

Monthly Business Review:

Workflow Transformation Assessment and Future-State Operating Model

Business Transformation Case Study

Shikha Khare · shikhakhare.netlify.app · Version 1.2 · August 2026

Publication note

This document presents an anonymised account of transformation work. Organisational names and identifying details have been generalised. Current-state content draws on direct operational experience. Future-state AI interventions are recommendations and have not been implemented — expected benefits are directional, not measured results.

Case study at a glance

Workflow	Monthly Business Review — infrastructure cost performance
Audience	Director and VP leadership
Stakeholders	25+ across multiple business units
Cycle	18–19 days from financial close to executive review
Primary constraint	Unstructured commentary intake and reactive exception management
Intervention sequence	Redesign and standardise → Automate → AI augmentation
Human-led throughout	Commentary quality validation, risk assessment, resource allocation, sign-off

1. Context and business problem

An MBR process managing infrastructure cost performance across a large, multi-business-unit organisation had scaled in complexity faster than its operating model. The process ran on a fixed 18–19 day cycle from financial close to executive review. Data extraction and variance reporting were largely automated. The problem sat downstream.

The review existed to help leaders evaluate risk, allocate resources, and act on what they saw. In practice, by the time the document reached the executive layer, the analyst team had spent most of their available window on coordination, collation, and reformatting rather than analysis and decision support. The MBR was producing a document. It was not producing the decisions it was designed to support.

The diagnostic question

Before reaching for automation or AI, ask where analyst effort is actually going. If the answer is coordination, collation, and reformatting rather than analysis, the constraint is workflow design — not technology.

2. What was actually broken

The data pipeline was not the bottleneck. The workflow around it was.

Commentary arrived via email, Slack, and Excel with no enforced format and no consistent structure. Different teams submitted at different levels of detail, with different supporting evidence, through different channels. Analysts had to locate, interpret, and reconcile all of it before any validation could begin.

Exception management was reactive rather than proactive. Non-responders were only identifiable at the deadline, by which point the window for follow-up had already compressed. Analysts spent significant time on reminders, chasing, and reformatting explanations that were not suitable for executive consumption.

The result was a structural capacity problem: the activities consuming the most analyst time were the ones that added the least value. The activities requiring the most judgment — quality validation, risk identification, decision support — were left with whatever time remained.

Root cause, operational bottleneck, business consequence

Root cause	Operational bottleneck	Business consequence
No structured intake: commentary arrived in multiple formats across three channels with no required fields	Analysts spent 4–5 days interpreting, validating, and reformatting submissions	Insufficient time for analysis, risk identification, and decision support before executive review
Reactive exception management: non-responders only visible at the deadline	Chasing and escalation consumed the development window	Late submissions compressed the time available for all other activities
Manual reconciliation: every submission checked individually regardless of complexity	Analyst effort applied uniformly rather than where it was most needed	Judgment capacity directed at routine cases rather than exceptions requiring interpretation

The bottleneck was not data extraction — that was already automated

The constraint was how commentary entered the process and how it was managed once it arrived. That is a workflow design problem, not a technology problem.

3. Diagnosis across six dimensions

Evidence basis: direct operational experience managing this process.

The assessment examined the workflow across six dimensions. The finding was consistent: this was a process design problem, not a data or technology gap.

Process: commentary collection had never been standardised. Three channels, no required fields, no consistency criteria — the process accumulated these gaps rather than being designed to prevent them. Exception management was entirely reactive: there was no mechanism to identify or escalate non-responders before the deadline.

Data: the reporting data itself was reliable. The problem was the unstructured business commentary arriving on top of it — multiple formats, variable detail, inconsistent supporting evidence. This was a workflow design gap, not a data quality problem.

Technology: the reporting stack functioned correctly. SQL queries ran on schedule, outputs populated a master workbook, variance tables updated automatically. The technology gap was in commentary collection: email, Slack, and Excel could not enforce completeness or structure.

People and capacity: analyst capacity was absorbed by coordination and reformatting. The judgment-intensive activities — assessing whether an explanation was analytically coherent, identifying risks, evaluating business driver context — were left with whatever time remained after the coordination overhead.

Governance: no structured escalation existed before the deadline. Validation criteria existed as institutional knowledge rather than documented standards. No audit trail connected submission to decision.

Capacity and economics: the 18–19 day window is a fixed organisational constraint. The commercial impact was not cycle length but the compression of judgment capacity within it — analysts spending time on low-value coordination rather than high-value analysis.

4. Intervention decisions

Evidence basis: structured recommendation using the Workflow Transformation Assessment Framework. AI components are proposed and have not been implemented.

Each activity was assessed against the intervention hierarchy: process redesign, standardisation, automation, AI augmentation, or continued human ownership. The sequence is not arbitrary — each intervention creates the conditions that make the next one reliable. AI applied to unstructured inputs does not improve the situation; it produces faster, more confident wrong answers.

01 Redesign Fix the structure	02 Standardise Enforce consistency	03 Automate Remove rule-based work	04 AI augment Interpret at scale	05 Human-led Judgment and accountability
----------------------------------	---------------------------------------	---------------------------------------	-------------------------------------	---

Phase 1 — Redesign and standardise

Activity	Current constraint	Intervention	Decision	Human role
Commentary intake	Multi-channel, inconsistent, no required fields	Structured intake form with required fields: driver, value, context, outlook, evidence	Standardise	Submitter owns the explanation; analyst reviews completeness
Exception management	Manual tracking; non-responders visible only at deadline	Automated three-tier escalation: reminder, management alert, director escalation	Automate	Analyst handles escalated cases requiring judgment

Phase 2 — Rules-based automation

Activity	Current constraint	Intervention	Decision	Human role
Variance reconciliation — standard cases	Every submission checked manually regardless of complexity	Tolerance-based rules engine: auto-clear within threshold, route exceptions to queue	Automate	Analyst reviews exception queue; signs off on reconciled cases
Data extraction and reporting	Already automated and functioning correctly	No change required	No intervention	System-owned throughout

Phase 3 — AI augmentation

AI components: proposed, not implemented. Benefits are directional.

Activity	Current constraint	Intervention	Decision	Human role
Structural completeness check	Manual checklist review of every submission	AI checks required elements against intake criteria; flags gaps for analyst	AI-assisted	Analyst reviews flags; approves or returns
Edge-case reconciliation	Narrative drivers do not map cleanly to numeric variance	AI flags possible gaps; suggests reconciliation paths	AI-assisted	Analyst validates each case; determines follow-up
Commentary quality validation	Contextual judgment: is the explanation analytically coherent?	Existing control retained unchanged	Human-led	Analyst fully accountable. Cannot be encoded.
Executive narrative preparation	Manual synthesis under time pressure	AI drafts from validated structured inputs; analyst reviews and approves	AI-assisted	Mandatory review before distribution to VP/Director audience

5. Future-state workflow

Evidence basis: structured recommendation. Not implemented.

The future-state model adopts an exception-driven approach. Automation and AI handle the initial processing layer. Analysts focus on what the system cannot resolve — cases requiring judgment, context, and

accountability. The fixed 18–19 day calendar is unchanged. What changes is how analyst effort is distributed within it.

01	Data close and variance identification	Automated
02	Structured commentary intake	Redesigned
03	Automated status tracking and three-tier escalation	Automated
04	Rules-based variance reconciliation — standard cases	Automated
05	Exception queue — analyst review of flagged cases	Human judgment
06	AI completeness and edge-case assessment	AI-assisted
07	Commentary quality validation — analyst	Human judgment
08	AI executive narrative draft — analyst review and approval	AI-assisted
09	Lead review and sign-off	Human judgment

6. What changes — current versus future state

Current state	Future state	Business effect
Commentary arrives via email, Slack, Excel — variable format, no required fields	Required-field structured intake form — completeness enforced at submission	Fewer clarification cycles; cleaner data for downstream processing
Non-responders only identifiable at the Day 20 deadline	Automated escalation from Day 14 — three tiers before deadline	Earlier visibility; reduced deadline pressure; fewer late submissions
Every submission checked manually regardless of complexity	Rules engine auto-clears standard cases; analyst reviews exceptions only	Analyst effort directed at cases requiring judgment, not routine checks
Commentary quality validation: informal, inconsistent criteria	Retained under human ownership — documented criteria, consistent application	More consistent quality; clearer accountability
Executive narrative: manual synthesis under time pressure	AI drafts from validated inputs; analyst reviews before distribution	More time for analysis; consistent narrative structure

Current state	Future state	Business effect
No audit trail connecting submission to decision	Structured intake and system-tracked approvals create a permanent record	Accountability and governance visible throughout the cycle

7. Human judgment versus automation versus AI

The boundary between AI and human ownership is not drawn by technology capability. It is drawn by the nature of the judgment required.

What AI does in this model	Why
Checks structural completeness against defined criteria	Pattern matching against a fixed checklist — no contextual interpretation required
Flags edge cases where narrative does not reconcile to number	Volume exceeds practical human capacity; output is a flag, not a decision
Drafts executive summary from validated structured inputs	Templatable synthesis once inputs are clean and consistent

What stays human-led	Why
Commentary quality validation	Requires business context, pattern recognition across months, and understanding of what leadership already knows — cannot be encoded
Exception resolution	Cases the rules engine cannot resolve require practitioner judgment
Risk assessment and escalation	Requires interpretation of business driver context beyond the reported numbers
Resource allocation and investment decisions	Accountability cannot be delegated to a system
Final review and sign-off	Human accountability for what reaches the executive layer

The test of a rigorous assessment

A well-executed assessment identifies activities that should remain human-led. Commentary quality validation is retained by design, not by default. The assessment found that applying AI to this activity would produce outputs requiring more correction than the manual process it replaces.

8. Governance and measures

Evidence basis: structured recommendation. Not implemented.

Governance is designed in from the start rather than appended at handover. A design that depends on the practitioner's continued presence to function has not been properly designed.

Ownership

Role	Owns	Escalation
Central analyst	Workflow execution, exception queue, document development	Central Team Lead
Business unit POC	Accuracy and timeliness of commentary	Business Head
Central Team Lead	Quality review, escalation support, lead sign-off	VP/Director
Technology owner	Workflow configuration, access, monitoring	Central Team Lead

Proposed measures

Baseline values to be established during implementation.

Measure	What it tracks	Threshold
Submission timeliness	% received by Day 17	To be set at baseline
First-submission completeness	% complete on first submission	To be set at baseline
Analyst coordination effort	Hours on chasing vs analysis	Directional improvement expected
Automation coverage	% of standard cases auto-cleared	To be set at baseline
AI override rate	% of AI outputs changed by analyst	Rising rate signals model drift
Decision quality	Leadership time on decisions vs explanation	Directional improvement expected

Conclusion

The MBR's performance constraint was not a technology gap. It was a workflow design gap. Commentary collection, exception management, and validation had never been standardised, and the effort required to compensate for that consumed the analyst capacity that should have gone to analysis and decision support.

The recommended sequence — standardise the intake, automate the rules-based work, then introduce AI to the interpretive tasks — is not a preference. It is a dependency chain. Each phase creates the conditions that make the next one reliable. Applying AI to an unstructured intake process would have accelerated the problem, not resolved it.

Commentary quality validation remains human-led throughout. Not because the technology cannot produce an output — it can — but because the output would require more correction than the manual process it replaces. The judgment stays where the context is.

This case study applies the Workflow Transformation Assessment Framework to a recurring reporting workflow. The hiring operating model case study applies the same framework to an approval and exception workflow. The framework overview is available at shikhakhare.netlify.app.