



HOW AGENTIC AI IS REDEFINING PUBLISHING

Executive Summary

The press, the web, the algorithm – each reshaped how journalism reaches people. Agentic AI reshapes something closer to the source: how decisions are made, who makes them, and what it means to be accountable for what reaches an audience.

Agentic AI is not the next feature update for publishing. It is a structural shift in where decisions get made – and who, or what, is accountable for them. This document traces that shift: from the forces driving it, through what governance looks like when it holds, to the coordination challenges that governance frameworks have not yet caught up with.

From Generative Tools to Agentic Systems

"AI is no longer simply a productivity enhancer. It is moving from assistant to intelligent orchestrator – weaving together human judgment and machine execution across the full arc of content planning, creation, distribution, and monetisation."

The first wave of AI in publishing was additive. Newsrooms adopted headline generators, summarisation tools, and writing assistants – valuable at the margins, but subordinate to existing workflow logic. Editors still edited. Journalists still filed. The tools accelerated what humans were already doing.

What is emerging is of a different order. Agentic AI does not

assist – it acts. It perceives context, pursues goals, and executes multi-step workflows with limited human intervention at each step. For publishing organisations, that distinction has material consequences downstream: how editorial work is structured, where human judgment is required, and whether the people responsible for that judgment are directing the system or simply operating within it.

The Shifts That Matter

The publishing industry has navigated technological disruption before. What makes this moment different is not the pace of change – it is the depth of it. Previous waves reshaped distribution, monetisation, and format.

Agentic AI reaches further: into editorial decision-making, institutional memory, and the governance of content itself. Understanding what that means requires looking at the forces driving it with some precision.

Four shifts stand out.

01

From features to orchestrated systems

Isolated AI tools have reached their ceiling. The value frontier has shifted to orchestrated agentic systems that act across the full content lifecycle – not within individual tasks.

02

Context, continuity & institutional memory

Publishing is inherently continuity-driven. AI systems that reset with every session cannot serve editorial workflows built on accumulated beat knowledge, source relationships, and evolving editorial positions.

03

Trust, provenance & explainability

As AI-generated content proliferates across the open web, publisher credibility rests on auditable provenance. AI that cannot explain or trace its outputs is not a productivity tool – it is a liability.

04

Convergence across editorial, audience & revenue

The traditional separation between editorial, audience development, and revenue operations is dissolving. Agentic systems that serve all three simultaneously require unified data, workflow, and governance architectures.

WHY THIS MATTERS

These shifts are mutually reinforcing. A publisher that responds to each individually – adding a memory-capable AI here, a provenance layer there – produces an architecture of patches. The response to four structural shifts is a structural redesign: an operating model built for agentic AI from the ground up, not retrofitted around it.

The Pyramid, Inverted

One place those structural forces are most visibly felt is inside the newsroom – in how editorial time is allocated, and how much of it ends up in the wrong place.

Every newsroom has an implicit value hierarchy. At the apex: original, investigative, deeply contextual journalism – the work that builds trust, commands subscriptions, and justifies the institution's existence. In the middle: analysis, explainers, continuity coverage. At the base: time-sensitive, routine content – earnings summaries, weather, sport

results, wire rewrites, the relentless churn of breaking updates.

In theory, editorial effort is weighted toward the apex. In practice, most newsrooms have unintentionally inverted the pyramid. Journalists spend the majority of their time at the base – not because editors have chosen this, but because the base is loud, urgent, and insatiable. Routine coverage demands constant attention. The structure of the day is shaped by what is time-sensitive, not what is high-value.




Fig. 1 – The newsroom content pyramid: today's inverted reality vs the agentic-enabled ideal

DESIGN, NOT DEFAULT

Restoring the pyramid is a design decision, not a consequence of deployment. The question is not how many journalists a publisher employs but how editorial work is allocated between humans and machines – and whether that allocation is governed by deliberate intent or left to default. Reclaimed capacity can flow into deeper, original journalism or out of the cost base entirely. The technology does not determine which. The publisher does, through the choices that follow deployment, not the ones that precede it.

Restoring the Pyramid

What agentic AI offers the newsroom is reallocation – the ability to systematically redirect editorial attention from the work that fills the day to the work that defines the publication. Reallocation is a capability, not an outcome: agents handle the structured, time-sensitive, repeatable work at the base; humans concentrate on the work that requires judgment, context, and accountability. What changes is not the editorial mission. What changes is how much of the working day is spent serving it.

Agentic systems have already demonstrated production-grade capability in constrained, structured contexts. The Associated Press and Reuters have deployed automated writing for financial

results and sports summaries at scale – output that is factually reliable, produced in seconds, and requires only light editorial oversight.¹ These are not experiments; they are operational realities at major news organisations.

The broader pattern is consistent: signal monitoring, routine drafting, metadata tagging, story continuity, workflow orchestration – tasks that fill the working day without advancing the editorial mission. An agent drafting within approved templates is carrying out editorial policy, not making editorial decisions. What returns to the journalist is attention: for investigation, source cultivation, the contextual understanding that turns a data point into a story, and the ethical judgment that decides what a story should and should not say.

CORE PRINCIPLE

Humans define editorial intent, set escalation thresholds, and retain accountability for every piece of content that reaches an audience. Agentic AI executes within those parameters, surfaces exceptions, and routes decisions to the right human at the right moment. The chain of editorial responsibility does not change – the distribution of workload does.


Designing for Coherence

Redistributing workload is one thing. Sustaining it at scale – across teams, beats, and content types – requires a different kind of response. Individual tools, however capable, do not add up to a coherent architecture. A publisher deploying a signal-monitoring agent in one team and a drafting tool in another, with no shared governance model or escalation

protocol, produces complexity without coordination and risk without oversight.

What forward-thinking publishers are beginning to design instead is a coherent editorial architecture organised across three dimensions of publishing value – each carrying a specific governance obligation:

CUSTOMER EXPERIENCE (CX)
How stories reach the right reader without the reader's feed becoming the newsroom's actual editor.
LEFT UNGOVERNED
An unguided personalisation agent can quietly narrow what a democracy's information diet looks like, one feed at a time – a failure mode with no analogue in retail personalisation.
REVENUE (RX)
How content generates sustainable commercial returns without yield becoming the signal that shapes coverage.
LEFT UNGOVERNED
A monetisation agent optimising yield structurally favours outrage and virality over accuracy, reintroducing at machine speed the exact pathology that reader-trust journalism has spent a decade trying to escape.
OPERATIONS (OX)
How workflows are orchestrated efficiently without speed becoming the criterion by which editorial decisions get made.
LEFT UNGOVERNED
An operations agent optimising for throughput will, without editorial governance, treat a correction and a retraction as equivalent production events – collapsing a distinction that carries legal and ethical weight no other content industry has to carry.



A Day in the Agentic Newsroom

Abstractions only carry so much weight. A vignette is more honest.

Consider a mid-size regional publisher on a typical morning.




Fig. 3 – Agent and human activity across a morning shift (illustrative)

ILLUSTRATIVE – A MID-SIZE REGIONAL PUBLISHER, 6AM

By the time the first editor arrives, the overnight work is already done. Agents have processed wire traffic, monitored financial filings across four key beats, and assembled a **triage board** – not an inbox – showing what is editorially significant, what is time-sensitive, and what has been escalated for review. The escalated item comes with a briefing note and source links already assembled. She does not need to go looking for context.

Routine content – earnings summaries, a weather update, sport results – has been drafted within pre-approved templates and queued for review. **The review is real: low-confidence items and anything touching a named individual are separated out and cannot be cleared alongside routine copy.** Each confirmation is logged. Less of the morning shift is spent on wire rewrites than it once was. More of it is spent on judgment.

One story has been escalated not because it is urgent but because the agent encountered a factual claim it could not verify against the archive and routed it rather than proceeding. **That is the system working as designed.** Separately, the analytics agent has surfaced an overnight signal – a three-day-old story driving unusual subscriber conversion – and notified the commercial desk automatically. No one had to look for it.

By 10am, two journalists are conducting interviews for a long-running investigation. One is working on a data explainer pre-populated with sourced statistics, each verified against its original source before use. One is editing the escalated story. It is the only piece today that required genuine editorial judgment to resolve. **The pyramid is, for once, the right way up.**

This morning is not a projection. The individual capabilities it describes – automated drafting, triage routing, escalation logic, analytics surfacing – are operational in newsrooms today. What is not yet common is the governance architecture that holds them together: the decisions, made

in advance, about what agents are permitted to do, where human judgment takes over, and what happens when the three pillars pull in different directions. That is what determines whether a morning like this is repeatable – or whether it unravels the first time the system is tested.

Humans in the Lead

The greatest risk in agentic AI adoption is not that systems will fail. It is that they will succeed – autonomously, at scale, without the governance architecture to account for what they produce. Publishers who discover that gap after the fact will discover it through damage.

Humans in the Lead is the design principle that prevents this. It is not a policy declaration – it is an architectural commitment: human judgment embedded in agentic systems not as a last-resort override, but as the primary source of editorial intent, ethical standards, and escalation logic that defines the framework within which AI systems are permitted to act.

This principle has both an ethical and a commercial dimension – and they are inseparable. In an environment where AI-generated content proliferates across the open web, publisher credibility rests not on volume but on verifiable human editorial oversight. That credibility has an internal face – the governance that holds – and an external one: what publishers tell their audiences about how AI was involved in what they read. Concealment is not a viable posture. The publications that audiences will trust most are those that can demonstrate, not just assert, that a human was accountable – clear about what AI did, what a human verified, and where the line between them falls. In practice, that accountability means something specific for each role in the editorial chain.

EDITORS

Set the guardrails, escalation triggers, and content standards within which AI operates. Agents execute editorial policy; editors set it.

JOURNALISTS

Retain narrative ownership – source relationships, interpretation of evidence, and the framing of every story that requires contextual or ethical judgment. Not all stories. All the ones that matter.

AI SYSTEMS

Operate within governed, auditable boundaries. Every consequential output is traceable. Provenance is the architecture.

A NECESSARY QUALIFICATION

"Human in the loop" is easy to assert and hard to sustain. Under deadline and volume pressure, the reality is simple: a reviewer facing hundreds of agent-drafted items will clear them – until oversight has quietly degraded into a signature.

*"Publishing of some information will be automated and carried out by agents, and stories will be automatically updated. The much-discussed 'human in the loop' might be quietly retired. Not everyone views this as a positive development."*²

The data supports the concern. The Reuters Institute's Journalism, Media and Technology Trends and Predictions 2026³ – drawn from 280 senior newsroom executives across 51 countries – found that only 44% of publishers have established protocols around human oversight and control of AI. More than half do not. A principle asserted without protocols to enforce it is precisely the governance model that does not confront this. It is not a safeguard; it is a label.

Confronting it requires design, not assertion. Oversight holds only when it is bounded by what a human can actually attend to – when review capacity is treated as a finite resource that limits how much an agent is permitted to queue for human clearance, when sensitive categories are structurally separated and cannot be bulk-approved, and when the system is audited not on whether a sign-off occurred but on whether it was real. The question to ask of any governance model is not whether humans are nominally in the loop. It is whether the loop is designed to remain meaningful when the volume is high and the deadline is close.

Even when oversight holds, questions remain. Agentic systems will produce errors – a confidently wrong confidence score, a fabricated detail, a verified-looking

claim that is not true – and the question of who is liable when that error reaches an audience does not disappear because the architecture is sound.

What happens when agents act correctly, oversight holds, and the system still produces an outcome no one authorised?

The Coordination Problem

The governance frameworks described so far address how individual agents are directed, bounded, and held accountable. There is a harder problem: what happens when agents act correctly, oversight holds, and the system

still produces an outcome no one authorised? This is the multi-agent coordination problem, and it is where even well-designed governance frameworks reach their limits.

The problem is not what each agent does. It is what they do together.

ILLUSTRATIVE – MULTI-AGENT COORDINATION FAILURE

A central bank signals an unexpected rate hold. The **editorial agent** flags it as high-significance, a human editor clears it, and it publishes within guardrails. The **personalisation agent** identifies that financially-engaged subscribers convert strongly on rate content and begins surfacing it prominently across that segment. The revenue agent simultaneously detects a pre-booked financial services placement and prioritises distribution to match. Neither agent has violated a rule. Neither is governed in relation to the other. The editor who cleared the story had no sight of either decision. **What reaches a financially-exposed audience is now shaped by commercial and personalisation logic the editorial process never contemplated – market-sensitive content amplified to investors, alongside financial services advertising, without editorial authorisation of that targeting.** Every agent did its job. The coordination between them did something no one authorised.

Every agent followed its rules. The failure was in the combination. Explainability – the industry's instinctive response – can describe what each agent did and why. It cannot tell you whether the coordination was legitimate, or who was accountable for what emerged from it. The audit trail proves compliance. It doesn't prove governance.

Publishers adopting multi-agent systems must treat inter-agent coordination logic as itself requiring editorial governance, not just the individual agents. This means pre-

defining – and where possible encoding – what editorial authority looks like when agents disagree. It is harder than it sounds. Most publishers haven't done it, in part because the tooling is immature, but more often because the answer given at policy level – editorial owns it – is not yet consistently reflected in the architectural decisions being made by the technology teams building these systems. Two approaches are emerging from the most considered deployments.

TWO APPROACHES EMERGING

Neither treats coordination as an afterthought. Each makes it a first-class governance object – though by different means and with different trade-offs.

Orchestration governance treats every inter-agent handoff as a discrete governance event – logging the reasoning state, the rule invoked, and the authority applied – so the coordination chain is replayable rather than reconstructed. It adds latency and works well for structured, high-stakes content; it has not yet been applied at full daily output volume. **Compositional governance** embeds editorial intent as structured metadata that travels with content – sensitivity scores, distribution constraints, escalation triggers – which downstream agents must acknowledge before acting. Its limitation: metadata can be stripped when content crosses to third-party platforms that do not support the schema, leaving governance constraints behind at the system boundary.

Both are early-stage. The harder question – how coordination governance scales across hundreds of simultaneous agent interactions in real time – remains open.

The governance layer has not caught up to the architecture. The publishers responding well are not those claiming it

has – they are those designing precisely for what it cannot yet cover.

The Window

Publishing leaders are not waiting to see how agentic AI reshapes their operations. The deployment decisions are

being made now – many without governance frameworks adequate to the scale of what is being released.

37%

of people globally trust the news⁴ – the lowest figure recorded since the Reuters Institute began measuring trust in 2015, falling in 29 of 48 markets surveyed. Trust in news on AI chatbots stands at 20%.

Three in four publishers expect agentic tools to have large or very large impact on the industry.³ Fewer than two in five are confident about journalism's overall prospects. These numbers sit alongside a trust figure that should give every publisher pause: audiences are extending their scepticism faster than publishers are building their governance. The architecture decisions made in the next 12 to 18

months – about how agentic AI is integrated, directed, and held accountable – will be structurally consequential. Organisations that treat this as a procurement cycle will inherit architectures that are difficult to audit, difficult to trust, and difficult to change. The window for deliberate design is not indefinite.

The future of publishing will not be defined by how much content AI can generate – but by how intentionally it is designed to serve human judgment.

Agentic AI reaches into the decisions that define what a publication stands for: what to cover, how to frame it, and who is accountable for what reaches an audience. The challenge ahead is not simply governing individual AI tools. It is ensuring that editorial judgment remains coherent across an ecosystem of agents acting throughout the publishing lifecycle. In that world, editorial intent cannot sit at a single approval point. It must be designed into the system itself – a shared space where editors and agents shape the same story together, with judgment and accountability built in, not bolted on. Neither policy nor oversight alone is sufficient. The future of trustworthy publishing will depend on architecture.





References

- 1** | **AP Automated Insights.** The Associated Press has deployed automated writing for financial results since 2014 and sports summaries at scale in partnership with Automated Insights; Reuters operates equivalent automation for financial data alerts via its FactGenie tool. ap.org/solutions/artificial-intelligence/automated-insights
- 2** | **Reuters Institute – How will AI reshape the news in 2026? Forecasts by 17 experts from around the world.** January 2026. reutersinstitute.politics.ox.ac.uk/news/how-will-ai-reshape-news-2026-forecasts-17-experts-around-world
- 3** | **Reuters Institute – Journalism, Media and Technology Trends and Predictions 2026.** Nic Newman. Survey of 280 senior news executives across 51 countries, January 2026. 75% expect large/very large impact from agentic tools; 38% confident about journalism's overall prospects; 44% have human oversight and control protocols in place. reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2026
- 4** | **Reuters Institute – Digital News Report 2026.** 100,000 interviews across 48 markets, June 2026. 37% global trust in news — the lowest recorded since 2015; 20% trust in news on AI chatbots. reutersinstitute.politics.ox.ac.uk/digital-news-report/2026

About the Author



Subrat Pati

VP- Group Manager - Client Services



Arindam Das

Senior Manager - Client Services



Amol Jain

Principal Technology Architect



Pranjal Dashora

Principal Consultant

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.