NANSEN

Operationalizing Opal AI

A Leadership Blueprint for Moving from Deployment to Durable Performance

Optimizely

Nansen Inc. | 300 W. Hubbard St. Suite 301, Chicago, IL 60654

Artificial intelligence is expanding what marketing and digital teams can accomplish. Experiments launch faster, insights surface earlier, and content scales with less manual effort. Eighty-nine percent of enterprise organizations are now actively using or testing AI applications across their operations.[1] The question is no longer whether organizations will adopt AI. It is whether their approach can sustain value at scale.

This guide addresses that challenge. It is designed for leaders who have moved beyond initial pilots and are now asking a more difficult question: **How do we turn early momentum into durable capability?**

The answer is not additional technology. Greater capability alone does not create competitive advantage.

As AI reduces the cost of generating ideas, analysis, and experimentation, organizations face a different operational reality. Opportunities multiply. Decisions increase. Teams can test more possibilities than ever before. Without a clear operating model, expanded capacity often creates fragmentation instead of progress.

Operationalizing Optimizely's Opal requires more than activating a new technology layer. Leaders must establish the systems that allow experimentation, governance, and decision-making to work together. Structure becomes the mechanism that converts speed into measurable outcomes.

Organizations that succeed do not simply deploy AI tools. They build the operating discipline that turns those tools into sustained advantage.

Table of contents

The Adoption Plateau

Across enterprise AI implementations, a consistent pattern is emerging.

During the first ninety days after deployment, momentum builds quickly. Teams adopt new tools, early gains reinforce confidence, and organizations begin to see how AI can reshape the way work gets done.

Then a more complicated picture begins to surface. Gartner predicted in June 2025 that more than 40% of agentic AI projects will be canceled by the end of 2027, not because the technology failed, but because organizations could not bridge the gap between a working proof of concept and a sustainable production system.[2]

The UK government's Copilot trial reflected a similar challenge. Across 20,000 civil servants over four months, participants reported an average of 26 minutes saved per day. Yet 17% experienced no meaningful benefit.[3] The most common uses were drafting documents, summarizing meetings, and writing emails. Useful productivity gains, but not operational transformation.

This is where many organizations stall.

AI expands the number of decisions an organization must process. When ownership remains informal, accountability becomes harder to trace, decisions slow, and governance begins to weaken.

There is another challenge that often emerges during this phase: assuming time savings automatically translate into return on investment.

Saving ten hours a week does not create value on its own. If that capacity simply gets redirected into more of the same work, little changes.

Value is created when organizations intentionally reinvest that time into higher-quality outputs, new initiatives, faster experimentation, or strategic work that was previously out of reach.

Organizations that move beyond the plateau strengthen the systems through which decisions flow. Decision authority becomes clearer, governance becomes more consistent, and leaders gain greater visibility into how AI contributes to broader business objectives.

When structure catches up to velocity, AI shifts from a productivity tool to a mechanism for sustained organizational learning.

The Shift:
From Production AI to Workflow AI

Much of the AI conversation over the past two years has focused on production AI: generating more content, faster and at a lower cost, often through a growing collection of disconnected tools.

Production AI creates value. But every new output introduces additional work: review, approval, distribution, and measurement. Without changes to the surrounding operating model, complexity grows faster than impact.

Organizations that move beyond the adoption plateau make a different shift. They move from production AI to workflow AI.

Rather than accelerating isolated tasks, workflow AI compresses entire cycles, from brief to published, from insight to decision, while embedding governance directly into the process.

Production AI often adds tools. Workflow AI prioritizes cycle time, decision quality, governance clarity, and learning velocity.

This distinction forms the foundation for the operating architecture that follows.

Cutting Through the Noise

Part of the challenge is that "doing AI" means something different to every organization. Teams often use the same language while describing entirely different levels of maturity and ambition.

In practice, AI adoption tends to appear in four forms:

1. Personal Assistant

Individual productivity gains such as drafting emails, summarizing meetings, or conducting quick research. Useful, but largely isolated to individual workflows.

2. Content Machine

Generating content and assets at scale. This is where many organizations begin and where many eventually stall.

3. The Optimizer

Running AI-assisted experimentation against defined business outcomes. Success at this stage depends on stronger data quality, measurement discipline, and clearer feedback loops than most organizations initially have in place.

4. Orchestration

Specialized AI agents coordinating end-to-end workflows, with people stepping in at defined moments for oversight and decision-making. This is where operational advantage begins to emerge.

Most organizations today operate between the first two stages. This guide is written for leaders working toward the fourth.

The distinction matters because each stage requires a different operating model. Productivity tools can operate independently. Orchestrated systems cannot. As AI becomes embedded across workflows, ownership, governance, and decision structures become increasingly important.

Treating orchestration like an individual productivity tool is often what causes organizations to stall before adoption becomes sustainable.

Redefining Success

The Five Dimensions of Maturity

Early AI adoption is often measured through activity. Logins increase. Queries expand. Drafts multiply.

These signals confirm activation. They do not predict competitive advantage.

Institutional maturity appears elsewhere. It is visible in how consistently organizations turn AI-generated insights into coordinated, accountable action.

At Nansen, we evaluate AI institutionalization across five dimensions that determine whether adoption becomes sustained advantage:

1. Decision Clarity

AI recommendations move through clearly defined ownership. Teams understand who evaluates insights, who approves action, and who is accountable for outcomes.

Measure: Time from insight to action; percentage of recommendations with a named owner.

2. Adoption Breadth

Usage extends beyond early champions and becomes embedded across the cross-functional workflows that shape everyday planning and execution.

Measure: Percentage of teams actively using AI within core workflows; repeat usage across functions.

3. Learning Velocity

Insights from one initiative improve the next. Knowledge compounds rather than remaining isolated within individual experiments or teams.

Measure: Experiment cycle time trending down; documented reuse of insights across teams.

4. Governance Stability

Operational guardrails create clarity around brand standards, legal requirements, and data boundaries, allowing teams to move faster without increasing risk.

Measure: Percentage of AI outputs passing review without manual intervention; time from draft to approved output.

5. Orchestration

Experimentation influences broader business decisions, including roadmap priorities, resource allocation, and investment focus.

Measure: Strategic decisions informed by experiment data; roadmap initiatives influenced by testing outcomes.

Organizations focused only on activity increase output.

Organizations focused on maturity build advantage.

The difference between the two is durability.

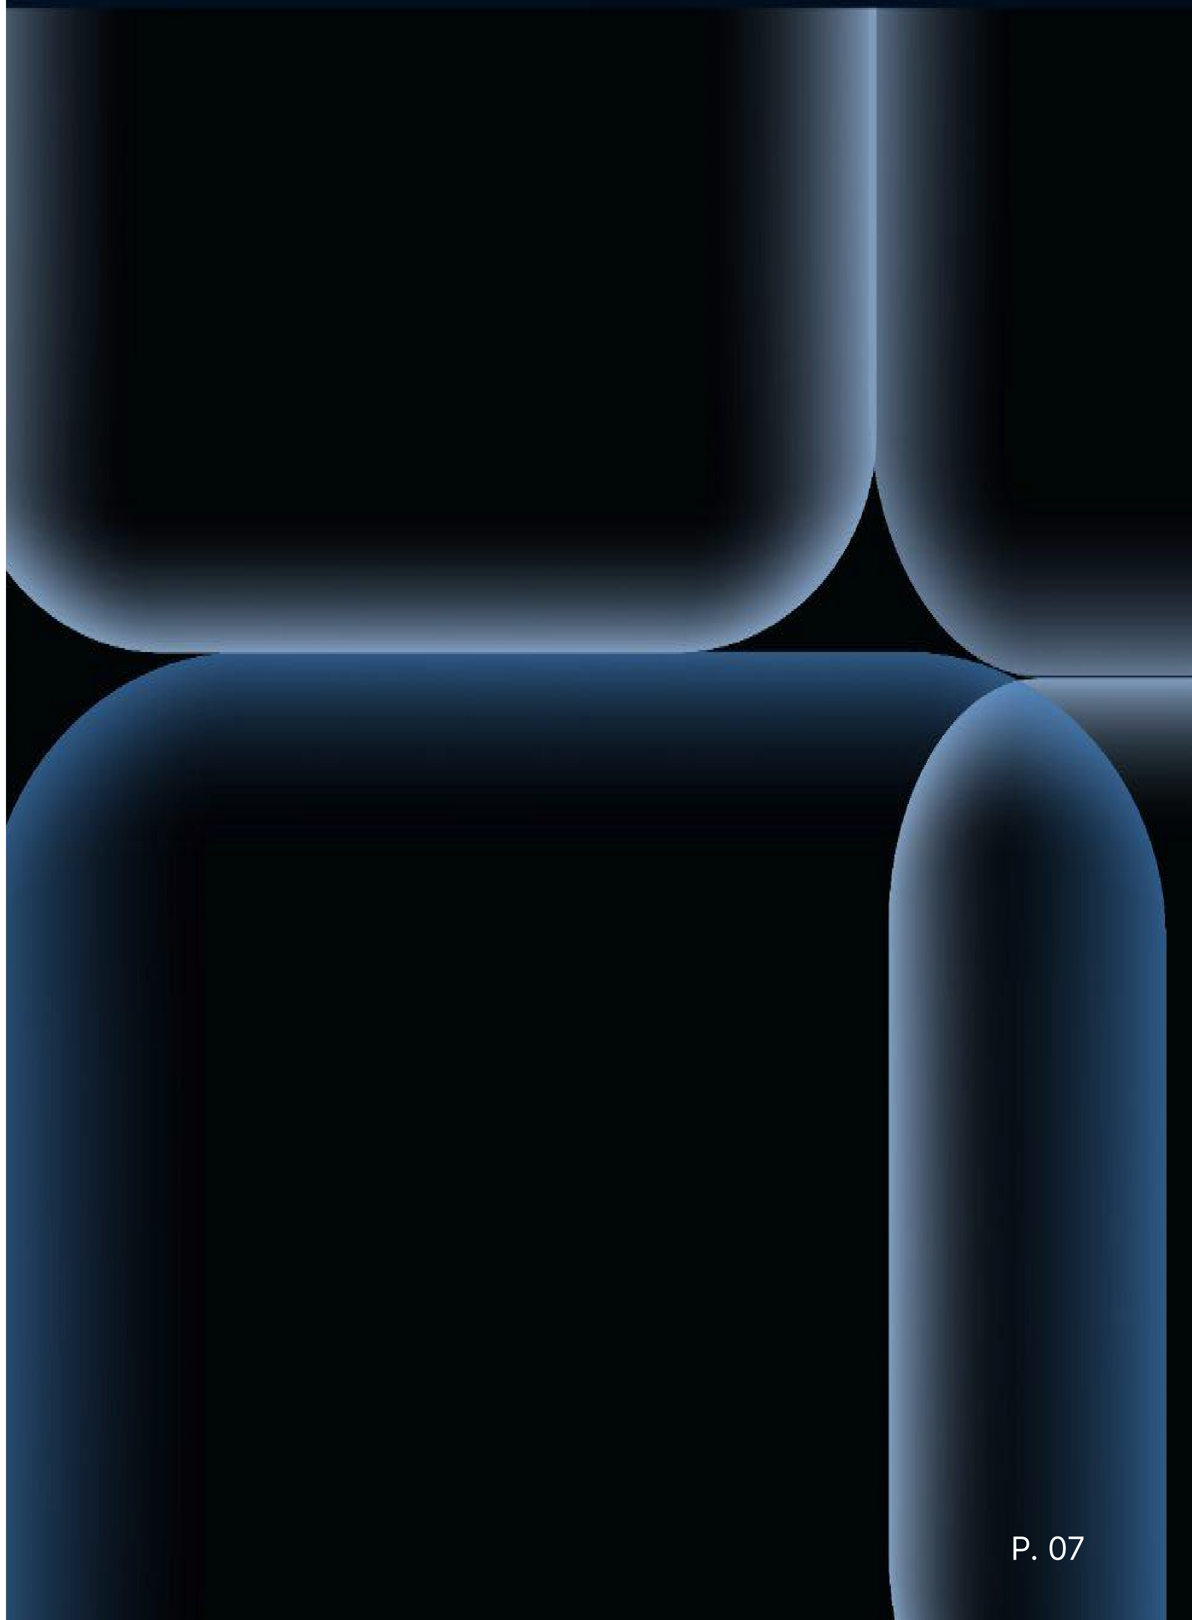

The Nansen Operating Architecture

Across enterprise Opal implementations, four structural domains consistently determine whether AI scales sustainably.

Together, these domains define how organizations convert AI deployment into institutional advantage. Each addresses a specific failure point that commonly emerges at the adoption plateau.

Strategic Anchoring

As hypothesis generation becomes inexpensive, prioritization becomes more important than ever.

Without clear direction, AI expands activity across too many fronts at once. Experimentation increases, but impact becomes diluted.

Organizations that scale successfully focus AI efforts on a small number of High-Impact Experience Zones: areas where improvement directly influences revenue, retention, operational efficiency, or competitive differentiation.

These zones often include high-value customer journeys such as checkout flows, lifecycle engagement programs, pricing and packaging strategies, or experiences tied to high-value customer segments.

Anchoring experimentation within these areas concentrates learning and increases the likelihood that AI acceleration translates into measurable business outcomes.

A practical starting point is identifying three operational patterns:

- Repetitive, rules-based tasks such as product descriptions generated from structured specifications or recurring reporting workflows
- High-volume content adaptation across channels, formats, and audiences
- Review and quality-control processes tied to brand, accessibility, legal, or compliance standards

These patterns reveal where AI can create immediate operational leverage without requiring wholesale process redesign.

Leadership's role is to determine where deployment should focus first and which initiatives represent the highest-leverage opportunities.

Decision Architecture

As AI adoption expands, many organizations discover that the primary constraint is not the technology itself, but the human systems responsible for evaluating and acting on new insights.

Road Scholar, a nonprofit that creates educational travel experiences, experienced this shift firsthand. Before establishing a formal decision architecture, AI agent requests arrived through informal channels: Slack messages, meeting sidebars, and scattered email threads. Prioritization varied by team, approval paths were inconsistent, and ownership remained unclear.

After implementing a weekly AI review process with named decision owners and explicit end-user sign-off before production, the organization reduced the path from concept to live agent to as little as two weeks.

The pattern extends well beyond a single organization. Sustainable velocity depends on clearly defining how Opal-generated recommendations move through decision authority:

- AI systems generate hypotheses and recommendations
- Leaders evaluate alignment with strategy, governance, and risk tolerance
- Named owners implement decisions and remain accountable for outcomes

Clear decision pathways reduce friction, shorten execution cycles, and prevent experimentation from turning into prolonged debate.

When decision architecture functions effectively, work moves through ownership rather than negotiation.

Guardrails as Infrastructure

Confidence enables speed.

Rather than requiring manual approval for every AI-generated output, leading organizations embed governance directly into the workflows that guide experimentation.

Brand standards, legal requirements, and privacy protections become part of prompt frameworks, operational playbooks, and system configurations.

Teams can then operate confidently within defined boundaries, knowing experimentation remains aligned with organizational expectations.

This creates a shift from approval-based governance to infrastructure-based governance.

Oversight decreases because clarity increases.

When guardrails function as infrastructure, organizations accelerate experimentation without introducing operational risk.

Learning Velocity

Experimentation creates value only when insights compound over time.

Organizations that sustain advantage establish a disciplined cadence for reviewing, synthesizing, and applying what experimentation reveals.

Weekly reviews identify emerging performance patterns. Monthly discussions connect insights across initiatives. Quarterly sessions recalibrate priorities based on accumulated learning.

A useful mental model is the flywheel:
Create → Review → Analyze → Experiment

Each cycle produces sharper insights, stronger standards, and more effective execution. Over time, the loop tightens.

That tightening is learning velocity made visible, and the operating architecture described throughout this guide is what keeps it moving.

Opal Institutionalization in Practice

Road Scholar approached Opal as an evolution of its operating model rather than a standalone productivity tool.

Before deployment, leadership identified a persistent operational constraint. The programs team spent nearly nine hours per person manually reviewing sustainability documentation across more than 5,000 travel programs: a highly structured, rule-based process constrained by scale and repetition.

Rather than layering AI onto existing workflows, the organization first established clear ownership, governance, and quality standards, then embedded Opal directly into that operating structure.

The impact was immediate.

A sustainability review agent reduced a nine-hour manual process to seven minutes across more than 5,000 programs, a 97% reduction in processing time. A second agent, designed to process activity notes across more than 940 travel programs, reduced runtime by 96%.

Both agents were approved for production by the end users responsible for relying on the outputs every day, not solely by leadership or development teams.

What distinguished the approach was the operational discipline surrounding deployment. Each agent moved through a defined build, testing, and approval process. Autonomy boundaries were established before development began: the agent handled structured extraction, subject matter experts validated outputs, and strategic decisions remained within the team.

Within six months, four departments were actively using Opal.

That level of structural clarity marked the shift from initial activation to true institutionalization.

Cultural Architecture

AI adoption ultimately becomes a leadership challenge.

Executives must balance competing priorities: speed and control, autonomy and accountability, efficiency and professional judgment.

Technology alone cannot resolve these tensions.

Leadership determines how experimentation operates within the organization by establishing clear expectations for ownership, transparency, and responsible use.

When expectations are explicit and decision authority is visible, teams gain the confidence to experiment without weakening trust or introducing unnecessary risk.

In that environment, AI stops functioning as a standalone initiative and becomes embedded within how the organization operates.

The 90-Day Institutionalization Framework

The Nansen Operating Architecture provides the structural model for scaling AI. The following framework outlines how organizations operationalize that model during the first ninety days after deployment.

Institutional maturity does not emerge automatically. It must be designed.

During this transition period, organizations either establish the discipline required for sustainable AI adoption or allow experimentation to fragment across teams.

Institutionalization occurs when structure, workflow, and accountability evolve alongside experimentation velocity.

Days 0—30: Structural Clarity

The first phase establishes the conditions required for experimentation to scale safely.

During early deployment, momentum accelerates quickly. Without clear authority and governance standards, however, teams hesitate, approvals slow progress, and accountability becomes diffuse.

Leadership removes ambiguity before velocity increases further.

Decision rights become explicit, governance guardrails are codified, and High-Impact Experience Zones are identified to ensure AI deployment aligns with core business objectives.

Teams use the three-pattern lens — repetitive tasks, high-volume adaptation, and review and quality control — to identify which workflows are best-suited for AI support.

At this stage, teams begin operating within defined boundaries rather than informal assumptions.

Clarity creates confidence.

Days 30—60: Workflow Integration

Once structural clarity is established, organizations move AI from isolated experimentation into the workflows that guide daily execution.

Opal becomes part of campaign planning, experiment design, and experience optimization.

AI-generated insights begin informing planning discussions. Documentation of tests and outcomes becomes standardized. Results are reviewed through a recurring operational cadence.

Autonomy boundaries — AI-driven, AI with review, and human-led — are defined before new agents or workflows are introduced.

At this stage, experimentation becomes part of how teams plan, execute, and evaluate work rather than a standalone innovation effort.

AI begins shaping operational strategy rather than simply accelerating execution.

Days 60—90: Institutional Commitment

The final phase converts experimentation into organizational learning.

By this stage, experimentation is visible across multiple teams. Leadership ensures insights accumulate across the organization rather than remaining isolated within individual initiatives.

Adoption breadth expands. Experimentation insights begin informing prioritization decisions. Decision bottlenecks become easier to identify and resolve.

AI-supported workflows start influencing roadmap discussions, resource allocation, and customer experience improvements.

The flywheel — Create → Review → Analyze → Experiment — is now in motion.

At this stage, Opal functions less like a tool and more like operational infrastructure.

Measuring What Predicts Durability

Short-term performance gains demonstrate potential. Structural indicators reveal whether experimentation is strengthening the operating system itself.

Leaders should monitor signals such as:

- Reduction in experiment cycle time, from hypothesis to measurable result
- Reuse of insights across initiatives and teams
- Frequency of decision bottlenecks caused by unclear ownership
- Reduction in governance friction as guardrails become embedded within workflows
- Expansion of cross-functional adoption beyond early champions

These indicators reveal whether experimentation is producing isolated improvements or compounding organizational learning over time.

Durability is what separates momentum from sustained advantage.

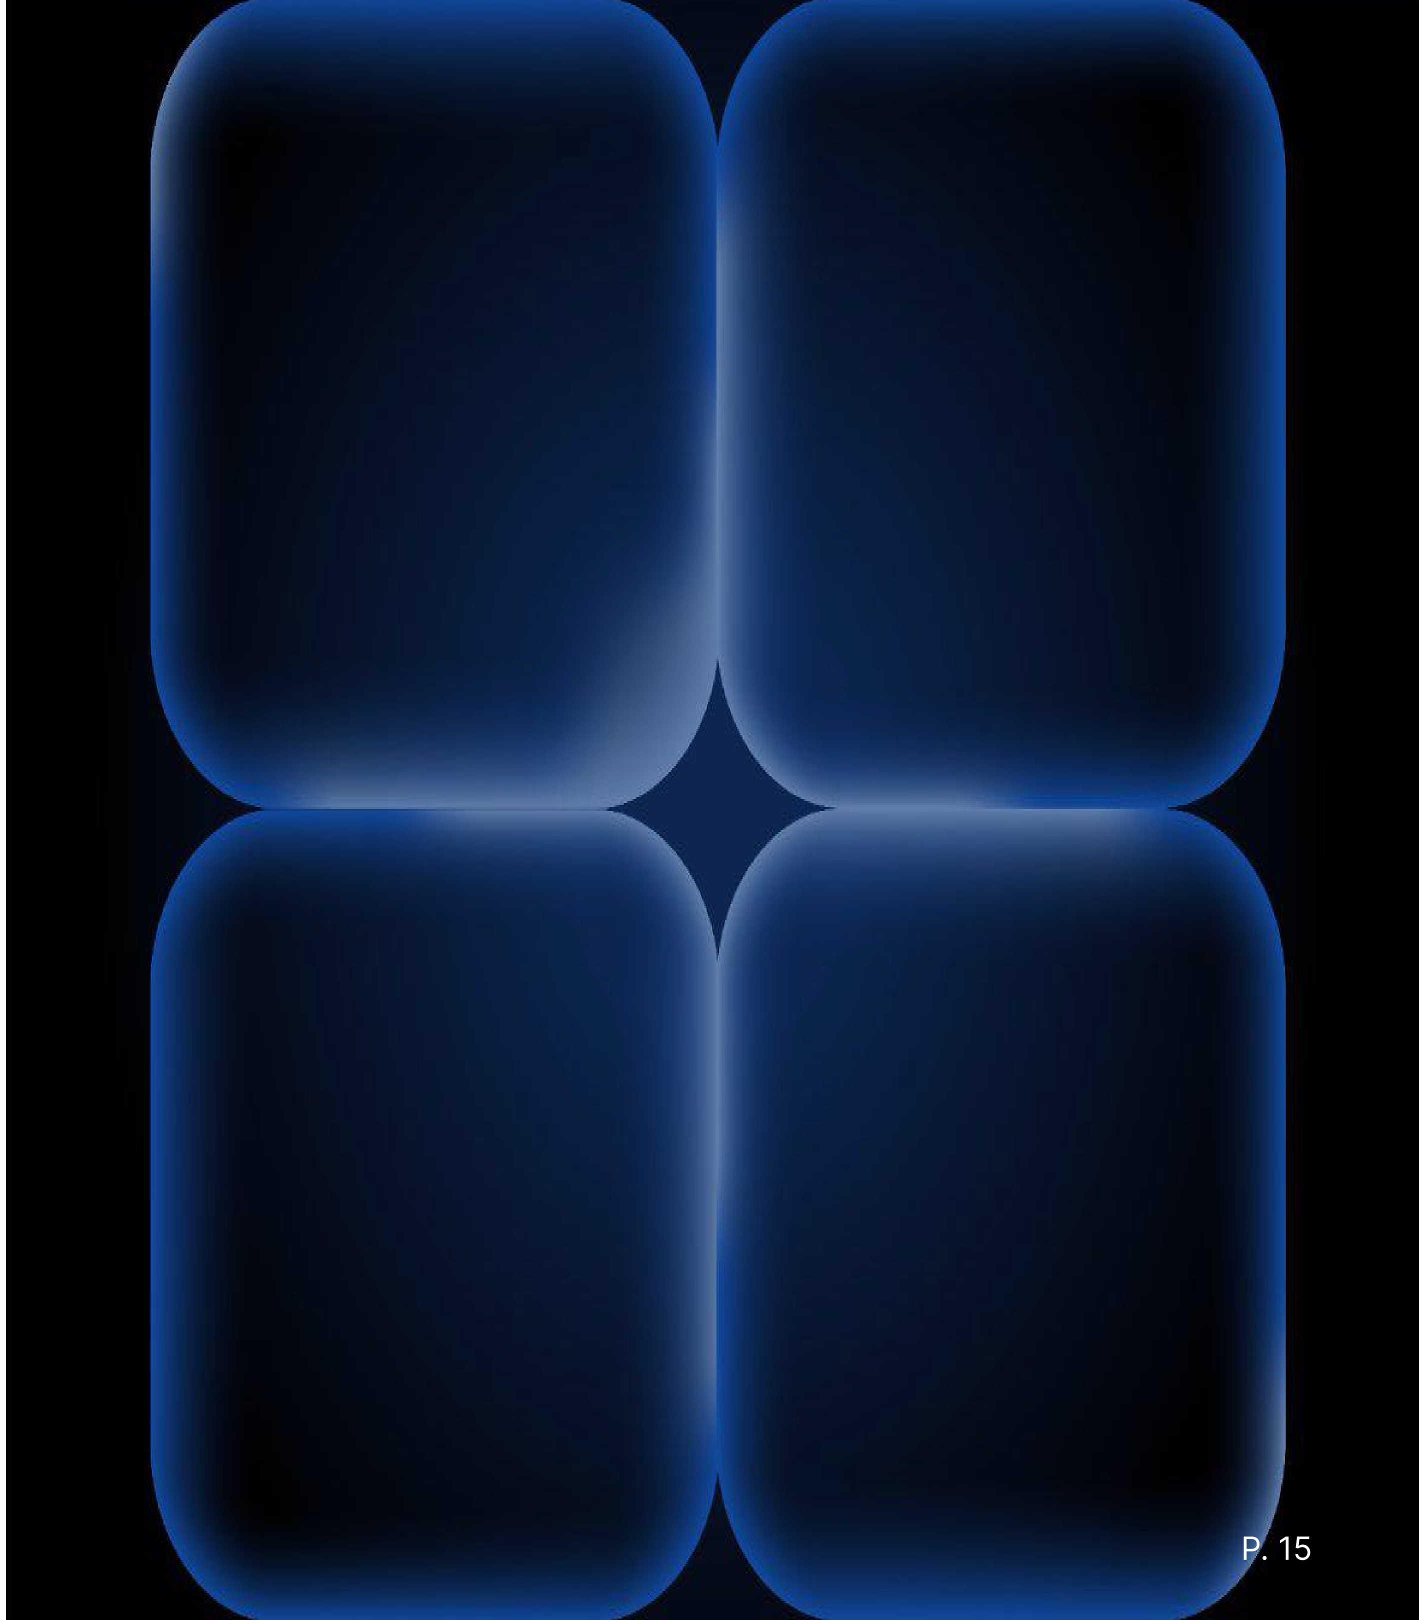

The AI Maturity Continuum

Enterprise AI adoption typically progresses through four stages.

Activation

Teams begin exploring AI, but usage remains fragmented and largely isolated to individual productivity or content generation use cases.

Integration

AI becomes embedded within planning and operational workflows across more teams. Agents begin moving from proof of concept into production environments.

Governance Discipline

Decision authority, review cadence, and governance guardrails evolve to support increased experimentation velocity. Autonomy boundaries are applied consistently across workflows.

Strategic Engine

AI begins informing roadmap decisions, investment priorities, and competitive positioning. Organizational learning compounds continuously, and AI begins influencing the business at a systemic level.

Many organizations plateau between Integration and Governance Discipline. Progress beyond this stage rarely depends on additional tooling. More often, it requires stronger leadership structure, clearer accountability, and operational discipline.

At full maturity, AI no longer feels experimental. It becomes foundational to how the organization operates.

Leaders assessing institutional maturity should ask:

- Are AI-generated recommendations routed through clearly defined authority?
- Does the organization review experimentation through a consistent operational cadence?
- Has adoption expanded beyond early champions into core workflows?
- Are governance guardrails codified and trusted by teams?
- Do experimentation insights influence strategic prioritization?

Unclear answers usually indicate incomplete institutionalization.

At that point, structural clarity becomes the next priority.

Designing for Durable Advantage

Operationalizing Opal ultimately reflects a leadership decision about how performance is designed, governed, and sustained.

Organizations that treat AI primarily as a productivity layer may achieve incremental gains in efficiency and output. Organizations that treat AI as operational infrastructure reshape how work is prioritized, how decisions are made, and how experimentation informs strategy.

Competitive advantage rarely belongs to the organizations that adopt new technologies first. It belongs to those that operationalize them with discipline and intentionality.

When leadership establishes clear decision authority, embeds governance directly into workflows, and aligns experimentation with strategic priorities, AI-supported learning begins to compound rather than fragment.

In that environment, experimentation stops functioning as a collection of isolated initiatives. It becomes embedded within the organization's operating system.

When experimentation becomes infrastructure, infrastructure becomes competitive leverage.

Operationalizing Opal with Nansen

Many organizations have already invested in AI. Far fewer have established the operating structure required to turn that investment into sustained performance.

Technology enables AI. Nansen operationalizes it.

Nansen designs the missing layer between AI deployment and measurable business impact: the decision architecture, governance frameworks, workflow design, and experimentation systems that allow Opal to scale with clarity, accountability, and operational discipline.

Without that structure, experimentation often fragments across teams, ownership becomes unclear, and early momentum stalls before it produces durable advantage.

Organizations that operationalize AI effectively will shape how their category evolves. Those that treat AI as a disconnected productivity initiative will struggle to convert activity into sustained competitive leverage.

If your organization is deploying Opal, the operating model you establish now will determine whether AI becomes a temporary productivity layer or a long-term source of competitive leverage.

Connect with Nansen to begin operationalizing Opal across your organization.
www.nansen.com/contact.

Sources

1. *ISG, Enterprise AI Adoption Trends Report, 2025.*
2. *Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," June 2025.*
3. *UK Government, "Microsoft 365 Copilot Experiment: Cross-Government Findings," 2025.*

Methodology Note

Road Scholar operational metrics and implementation observations are based on Nansen client engagement data and internal deployment analysis.