




HUMAN IN THE LOOP NORMALIZATION FLOW:

THE HIDDEN TIME TAX IN ENTERPRISE
AI, AND HOW HUMAN-IN-THE-LOOP
NORMALIZATION FLOW FIXES IT



Overview

Enterprise AI often appears fast on the surface, yet real work still slows down after the output is generated. The Human-in-the-loop Normalization Flow (HNF), introduces a practical way to measure and reduce the hidden time

spent verifying outputs, navigating approvals, switching tools, and building trust in AI assisted workflows. It provides a minimal, privacy preserving method to make this hidden delay visible and manageable at scale.

The gap everyone feels but few measure

AI can produce answers in seconds. But in many enterprises, the work still takes minutes.

That gap is where a large share of expected AI value quietly disappears. A user receives an AI generated draft, summary, recommendation, or code suggestion quickly, but the task is not complete. The user still needs to review the output, check supporting information, verify policy alignment, gather evidence, wait for approvals, or move across systems before taking action. The model is fast. The work around is not.

This hidden delay is what HNF calls the Hidden Time Tax. It describes the time spent after AI generation and before business completion. In practice, this tax often takes the form of verification time, governance waits, trust formation time, and context switching. Across thousands of tasks, these minutes accumulate into a major operating drag that suppresses throughput and blurs accountability across teams.

The critical issue is that most organizations do not measure this part of the workflow. They see that AI has responded quickly, but they do not see how long it took people to use that output with confidence. That means a major share of the AI experience remains unobserved.

Insights

The main bottleneck in enterprise AI often begins after the answer appears, not while it is being generated.


Why do AI metrics fail to reflect reality?

Most AI programs measure what is easiest to count. They track adoption, prompt volume, output generation speed, number of completions, or number of drafts produced. These are useful indicators, but they do not tell leaders how long real work takes from start to finish.

In many workflows, AI reduces drafting time sharply, yet total task duration remains almost unchanged. The reason is simple. The remaining effort sits in verification, approvals, evidence gathering, and context switching. If a response takes seconds to generate but several minutes to validate, the visible gain is only part of the picture.

WHERE TIME ACTUALLY GOES IN AI ASSISTED WORK

The hidden time after generation often dominates the workflow.



Model response time

Time taken by the AI model to generate the output.



Verify time

Time spent reviewing, validating and confirming the AI generated output.



GateWait

Time lost waiting for approvals, validations or other governance clearances.



SwitchCost

Time spent switching between tools and systems to get the work done.




INSIGHTS

The main bottleneck in enterprise AI often begins after the answer appears, not while it is being generated.


This creates a common disconnect. Teams report higher usage and faster generation yet struggle to explain why throughput has not improved as much as expected. AI looks successful at the surface

level, but the real operating gain remains unclear because the human effort after generation is not being measured.

WHAT MOST TEAMS MEASURE



WHAT DETERMINES REAL PROGRESS



INSIGHTS

If what happens after AI output is not measured, the real value of AI cannot be understood with confidence.

HNF addresses this blind spot by shifting the focus from how fast AI produces output to how quickly that output becomes usable in actual work. That shift is important because it changes the unit of value from model speed to workflow completion.

This is also why HNF is different from generic AI value reporting. It does not stop at AI assistance. It measures what happens next

and whether the workflow becomes faster, smoother, and more dependable.

Insights

If what happens after AI output is not measured, the real value of AI cannot be understood with confidence.

HNF at a glance

HNF is a practical method for making hidden workflow friction visible and improvable. It focuses on what happens after AI produces an output and before the work is finally completed.

HNF uses a minimal, privacy preserving set of signals to reconstruct real workflow behavior. It captures when an output is ready, when it is viewed, and when the work is accepted or completed.

It also observes whether users move outside the primary tool to gather evidence, validate information, or clear governance steps. From this, it calculates where time is going.



SIGNALS

METRICS

OUTCOMES



Output ready

AI produces an output



Viewed

User views the output



Approved / Completed

Work is accepted or task completed



Focus transitions

User moves outside and returns to the workflow



Gate events

Approvals, checks or policy validations occur



FLOW METRICS

VerifyTime	Time from output ready to completion
GateWait	Time lost in approvals and checks
SwitchCost	Time lost in context switching



TRUST METRICS

TrustCurve	How reliance changes over time
TrustTime	Time to reach a trust threshold



GUARDRAIL METRIC

Risk Adjusted Reliance	Reliance weighted by risk sensitivity and governance context
------------------------	--



Visibility

Make hidden friction visible across workflows



Targeted improvement

Focus on the right bottlenecks and act with precision



Safer scale

Improve trust and maintain risk-aware usage



Measurable ROI

Convert minutes saved into real operational and business impact



Privacy by design

Work is accepted or task completed



No Content

We do not capture or store content.



No Prompts

We do not collect prompts.



No Keystrokes

We do not track keystrokes.



No rater burden

We do not ask users to rate outputs.



INSIGHTS

HNF turns invisible human and governance delay into measurable operating signals without creating new burden for users

HNF does this without collecting prompts, screenshots, keystrokes, or raw content. It does not require users to rate or label AI outputs. It is a no rater model built to reduce friction rather than add more process.

HNF then converts these signals into a compact set of measurable indicators that can be governed through service level indicators and service level objectives. These indicators do not just show whether AI is being used.

They show whether AI is reducing human effort, where trust is slow to form, where governance creates delay, and whether faster use remains safe.

Insights

HNF turns invisible human and governance delay into measurable operating signals without creating new burden for users.

Start with what is already there

HNF is designed to begin quickly and with minimal disruption. Organizations do not need full integration or workflow redesign before they start seeing value.

The first phase relies on three simple timestamps. When the AI output is ready. When the user first reviews it. When the work is finally completed or accepted. This is enough to begin reconstructing the flow of the task and measuring how long the human part of the process takes.

If an explicit approval signal does not exist, HNF can infer completion using safe workflow patterns. This may include existing actions such as sending, saving, resolving, submitting, or conservative matching logic between the generated output and the final artifact. This makes it possible to begin in fragmented environments without waiting for a perfect technical setup.

In this phase, HNF also estimates time spent outside the primary application. This helps surface where users pause to gather evidence, check policy, or work through approval steps. Such delays often remain invisible in standard reporting, yet they are central to the true experience of AI assisted work.

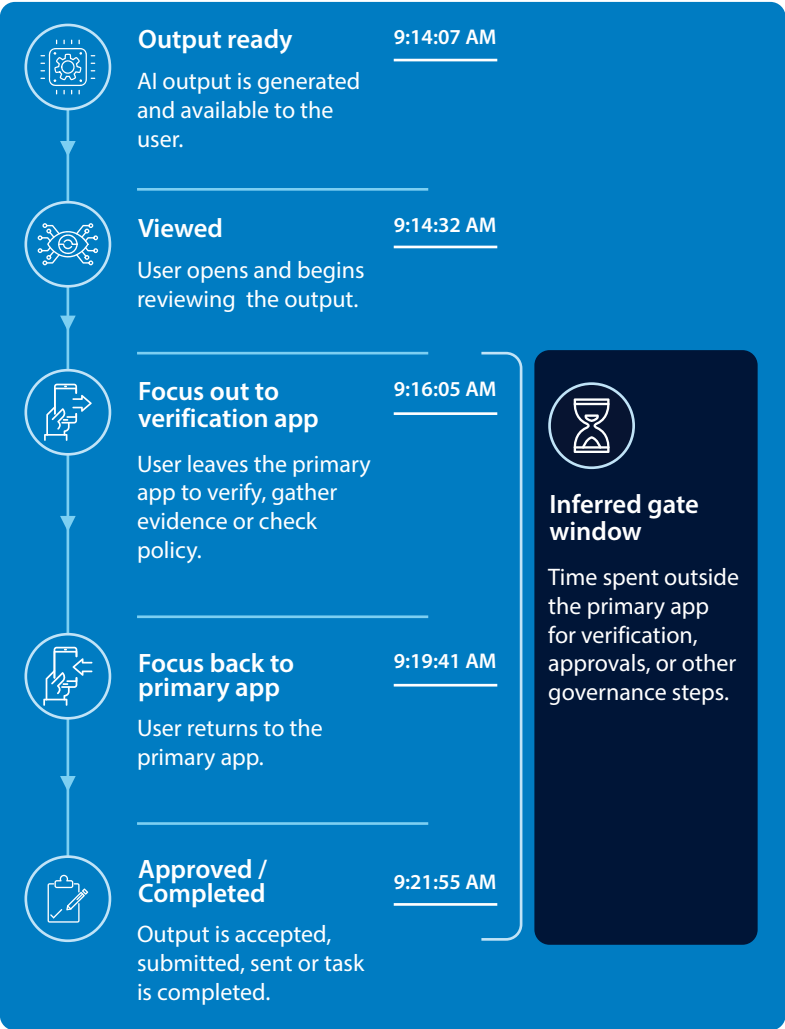
Within roughly two weeks, teams can begin to see where most effort is going. In some workflows, verification dominates. In others, approval waits or switching costs become the main bottleneck. The value of this phase is not precision first. It is visibility first.

Insights

You do not need a heavy integration model to begin measuring meaningful workflow friction.

START WITH WHAT IS ALREADY THERE

Phase 1 uses minimal signals to make hidden workflow friction visible.



VerifyTime	GateWait (linferred)	SwitchCost (basic)
Reveals how long it takes from output ready to approval or completion	Reveals how long it takes from output ready to approval or completion	Reveals time lost when users leave the primary app for verification or context

Value appears quickly with just three timestamps and an inferred gate window.

Privacy by design

- No content captured
- No prompts captured
- No keystrokes captured
- No workflow redesign

INSIGHTS

You do not need a heavy integration model to begin measuring meaningful workflow friction.

From visibility to proof

Once friction becomes visible, the next step is to validate what reduces it. This is where HNF moves from directional insight to stronger measurement confidence.

Phase 2


FROM VISIBILITY TO PROOF

Replace key inferences with trusted signals. Test one change. Prove real impact

Why this phase matters

- Directional insight is powerful, but proof drives change.
- Trusted signals reduce guesswork and increase confidence.
- One focused experiment reveals what truly works.

Phase 2 strengthens measurement confidence. We introduce a small set of trusted workflow signals, make one gate explicit, and test one improvement at a time to prove what reduces hidden effort.



What we add in phase 2	What improves	Example focused improvements
One reliable completion signal Use trusted events like ticket closure, merge, submit, send, or status transition to replace inferred completion	Higher measurement confidence Less reliance on inference and stronger reliance on real workflow events.	Add inline evidence Provide source links or citations to reduce verification effort.
One explicit gate signal Capture a real governance or approval event (for example, DLP check, compliance step, policy approval) to replace inferred waits.	Clear attribution Isolate the impact of one change at a time with A/B or controlled testing.	Simplify approvals Clarify policies, reduce back-and-forth, or automate checks.
Early reliance detection Measure how much AI output is used with minimal rework - behavior based, not survey based.	Evidence of true effort reduction See real changes in VerifyTime, GateWait, SwitchCost, and trust metrics.	Reduce context switching Surface information in flow, integrate tools, and remove duplicate steps.

Experiment lifecycle: test one change at a time



Form hypothesis

What part of the workflow is creating the most tax?



Run controlled test

Introduce the change to one group and hold others constant.



Measure impact

Compare metrics before and after the change.



Adopt or iterate

Compare metrics before and after the change.



INSIGHTS

Phase 2 moves HNF from visibility to proof. By replacing key inferences with trusted signals and testing focused improvements, teams can confidently reduce hidden effort.



Phase 3

FROM PROOF TO OPERATING MODEL

Make HNF part of how teams run and improve work

As measurement matures, HNF evolves from experiment to operations. We expand precision, set targets and establish governance so teams can manage flow, trust and safety at scale.

What changes in Phase 3

- 

Stronger precision
More explicit signals across workflow improve accuracy of VerifyTime, GateWait and SwitchCost.
- 

Deeper governance insight
Break down GateWait by gate type (policy, compliance, manual review) to address the right bottlenecks.
- 

Trust and safety by design
Reliance 2.0 and Risk Adjusted Reliance ensure speed improvements never compromise safe use.
- 


Managed by targets
Service Level indicators (SLIs and Service Level Objectives (SLOs) provide clear threshold and early alerts.
- 

Ownership and rhythm
Product, Operations and Risk share accountability with a cadence of review, action and follow up.

HNF Executive Scorecard (HBS)


VerifyTime (p95)	GateWait (p95)	SwitchCost (p95)	TrustTime (p50)	Risk Adjusted Reliance
				
4.6 min	3.1 min	1.7 min	5.2 days	0.82
-22%	-18%	-20%	-27%	(+0.07)

HBS Composite Score



▲ 9pts vs last 30 days


Trends (Last 12 weeks)



SLO Status

VerifyTime (p95)	GateWait (p95)	SwitchCost (p95)	TrustTime (p50)	Risk Adjusted Reliance
< 6 mins	< 4 mins	< 2 mins	< 10 days	> 0.75
✓ On Track	✓ On Track	✓ On Track	✓ On Track	✓ On Track

The HNF Operating Loop





Product Teams

- Improve experience and flow
- Reduce VerifyTime and SwitchCost
- Design and deliver interventions



Operations Teams

- Monitor performance and SLOs
- Drive adoption and process changes
- Ensure operational stability



Risk & Compliance

- Define guardrails and gates
- Monitor Risk Adjusted Reliance
- Ensure safe and responsible scale



INSIGHTS

HNF becomes a true operating model when it is measured, governed and improved in a continuous loop. From insight to action. From action to impact.

The metric system that matters

HNF uses a compact set of measures, but they become easier to understand when grouped by purpose.

Flow metrics

These metrics show how work moves.

VerifyTime measures how long it takes to move from AI output being ready to the work being accepted or completed. It captures the human verification loop directly.

GateWait measures delay caused by governance requirements such as approvals, compliance checks, or validation steps.

SwitchCost measures time lost when users leave the primary tool to gather context or verify information elsewhere.

Trust metrics

These metrics show how confidence forms over time.

TrustCurve describes how reliance changes as users gain repeated exposure to AI outputs in a given workflow.

TrustTime reflects how long it takes to reach a meaningful level of trust or reliance. These are behavior-based signals, not survey based impressions.

Guardrail metric


This metric ensures that speed gain remains safe.

Risk Adjusted Reliance connects usage to risk sensitivity. It helps ensure that faster or more frequent use of AI does not come at the expense of reliability, governance, or safe judgment.

These measures together create a more balanced view of AI performance. They move the conversation from surface activity to workflow behavior, trust, and operating safety.

Insights

HNF measures not only how work flows, but also how trust forms and whether improvement remains safe.



THE METRIC SYSTEM THAT MATTERS

A compact set of metrics to measure flow, trust, and safe AI adoption.

HNF translates minimal workflow signals into a balanced set of metrics that capture how work moves, how trust develops and whether faster use remains safe. Together, these metrics help teams diagnose friction, validate improvements and manage AI adoption with confidence.

FLOW METRICS

Measure how work moves from AI output to completion.



VerifyTime

Time from output ready to completion.



GateWait

Delay due to approvals and checks.



SwitchCost

Time lost in context switching.

TRUST METRICS

Measure how confidence in AI builds over time.



TrustCurve

How reliance builds with use.



TrustTime

Time to reach a trust threshold.

GUARDRAIL METRIC

Ensure faster use of AI remains safe and responsible.



Risk Adjusted Reliance

Usage adjusted for risk and governance sensitivity.

Not just speed –
measure how work
flows, how trust
forms, and whether
it remains safe.



End-to-end visibility



Actionable insights



Measurable improvement



Safe and responsible scale



INSIGHTS

HNF metrics move beyond activity to reveal real workflow behavior.



How HNF creates business value

HNF matters because it helps organizations focus on the delays that suppress throughput. Instead of improving AI outputs in isolation, teams can identify whether the real constraint is verification, approval, or fragmented context gathering. That shifts effort from generalized optimization to targeted workflow improvement.

Even small reductions at the task level can become meaningful at scale. If teams save a minute of verification effort in a high-volume workflow, that gain compounds across hundreds or thousands of tasks. What appears small in one interaction becomes operationally significant across a business process.

HNF also improves decision quality. It helps teams validate which

interventions are genuinely reducing effort and which are not. This keeps AI investments grounded in evidence. Leaders can prioritize changes based on real workflow outcomes rather than relying on assumptions or anecdotal success.

In this way, HNF transforms AI adoption from a question of visible usage into a question of measurable workflow gain. It creates a more credible connection between AI investment, operational improvement, and sustained value.

Insights

AI creates value when hidden effort is reduced in measurable, repeatable ways.

HOW HNF CREATES BUSINESS VALUE

From hidden friction to measurable impact



HNF helps organizations focus on the delay that suppress throughput. Instead of improving AI outputs in isolation, teams can identify whether the real constraint is verification, approval, or fragmented context gathering. That shifts effort from generated optimization to targeted workflow improvement.



Even small reductions at the task level can become meaningful at scale. A minute saved per task in a high-volume workflow compounds across hundreds and thousands of tasks, creating significant operational gains.



HNF improves decision quality. It helps teams validate which interventions are genuinely reducing effort and which are not. This keeps AI investments grounded in evidence and enables leaders to prioritize changes with the highest workflow impact.




Example: Service desk response drafting

Before HNF

6:00 mins average
end-to-end time

After HNF Intervention

3:15 mins average
end-to-end time



Impact



~46% reduction in total;
task time



Capacity released at
scale



Higher CSAT and first-time
reduction



INSIGHTS

AI creates value when hidden effort is reduced in measurable, repeatable ways.



How do teams work with HNF?

HNF is designed to be used, reviewed, and improved regularly. It is not a reporting artifact. It is a working control surface for AI enabled workflows.

Teams use dashboard views to understand where friction is concentrated, which metrics are moving in the right direction, and where targeted action is needed. Operational reviews are deliberately simple. Trends are reviewed. Guardrails are checked. One or two high value actions are selected. Impact is assessed in the next cycle. This keeps the HNF grounded in action rather than observation alone.

Leadership uses a smaller set of views and targets. Instead of seeing raw telemetry, leaders see how flow, trust, and safe reliance are changing across important workflows. This enables

better prioritization and gives governance teams confidence that improvement is happening without weakening control.

The strength of HNF lies in discipline. Measurement is useful only when it becomes part of a routine decision loop. That is why the HNF is designed to fit within existing operating rhythms across product, operations, and risk.

Insights


HNF creates value when it becomes part of how teams review and improve work every week.

Phase 3

HOW TEAMS WORK WITH IT

Turn measurement into action through simple dashboards and weekly review cycles.

HNF translates minimal workflow signals into a balanced set of metrics that capture how work moves, how trust develops and whether faster use remains safe. Together, these metrics help teams diagnose friction, validate improvements and manage Ai adoption with confidence.



How teams operate



Product

Improve flow and experience by reducing friction in AI-assisted workflows.



Operations

Monitor performance and cadence to ensure consistent improvement and throughput gains.



Risk

Maintain guardrails and ensure safe, responsible and compliant AI use.



INSIGHTS

HNF creates value when measurement leads to structured, repeatable action.



Practical scenarios

HNF applies wherever AI output is followed by human checking, waiting, or trust building. That makes it relevant across a wide range of enterprise workflows.

Service desk drafting

A response may be generated quickly, but agents still verify details, consult past cases, or wait for policy checks before resolving the ticket. HNF reveals whether the main delay sits in review, in governance, or in switching between systems. Once visible, those delays can be reduced more precisely.

Change management

AI can help prepare change records quickly, yet completion remains slow because approvals, evidence gathering, and scheduling windows still dominate the process. HNF separates drafting speed from completion speed and helps teams focus on the true bottleneck.

Document and email review

A draft may appear instantly, but users still check wording, compliance, supporting information, and prior context before acting. HNF helps organizations understand how much hidden effort sits in that review process and where experience changes can make the workflow smoother.

Across these workflows, the pattern is the same. The delay is not in generating an answer. It is in making that answer dependable enough to use.

Insights

HNF is useful wherever AI output must become trusted, usable work.



PRACTICAL SCENARIOS

HNF applies across workflows when AI output is followed by human checking, waiting, or trust building

Different workflows. Same pattern. The real delay is rarely in generation. It's in making the output dependable enough to use.



Service desk drafting

AI drafts the response, but agents still verify, check policy, and resolve.



AI output generated (seconds)



Agent reviews the output



Checks policy / past cases (focus out)



Waits for governance (if any)



Finalizes and resolves ticket

Where hidden time occurs

- Detailed verification
- Searching prior tickets / knowledge
- Policy or compliance check
- Waiting for approval (if required)
- Switching across tools

What HNF reveals

- High VerifyTime due to multi-stop review
- Significant SwitchCost across systems
- GateWait is policy or approval steps

How teams improve

- Inline knowledge and context
- Better recommendation with references
- Policy checks earlier in flow
- Reduced tool switching
- Faster, safer resolution



Change Management

AI prepares change records quickly, but approvals and evidence gathering slow completion.



Where hidden time occurs

- Collecting supporting evidence
- Multiple approval layers
- CAB / change window wait time
- Rework after review
- Tool switching for validation

What HNF reveals

- GateWait dominates total task time
- High SwitchCost during evidence collection
- Long end-to-end time despite fast drafting

How teams improve

- Standardized evidence templates
- Parallel approvals where possible
- Better scheduling visibility
- Early policy validation
- Shorter cycle, fewer delays



Document and email review

AI creates drafts instantly, but review, compliance and context checks take time.



Where hidden time occurs

- Detailed content verification
- Compliance and accuracy check
- Checking supporting information
- Rework and refinement
- Switching across documents/ tools

What HNF reveals

- High VerifyTime due to deep review
- SwitchCost across multiple sources
- Hidden effort after AI generation

How teams improve

- Citation and source visibility
- Embedded compliance checks
- Better summarization with context
- Reduced switching
- Faster, confident outcomes



INSIGHTS

HNF is useful wherever AI output must become trusted, usable work.

Conclusion

Enterprise AI delivers real value only when human time is treated as a first-class operating constraint. The biggest delays in AI assisted work often occur after the answer is generated. Verification, approvals, context switching, and trust formation define how quickly work moves.

HNF addresses this directly. It measures what was previously invisible and turns it into actionable, governable signals. The method is deliberately practical. Start with minimal signals. Make hidden friction visible. Improve confidence through better workflow events. Test one change at a time. Scale what works with clear ownership and guardrails.


Within Infosys Topaz, **HNF** provides the measurement layer that helps AI capability translate into real workflow acceleration. The point is not simply to make AI look fast. The point is to make work itself move faster, with evidence teams can trust.

Final takeaway:

AI is fast. HNF helps organizations ensure that work becomes faster too.



About the Authors



Nitin Mahendru

Sr. Project Manager, Infosys Topaz Impact Team



Harini Kumar Goud Sesam

Principal Consultant, Infosys Topaz Impact Team

Infosys Topaz is an AI-first set of services, solutions and platforms using generative AI technologies. It amplifies the potential of humans, enterprises, and communities to create value. With 12,000+ AI assets, 150+ pre-trained AI models, 10+ AI platforms steered by AI-first specialists and data strategists, and a 'responsible by design' approach, Infosys Topaz helps enterprises accelerate growth, unlock efficiencies at scale and build connected ecosystems. Connect with us at infosystopaz@infosys.com.

For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.