final workflow prompt used during implementation was:

The Q1 2026 data is staged in S3 (using the file Current_Portfolio.csv). Conduct a fragility benchmark against the 2018 stress period and generate a PDF strategic rebalancing proposal for Q2. I need to see the chart of recommended sector shifts both here and inside the PDF; please render all charts at 100 DPI to ensure the PDF is optimized for archival. Finally, archive that PDF to S3 via the ReportArchiver tool and provide me with the exact secure download link.


The Bedrock Agent evaluates volatility expansion, market breadth deterioration, return dispersion, governance risk scores, sector concentration risk, and similarity to historical stress conditions.

The resulting response contains governance classifications, confidence scores, rebalance recommendations, and executive governance summaries.

Example output:

```
{
  "market_state": "HIGH_RISK",
  "confidence": 0.91,
  "recommended_action": "STRATEGIC_REBALANCE",
  "rebalance_summary": "Trim Technology exposure by 4-5% and redistribute to Staples, Utilities, and Cash.",
  "drivers": ,
  "executive_summary": "Market fragility indicators suggest reducing concentration risk and increasing defensive allocation."
}
```

This becomes the governance decision point within the workflow. The Bedrock Agent is not simply generating a classification score — it is producing a portfolio governance recommendation grounded in multi-factor market analysis.



Step 4: Confirm the report has landed in S3

After Bedrock Flow completes, the generated PDF report may still be uploading through the downstream Lambda process.

This creates a synchronization challenge.

If the SES email job starts immediately after Bedrock execution completes, it could attempt to retrieve a report that has not fully landed in Amazon S3 yet.

Control-M solves this using an MFT watch-only job.

The watcher monitors:

s3://sp500-lambda-bedrock/reports/

for reports matching:

Strategic_Rebalance_*.pdf

The workflow only proceeds once the report physically exists in Amazon S3 as a confirmed, completed governance artifact. This transforms the S3 watcher into a synchronization gate rather than a simple monitoring task. It ensures downstream email distribution and dashboard refreshes always operates.

Amazon S3

Buckets

General purpose buckets

Directory buckets

Table buckets

Vector buckets

Files

File systems

Access management and security

Access Points

Access Points for FSx

Access Grants

IAM Access Analyzer

Storage management and insights

Storage Lens

Batch Operations

Account and organization settings

AWS Marketplace for S3

reports/

Objects

Properties

Objects (36)

Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

Find objects by prefix

	Name	Type	Last modified	Size	Storage class
☐	Strategic_Rebalance_2026-04-30_1835.pdf	pdf	April 30, 2026, 13:35:16 (UTC-05:00)	62.9 KB	Standard
☐	Strategic_Rebalance_2026-04-30_1855.pdf	pdf	April 30, 2026, 13:55:07 (UTC-05:00)	62.4 KB	Standard
☐	Strategic_Rebalance_2026-04-30_2339.pdf	pdf	April 30, 2026, 18:39:48 (UTC-05:00)	62.6 KB	Standard
☐	Strategic_Rebalance_2026-04-30_2352.pdf	pdf	April 30, 2026, 18:52:12 (UTC-05:00)	62.4 KB	Standard
☐	Strategic_Rebalance_2026-04-30_2357.pdf	pdf	April 30, 2026, 18:57:54 (UTC-05:00)	62.7 KB	Standard
☐	Strategic_Rebalance_2026-05-05_1320.pdf	pdf	May 5, 2026, 08:20:50 (UTC-05:00)	62.4 KB	Standard
☐	Strategic_Rebalance_2026-05-05_1505.pdf	pdf	May 5, 2026, 10:05:20 (UTC-05:00)	62.4 KB	Standard
☐	Strategic_Rebalance_2026-05-05_1612.pdf	pdf	May 5, 2026, 11:12:41 (UTC-05:00)	62.6 KB	Standard
☐	Strategic_Rebalance_2026-05-05_1612.pdf	pdf	May 5, 2026, 12:49:53 (UTC-05:00)	62.7 KB	Standard

Strategic Asset Allocation Report

Q2 2026 - Hollow Rally Fragility Analysis

Generated by Sentinel-1 AI - April 30, 2026 18:35 UTC

Executive Summary

Our Q2 2026 Hollow Rally fragility analysis confirms that both key systemic risk thresholds have been breached: market breadth has fallen to 18.2% (below the 20% warning level) and the risk score stands at 0.38 (above the 0.30 critical threshold), signaling a narrow, fragile rally driven by a small cohort of mega-cap technology names. With Information Technology now comprising 34.4% of our \$26B portfolio — a 14.3 percentage-point overweight versus the 2018 benchmark of 20.12% — concentration risk is at an elevated level not seen since the pre-correction period of that cycle. Sentinel-1 recommends an immediate strategic rebalance: trim IT by 4.5% (\$1.12B) and rotate proceeds into defensive sectors (Consumer Staples +2.0%, US9048 +1.3%) and Cash (+1.8%) to reduce fragility and improve downside resilience.

Hollow Rally Analysis

A Hollow Rally is defined as a broad market advance that is disproportionately driven by a narrow subset of large-cap stocks, masking underlying deterioration in market participation. Two quantitative conditions must both be met to declare a Hollow Rally active: (1) Market Breadth below 20%, and (2) a Risk Score above 0.30. As of Q2 2026, both conditions are simultaneously breached. Market breadth currently reads 18.2%, meaning fewer than one in five S&P 500 constituents are meaningfully contributing to index gains — a historically low participation rate that mirrors conditions observed in late 2018 prior to the Q4 correction. The risk score of 0.38 exceeds the 0.30 threshold, reflecting elevated cross-asset volatility, tightening credit spreads divergence, and momentum concentration. Comparing against the 2016 historical benchmark (breadth: 18.2%, risk score: 0.35, IT weight: 20.12%), the current environment is structurally analogous but materially more concentrated: our portfolio's IT weight of 34.4% is 14.3 percentage points above the 2018 benchmark, amplifying the potential drawdown impact should the rally reverse. This analysis confirms a Hollow Rally is ACTIVE and warrants immediate defensive repositioning.

Key Risk Indicators

Indicator	Current Value	Threshold	Status
Market Breadth	18.2%	Less than 20%	BREACHED
Risk Score	0.38	Greater than 0.30	BREACHED
IT Overweight vs 2018	34.4%	20.12% in 2018	DISCREPANT

Indicator	Current Value	Threshold	Status
Market Breadth	18.2%	Less than 20%	BREACHED
Risk Score	0.38	Greater than 0.30	BREACHED
IT Overweight vs 2018	34.4%	20.12% in 2018	DISCREPANT

Step 5: Distribute the report via Amazon SES

Once the report is confirmed in Amazon S3, Control-M invokes the downstream email distribution workflow.

A Lambda function retrieves the latest governance PDF report from Amazon S3 and distributes it through Amazon SES to the configured recipients.

The distributed report includes:

- Governance risk analysis
- Historical stress benchmarking
- Current versus target allocation charts
- Strategic rebalance recommendations
- Executive governance summaries
- AI-generated portfolio rationale

This turns AI reasoning into an operational governance deliverable rather than an isolated analytical result.

The investment committee receives an executive-ready governance report immediately after the workflow completes, without requiring any manual intervention.

Step 6: Refresh the QuickSight governance dashboard

At the same time the governance report is distributed, Control-M refreshes the Amazon QuickSight dashboard connected to the Snowflake governance tables. The dashboard is automatically refreshed after each successful workflow execution, providing operations teams with the latest portfolio analysis and governance metrics.


The dashboard provides an operational view of the portfolio by displaying:

- **Current vs. Target Allocation** by sector, allowing users to compare existing portfolio weights against the AI-recommended allocation.
- **AUM Impact by Sector (\$M)**, highlighting the projected increase or decrease in assets under management for each sector based on the recommended rebalance.
- **Key portfolio risk indicators**, including **Risk Score**, **Market Breadth KPI**, and **IT Weight Percentage**, which summarize the current market conditions used during the portfolio analysis.
- **Run History**, showing each workflow execution, generated report filename, execution timestamp, and key market metrics for historical tracking and auditability.

The dashboard complements the generated PDF report by providing an interactive operational view of portfolio allocation changes, market indicators, and workflow execution history.

While the PDF serves as an executive-ready governance report for quarterly portfolio review meetings, the QuickSight dashboard enables operations teams to monitor portfolio metrics, validate workflow executions, and review historical results.

Together, they transform AI-generated portfolio recommendations into actionable operational intelligence.



Key Architectural Decisions

Decision 1: Separate Bedrock reasoning from PDF generation

Earlier versions of the workflow attempted to pass PDF content directly through the Bedrock Agent response. This caused binary data corruption when handled as text. The final architecture keeps the Bedrock Agent focused exclusively on reasoning and structured JSON generation, while a dedicated Lambda function handles PDF generation, chart rendering, S3 archival, and Snowflake persistence. This separation makes each component independently testable – a critical advantage when iterating on a production governance pipeline.

Decision 2: S3 watcher as a synchronization gate

Even after Bedrock Flow completes, the generated PDF may still be uploading through the downstream Lambda process. Triggering email distribution immediately after Bedrock execution would cause Lambda to attempt to retrieve a report that has not yet physically landed in S3. The Control-M MFT watch-only job solves this by monitoring the S3 reports path for the matching PDF filename pattern and only releasing downstream jobs once the file actually exists. This transforms a race condition into a deterministic orchestration gate.

Decision 3: DataBrew as an AI-ready data gate

AI agents reasoning over incomplete or malformed governance data do not always fail obviously. They produce confident-looking recommendations grounded in unreliable inputs. Placing AWS DataBrew as an explicit AI-ready data validation layer before Bedrock execution ensures that only complete, structurally valid governance snapshots reach the reasoning layer. If validation fails, the workflow halts before the agent is invoked, preventing the pipeline from producing and distributing a governance recommendation based on bad data.

Security Considerations

This solution implements defense-in-depth security across the orchestration pipeline:

IAM least-privilege access

- Bedrock Agent role: Scoped to *bedrock:InvokeAgent* and *bedrock:InvokeModel* for Claude 3.5 Sonnet only
- Lambda execution roles: Separate roles per function with least-privilege access to S3, SES, and Snowflake
- Control-M service role: IAM role with *sts:AssumeRole* permissions, scoped to specific resource ARNs

Data protection

- All data in S3 encrypted at rest using AWS KMS customer-managed keys
- Data in transit encrypted via TLS 1.2+ between all service endpoints
- Snowflake connection uses encrypted JDBC with private link where available
- Pre-signed S3 URLs for report access expire after a configurable TTL (default: 72 hours)

Amazon Bedrock Guardrails

For production deployments, we recommend configuring Amazon Bedrock Guardrails to:

- Filter outputs that do not conform to expected JSON schema
- Block responses containing personally identifiable information (PII)
- Apply content filters to prevent model hallucinations about specific financial instruments
- Log all guardrail interventions for audit purposes

Monitoring and Observability

The solution provides multi-layer observability:

Control-M monitoring

- Real-time workflow execution visibility in the Control-M Web interface

- Dependency-aware alerting: if any upstream step fails, downstream steps are held
- SLA management: configure time-based alerts if the full pipeline exceeds expected duration
- Historical execution reports for trend analysis

AWS CloudWatch integration

- Lambda function logs and metrics (duration, errors, throttles) via Amazon CloudWatch
- Bedrock invocation metrics (latency, token usage, throttling)
- S3 event notifications for object creation tracking
- Custom CloudWatch alarms for anomalous Bedrock response times

Governance audit trail

Every governance decision is persisted to the Snowflake DECISION_LOG table with timestamps, confidence scores, model version, and the input snapshot hash. This provides a traceable audit trail for regulatory compliance and model performance tracking.

What the complete workflow delivers

When the workflow executes at the end of a quarter, the investment committee receives a governance report containing market fragility analysis, historical stress benchmarking, current versus target allocation recommendations, and AI-generated governance summaries.

At the same time, QuickSight dashboards refresh automatically, governance history is persisted into Snowflake, and Control-M captures the workflow lifecycle end-to-end.

No manual intervention is required.

And if anything fails – stale governance data, validation issues, Bedrock quota problems, synchronization failures, or downstream delivery issues – the workflow stops before unreliable outputs reach stakeholders.

That is what production-grade AI orchestration looks like.

Conclusion

In Phase 3, we built an automated AI governance pipeline spanning governance data preparation, AI-ready validation, Bedrock reasoning, PDF report generation, operational notification, dashboard visualization, and Snowflake persistence – all orchestrated reliably through Control-M.

By integrating Amazon Bedrock Agents, Bedrock Flow, AWS DataBrew, Lambda, Amazon SES, Snowflake, Amazon QuickSight, and Control-M, AI governance becomes a dependable operational capability rather than an experimental AI exercise.

The result is a scalable orchestration pattern for enterprise AI governance workflows: agents reason, cloud services execute, and Control-M ensures the entire process runs in the correct order with reliability, visibility, auditability, and operational control.

Try it yourself

To help you get started, we've published the Control-M workflow used in this solution in a public GitHub repository:

<https://github.com/mol-bmc/bedrock-portfolio-rebalance>

The repository currently includes:

- The exported **Control-M job definitions (jobs.json)** for the end-to-end orchestration workflow.
- A **README** with basic information about the project.

These assets provide a starting point for understanding how the workflow is orchestrated with Control-M. You can import and review the workflow definitions, then adapt them to your own AWS, Snowflake, and Amazon Bedrock environment by configuring your own connection profiles, IAM roles, cloud resources, and datasets.

As the project evolves, additional implementation assets and supporting examples may be added to the repository.

Although this implementation demonstrates an AI-governed portfolio rebalancing use case, the orchestration pattern can be applied to many enterprise AI workflows that require trusted data preparation, agent reasoning, downstream automation, monitoring, and auditability.