



VERNDALE


F I S H T A N K

A Verndale Company


AI IN YOUR DXP

A Sitecore Innovation Report

7 1ST EDITION: PUBLISHED 2026



SitecoreAI™





This is what enterprise AI was always supposed to do

A guide from Sitecore's #1 Platinum Partner.

The gap between what organizations expected from AI and what they actually got isn't a technology problem. It's a structural one. AI got added on top of fragmented stacks, confined inside individual platforms, and handed to teams without changing the model underneath. The result is what you already know: more investment, flatter returns, and a growing distance between ambition and execution.

We're Fishtank, a Verndale company, an award-winning Sitecore Platinum Partner with over 27 years building enterprise digital experiences for some of the most complex organizations in North America. We wrote this guide to close that gap, with the research that explains it, the framework that fixes it, and the platform and intelligence layer that makes it real.

01 Act one: The problem

Page 2

02 Act two: The agentic frontier

Page 8

03 Act three: The solution

Page 13

04 Act four: The benefits

Page 19



01

Act one: The problem

➤ The hidden tax of context switching

The proliferation of disconnected tools creates a hidden operational burden across the organization.

➤ More AI investment with no impact

Teams keep spending on AI but aren't seeing real ROI or transformational change at the enterprise level.

➤ AI isolated in every tool can't scale

AI tools are siloed within each platform, preventing any orchestrated, organization-wide impact.

➤ It's the model, not the people

The broken cycle of training everyone, slow consolidation, someone leaving, starting over — this is a structural failure, not a lack of talent.

The majority of enterprise teams are running the same four-problem pattern right now. The tools vary, the budgets vary, the org charts vary, but the underlying conditions are identical across every organization we work with.

McKinsey surveyed nearly 2,000 people across organizations in mid-2025 and found that while almost nine in ten are regularly using AI, most have yet to embed it deeply enough into their workflows to realize material enterprise-level benefits.¹ Deloitte's global survey of 3,235 business and IT leaders found the same: access to AI is rising faster than the ability to put it to work.² The gap between availability and adoption is now the primary barrier to value.

These four problems explain why. And after decades of solving them at scale, we know they require more than a new tool or another training cycle to fix.

¹ QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (McKinsey & Company, November 2025)

² Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge* (Deloitte, January 2026)



Too many systems, one hidden tax

Context switching is the hidden cost holding back AI adoption within every enterprise team right now.

Your team’s work doesn’t live in one place. Conversations happen in Teams or Slack, work gets assigned in Jira, Monday, or Asana, knowledge lives in SharePoint or Confluence, and data sits in Salesforce and Analytics. The CMS is where content is supposed to land, but it’s rarely where the work to create it actually starts. Every move between these systems carries a cost that never shows up on a budget line, and the friction of rebuilding context quietly kills momentum before it builds into anything meaningful.

Breaking the pattern

Connect where work already happens.

Deloitte identifies activation, alongside of access, as the primary barrier to value.² That means meeting teams in the tools they already use daily, not adding another system to the stack.

Embed AI where the work already happens.

McKinsey’s high performers embed AI solutions directly into business processes¹, changing how frontline employees actually work rather than asking them to adopt a parallel workflow.

Design for deployment from the outset.

Organizations that treat scale as an afterthought see far lower adoption. Role-specific integration and visible advocacy are what shift behavior, not broader AI fluency training alone.

Figure 1 — Where enterprise work actually lives

This fragmentation was already expensive before AI entered the picture. The problem is that most organizations responded to it by adding AI tools on top of the same disconnected stack, one per platform, each smart in isolation, none of them connected to the others. More access has not translated into more utilization because the tools landed inside the same fragmented structure that made adoption difficult in the first place.




Figure 2 — Sanctioned AI access jumped 50% in one year from 2024 to 2025

When the gap is baked into the structure, there’s nothing the workforce can do about it. And adding more tools to a disconnected stack has never been the answer. Deloitte found that while 50% more workers now have access to sanctioned AI tools compared to just a year ago, fewer than 60% of those with access actually use AI in their daily workflow, and that number has barely moved.



Note. Figure 2 from Deloitte AI Institute, State of AI in the Enterprise. Copyright 2026 by Deloitte.

¹ QuantumBlack AI by McKinsey, The State of AI in 2025: Agents, Innovation, and Transformation (McKinsey & Company, November 2025)
² Deloitte AI Institute, State of AI in the Enterprise: The Untapped Edge (Deloitte, January 2026)



The gap between AI investment and enterprise impact is growing

AI impact and results are rapidly being outpaced by ever-increasing AI investment.

The budgets are real and the enthusiasm is real, but so is the growing distance between what organizations are spending on AI and the enterprise-level results they are building toward. McKinsey's 2025 global survey is unambiguous: while teams across every function report use-case level cost and revenue benefits, only 39% of organizations report enterprise-level EBIT impact from their AI investments.¹ The majority are spending, seeing pockets of improvement, and treating that as sufficient. Enterprise-wide transformation of the kind that actually changes how the business runs is reaching fewer than four in ten. The problem isn't in the budget, it's the repetitive cycle that continues to spiral, and the model producing it stays the same.

Figure 3 — The vicious AI pilot cycle

That pattern has a shape. A pilot runs successfully in controlled conditions with a small team and clean data. Production requires infrastructure investment, system integration, security reviews, and compliance checks, each demanding far more resources than the pilot ever did. The organization has no reliable way to know which pilots will scale, so it funds another one instead.

7 01 Invest in AI

Tools deployed, budgets spent, teams trained.

7 02 Pilot succeeds

Positive results in controlled conditions.

7 03 Scaling stalls

Can't predict ROI at scale, integration complexity hits.

7 04 Fund a new pilot

No enterprise impact, create new pilot, repeat.

¹ QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (McKinsey & Company, November 2025)

F I S H T A N K

AI still hasn't changed how the business operates.





Every tool gets smarter in isolation, but the business stays fragmented


Current AI features are contained within individual tools and platforms, with no connective layer to carry intelligence across the business as a whole.

Every major platform your team uses has added an AI feature in the last eighteen months. Salesforce, Jira, your CMS, Teams with Copilot, all have one. Each tool gets more capable within its own walls while the organization as a whole stays exactly as fragmented as it was before, because there is nothing connecting the intelligence across systems. Every AI feature delivers incremental lift inside the platform it came with, and generates friction at every boundary between them.

Deloitte found that only 43% of leaders feel their technical infrastructure is highly prepared for AI, and that number is falling year over year despite confidence in AI strategy being on the rise.² The foundation required to carry intelligence across systems is being left behind, and as Deloitte's research confirms, poor or fragmented data compounds risk and undermines every AI initiative it touches. The platforms are only part of the equation. The connective layer that makes intelligence flow across all of them is what's actually missing.

Figure 4 — Value dies at the boundary between tools

Even when intelligence is added to each system tool, there's still no system-wide integration. This friction that still exists at every tool boundary is what kills enterprise AI adoption.



★ Added AI layer — a lift inside each system, but still friction at every boundary between them.


What actually happens?

Excitement wanes, pilots stall, old workflows creep back, and AI becomes another unused tool.

² Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge* (Deloitte, January 2026)

Figure 5 — Agent scaling by function

The data reflects what teams are confirming on the ground. AI pilots build momentum in controlled conditions, then stall when they hit the edges of the platform they were built inside. McKinsey's data puts the scale of that boundary problem in concrete terms. No function connects to another, the enterprise as a whole never scales.



Note. Figure 5 from QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (n = 1,993). Copyright 2025 by McKinsey & Company.

Without the proper infrastructure to support it — AI fragmentation is deepening even as confidence in AI strategies continues to rise.



It was never a talent problem

The cycle keeps repeating because the model producing it stays intact.

This is the problem most organizations misdiagnose, and it tends to be the most expensive mistake they make. When AI adoption stalls, the instinctive response is to train more people. Run workshops, build learning paths, get everyone up to speed on the new tools, until someone leaves and the whole process starts over again.

84% of companies surveyed by Deloitte have added AI to an existing model that was never built to support it, then trained people to use it within that same structure². When someone leaves, the knowledge goes with them. When a consolidation point gets overloaded, delivery stalls. When more investment arrives, it gets directed into more training. The problem was always the structure, and the people inside it were doing exactly what the structure asked of them.

The workers are ready. Deloitte found that 68% of non-technical workers are open to or enthusiastic about using AI. McKinsey found that the single biggest differentiator between AI high performers and everyone else is operating model. While most are focused on educating employees, far fewer are rearchitecting company roles, daily workflows, and career pathways. High performers are nearly three times as likely to have fundamentally redesigned their workflows in deploying AI, and 65% of them have senior leaders who demonstrate genuine ownership of AI initiatives, compared to just 25% of all other respondents.

¹ QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (McKinsey & Company, November 2025)

² Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge* (Deloitte, January 2026)

Figure 6 — The broken cycle every organization keeps repeating

The cycle feels so familiar it can seem inevitable. It is a structural failure that has been misread as a talent gap.





Figure 7 — Enterprise confidence in strategy versus infrastructure readiness

Strategy confidence is rising. While infrastructure and data readiness are falling. The foundation needed to connect AI across systems is getting left behind. Closing this gap requires more than a new platform or a better pilot. It requires a connective layer that carries intelligence across every system where work already happens.



Note. Figure 7 from Deloitte AI Institute, State of AI in the Enterprise: The Untapped Edge. Copyright 2026 by Deloitte

The gap between AI readiness and AI reality

84%
haven't redesigned
job workflows around
AI integration.²

68%
of non-technical workers
are open or enthusiastic
about using AI.²

20%
of company leaders
feel their talent
strategy is AI-ready.²

65%
of industry high performers
have senior-level ownership of
AI — compared to just 25% of
everyone else.¹

What the highest performing organizations are doing differently

The organizations generating real enterprise-level returns from AI share a structural profile that separates them from the majority.

The research is consistent across McKinsey and Deloitte: the organizations generating real enterprise-level returns from AI share a structural profile that separates them from the majority. They are making three moves that most organizations have yet to prioritize.

- **A governance layer that coordinates AI across the whole business.**
Having a centralized team that connects and links AI efforts across the organization is among the strongest predictors of high performance. Intelligence working in one function in isolation produces incremental gains. Intelligence coordinated across functions produces compound ones.
- **Workflows redesigned around AI at the operating model level.**
High-performing organizations treat AI as a structural shift, rebuilding processes around its capabilities from the ground up. The distinction between embedding AI into how work gets done and adding AI features on top of how work already gets done is where the performance gap opens up.
- **A unified data and infrastructure strategy.**
Investing in platforms that connect data flows, enforcing enterprise standards for interoperability, and establishing domain-owned data products are the infrastructure conditions that allow coordinated AI to scale. Deloitte found that only 43% of leaders believe their AI infrastructure is highly prepared, and only 40% believe the same of their data management.² Strategy confidence is rising. The foundation required to act on it is still catching up.

¹ QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (McKinsey & Company, November 2025)


² Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge* (Deloitte, January 2026)




Figure 8 — How most companies are responding versus what actually works

Deloitte's data exposes the misdiagnosis directly. Most companies respond to AI adoption challenges by educating the broader workforce to raise AI fluency. Few are reimagining organizational structures around AI, and even fewer are redesigning career paths and workflows.

The most common responses address the people (yellow), while the structural fixes that actually move the needle (black) remain the least common.



Note. Figure 8 from Deloitte AI Institute, State of AI in the Enterprise: The Untapped Edge. Copyright 2026 by Deloitte.





02

Act two: The agentic frontier

➤ The rebrand everyone saw

XM Cloud became SitecoreAI: six products, one platform, real benefits.

➤ The shift almost nobody covered

The same launch shipped the Marketer MCP and the Agent API.

➤ A new frontier, not just a feature

Agentic value lives inside SitecoreAI and in the open space beyond it.

➤ The question the frontier asks back

Everyone gets access to the same doors, so how do you actually win?

Act One defined a structural problem: AI bolted onto fragmented stacks, with value dying at the boundary between tools. Every platform got smarter on its own, and the organization stayed just as disconnected as before. This act is about the first platform to take that problem seriously, and the move most of the market missed entirely.

In November 2025, Sitecore retired the XM Cloud brand, consolidated six separate products into one unified platform, and quietly opened that platform to the agentic frontier. That space, inside the platform and in the open landscape beyond it, is where the next decade of competitive advantage actually gets built.

Welcome to the era of SitecoreAI.



SitecoreAI™



SitecoreAI: The rebrand everyone saw


Six products became one platform, and the upgrade is real.

On November 4, 2025, Sitecore retired the XM Cloud brand and relaunched as SitecoreAI, consolidating six separately licensed products, content management, digital asset management, search, personalization, customer data, and Stream, into a single platform. They now live as four modules on one foundation: CMS as the anchor, Conversion Optimization handling search and personalization, DAM, and Agentic Studio. Existing XM Cloud customers moved over with no migration required, which is rarer in enterprise software than it should be.

Agentic Studio: The platform's agentic engine

Agentic Studio is a governed workspace where AI does real work, built on Agents, Flows, Spaces, and Signals, and launched with roughly twenty prebuilt agents plus the option to build your own with no code. Out of the box, it writes blogs, emails, and bulk content at production scale, translates and localizes while holding voice and format, researches competitors and audiences for AEO and SEO, spins up variants wired directly to A/B/n testing, audits work against brand, persona, and accessibility standards, and maps source documents straight into your templates so a brief becomes a publishable page.

Figure 9 — The immediate impact of SitecoreAI



Key benefits to highlight

- Easier authoring**
 Drag-and-drop WYSIWYG Page Builder; Fishtank clients report 4x–10x faster than the old Experience Editor. (Page-creation time down ~60%)³
- Faster development**
 Fully managed SaaS, built-in DevOps, one-click deploys, React/Next.js. (Integration productivity up ~50%)³
- Easier to add technologies**
 API-first and headless; connect modern frontends and AI apps, and hire from a far broader talent pool.
- Higher approval rates**
 WYSIWYG previews and fewer back-and-forth loops mean faster sign-off.
- Happier teams**
 IT stops babysitting infrastructure; marketing stops waiting in the dev queue.
- Results that earn buy-in**
 371% ROI (payback under six months) and a 50% lift in conversions (~\$2.1M).³

³ Forrester, TEI of the XM Cloud foundation now powering SitecoreAI (Forrester, July 2025)

Note. Figure 9 from Forrester, TEI of the XM Cloud foundation now powering SitecoreAI. Copyright 2025 by Forrester.

The massive shift almost nobody covered

In the same launch, with far less noise, Sitecore shipped two things that matter more to your operating model than any rebrand. They're two halves of one experience, and here's each half.



The Marketer MCP (the doorway)

The Marketer MCP is Sitecore's implementation of the Model Context Protocol, the open standard that lets an AI assistant use a system's tools safely. In practice, you work inside the AI tools you already use, Claude, Copilot Studio, Cursor, and simply ask: create a French landing page with an A/B test. No menus, no API calls. Because MCP is an open standard, the same doorway works both directions: your assistant can pull a task from Asana and act on it in SitecoreAI in one motion. It's probabilistic, meaning an assistant interprets your intent and decides what to do.

The Agent API (the engine)

The Agent API is the engine underneath, built for developers and partners. Skip the function names and think in terms of what you can ask it to do: create or update a page, drop a component onto it, spin up a personalization variant, launch an A/B test, find and place an asset. Each is a defined, governed action. Trigger it directly, or let the Marketer MCP trigger it for you, and every action runs on enterprise rails, including permissions, approval workflows, and rollback.

It's deterministic, meaning a defined action produces a defined result.

Two halves of one experience

This isn't two competing options or two different outcomes. The probabilistic doorway and the deterministic engine are two halves of the same thing: natural language intent up front, reliable execution underneath, with the MCP calling the API. Together they make SitecoreAI fully reachable, and that completeness is exactly what opens the agentic frontier on the next page.



The new frontier: The boundary is where advantage compounds

Two sets of capabilities. Two doors between them. That's where the advantage lives.

Everything on the last page, the consolidated platform, Agentic Studio, the built-in agents, is powerful, governed, and genuinely useful. It's also one platform's capability. The work your teams do every day happens across far more surfaces than that.

Capabilities in SitecoreAI

What Sitecore consolidated is real: Agentic Studio and the roughly twenty agents you can build on, plus research, testing, writing, personalization, translation, and DAM. Powerful and governed, and by definition table stakes, the baseline every competitor on the same platform can claim.

Capabilities in your enterprise


Your teams work across many systems: Teams, Slack, Salesforce, Jira, Google Analytics, FishtankOS, and Asana, the surfaces where decisions get made and work actually happens. That's where differentiated advantage is built.

The two doors connecting everything

The Marketer MCP and Agent API sit between the two. One lets you reach SitecoreAI's capabilities from the tools you already use; the other lets SitecoreAI's governed actions be triggered by the systems around it. Natural-language intent in, reliable execution underneath, carrying intelligence both ways. The connection between them is where advantage compounds.

Figure 10 — How the Marketer MCP and Agent API connect SitecoreAI to your enterprise

On one side, the systems your teams already work in. On the other, SitecoreAI's governed capabilities. Between them, the Marketer MCP and Agent API, the connective layer that carries intelligence in both directions, so work happens inside the tools your teams already use.





The real question is what will you do next?

Here's the catch on the far side of a great launch: the Marketer MCP and Agent API are open, so your competitors get the same doors you do. The moment everyone can connect SitecoreAI to everything, connecting stops being a differentiator. You won't always want to work inside SitecoreAI (that's the whole point of the frontier) but a surface is not a strategy.

So the frontier asks its question back at you: not whether you can connect, but how SitecoreAI's frontier plugs into the way you actually work, who governs agents now acting across your stack, and what turns standard-issue doors into an operating model that compounds. Those have an orchestration answer, not a product one.

7 Reach without judgment

Open standards route actions efficiently, but they don't decide which ones actually serve the business. The doors are open to everyone; what you do with them is still on you.

7 Access without governance

Agents operating across your stack can move fast, but the platform has no visibility into the guardrails your organization needs. Access without governance is just exposure.

7 Connection without context

SitecoreAI knows its own systems and standards deeply, but it doesn't know yours. Connecting is the easy part; making that connection intelligent is the harder problem.



03

Act three: The solution

➤ Introducing: FishtankOS

Fishtank's customized intelligence platform that connects the systems your team already runs and delivers output where work actually happens.

➤ Work where you already work

Your team keeps working exactly as they do today. FishtankOS brings intelligence into the tools they already use, zero behavior change required.

➤ A new operating model

FishtankOS sits at the center of the operating model, reading intent and routing execution across every connected system simultaneously.

➤ The intelligence layers

Every connected system gets two layers of intelligence: foundational knowledge, and organizational knowledge.



Introducing FishtankOS



FishtankOS is the operating layer for the agentic frontier: it connects the systems an organization already runs (SitecoreAI among them) and adds an intelligence layer that reads intent, routes execution across those systems, and delivers output where work is already happening. It is the connective layer that makes those tools work together as a unified, intelligent operating model — extending the value of every platform already in use, rather than leaving them a fragmented stack of isolated capabilities.

The distinction matters because most AI deployments today stop at the platform level, delivering value inside a single tool while leaving the coordination work between tools completely unaddressed. A CMS gets an AI writing assistant. A CRM gets predictive lead scoring. A project management tool gets smart task suggestions. Each of these may improve individual productivity in isolation, but they do nothing to address the coordination work that happens between platforms, and that coordination work is where the majority of enterprise team time is actually spent.

FishtankOS is designed to address the coordination layer directly.


Why we made FishtankOS

Execution, assembly, and repetitive data entry are the tax on getting to the work that actually matters.

At its core, FishtankOS is an organization's operating system: a purpose-built intelligence model shaped around the specific platforms, processes, and outcomes that define how that organization operates. A team that runs SitecoreAI, Salesforce, Jira, Microsoft Teams, and Google Analytics gets an intelligence layer that understands how all five of those systems work, how they relate to each other, and how to coordinate across them to complete real tasks. The repetitive work gets handled, the institutional knowledge becomes operable, and the team's capacity shifts toward the strategic work that drives outcomes.

The work that consumes most of a marketing team's day is the work that gets in the way of the work they were actually hired to do.

FishtankOS is the fix.





Let your teams work where they already work



One of the most consistent failure modes in enterprise AI adoption is the requirement to change how people work in order to access the technology. New platforms require new training cycles, new habits layered on top of existing ones, and a sustained behavioral change effort that most organizations do not have the bandwidth to support. Even when teams push through that initial friction, the adoption ceiling is low, because people will always default to the tools and patterns that feel natural when the pressure is on.

FishtankOS is designed so that the intelligence comes to the team rather than the other way around. It connects to the platforms people are already working in every day, including Microsoft Teams, Slack, SitecoreAI, Jira, Confluence, Salesforce, Monday, and Google, and makes those platforms responsive to plain-language intent. A team member describes what they need and FishtankOS interprets that intent, navigates the relevant systems, applies organizational standards, and delivers the outcome, all without the person leaving the tool they were already working in.

The goal is connecting intelligence to what already exists and putting it to work where teams already operate, inside the platforms and workflows they rely on every day.

This approach has a strategic implication beyond convenience. The highest-friction part of any AI deployment is the gap between where the AI lives and where work actually gets done. When that gap is eliminated, the investment starts delivering value immediately rather than after a lengthy change management effort, and it sustains that value because it's embedded in the natural flow of how a team already operates.

In practical terms, this means a content strategist can request a fully authored article from a source document without opening the CMS. A marketing manager can pull an analytics report on underperforming pages without navigating Google Analytics. A campaign lead can trigger a multi-step workflow spanning content creation, CRM updates, and project tracking from a single message in Teams. The expertise required to operate the underlying systems is handled by the intelligence layer, so the person making the request only ever needs to describe what they need.



The operating model has changed


Welcome to your organization's new operating system.

Enterprise technology has historically been organized around the assumption that people are the integration layer. Digital platforms sit on the edges of an organization, and the humans in the middle carry the responsibility of moving information between them, making judgment calls about how one system's output should translate into another system's input. This model made sense when automation meant completing individual tasks within a single platform. Intelligent orchestration across an entire connected environment requires something fundamentally different.

The architecture that FishtankOS introduces redraws that model. Each connected system carries its own intelligence layer, encoding both how that platform works and how this specific organization uses it. FishtankOS sits at the center, reading intent and routing execution simultaneously across whichever systems are needed to complete it. A single instruction can activate multiple platforms at once, completing in seconds what would otherwise require a person to move through four or five tools in sequence, reconstructing context at every step.

The four principles on the right hold this model together, and each one breaks a rule that enterprise software has lived by for decades.

Figure 11 — FishtankOS sits at the center of the model, intelligence embedded in every connected system



- 7 Every connected system carries its own intelligence layer, so context travels with the task across every platform.
- 7 FishtankOS acts as the central coordination layer, reading intent from a plain-language request and routing execution across whichever systems are needed.
- 7 A single trigger can activate multiple systems simultaneously, collapsing multi-step workflows into one action.
- 7 The output is delivered directly to the user inside the tool they are already working in, as part of their natural workflow.



What the intelligence layer actually contains

This intelligence layer is what separates FishtankOS from any other generic AI integration.

A large language model connected to a CMS without meaningful context produces output that sounds plausible but requires significant human review before it's usable, absorbing the productivity gain in the editing process. The intelligence layer closes that gap, giving the system enough context about how a platform works and how a specific organization uses it to produce output that is ready to act on from the start. Every system connected to FishtankOS gets two complementary layers of intelligence that work together to make execution both accurate and organizational.

When both layers are active, the gap between AI output and publishable, actionable content closes substantially. FishtankOS produces work that arrives correctly structured, voiced, and tagged, ready to publish the moment it's generated and eliminating the cleanup cycle entirely.

The intelligence layer works the same way across every system in the FishtankOS ecosystem, pairing two complementary layers that always operate together. Salesforce gets foundational intelligence about how pipeline data is structured and organizational intelligence about what signals matter to its sales motion. Jira gets foundational intelligence about workflow states and routing and organizational intelligence about how the team prioritizes work. The pattern holds across every connected platform, making each one genuinely useful. What each layer actually holds breaks down like this.

Foundational intelligence


This is the structural knowledge of how a given platform operates. For SitecoreAI, that spans site architecture, component libraries, workflow states, field types, image storage conventions, metadata schemas, and the mechanics of finding, creating, updating, and publishing content at scale. It stays consistent across every organization on the platform, letting FishtankOS navigate and operate the tool with the precision of an experienced administrator. That expertise is carried by the system itself, rather than by the person making the request.

Organizational intelligence

This is where a specific organization's own identity and operating standards are encoded. This layer holds its brand voice guidelines, image specifications, content templates, landing page structures, metadata conventions, press release formats, tagging taxonomies, and the patterns that set this organization's output apart from anyone else's. It is a working representation of institutional knowledge, one that makes every output feel like it came from this team, reflecting this brand and meeting this organization's own standards for quality.

Figure 12 — The FishtankOS intelligence layer

Your team doesn't change a thing. FishtankOS connects intelligence to the tools they already use every day.





Skills across the ecosystem

The bottleneck has shifted from execution to imagination, and that changes what enterprise digital teams are capable of entirely.


The practical expression of FishtankOS intelligence is what we call skills. Specific, trained capabilities that operate within and across each connected platform, informed by the foundational and organizational intelligence layers described above. Each skill is a purpose-built action designed around the real tasks that enterprise digital teams perform every day, trained with enough context to complete those tasks to organizational standards without requiring deep expertise in the underlying systems.

FishtankOS is the agentic layer that orchestrates across all of these systems simultaneously, powered by Anthropic and purpose-built for the complexity of enterprise digital operations. The significance of that orchestration capability is difficult to overstate. When a team member makes a single request that touches content creation, CRM data, project tracking, and analytics in one workflow, they are not simply saving time on a handful of individual tasks. They are collapsing entire processes that previously required multiple people, multiple tools, and multiple rounds of back-and-forth communication into a single, coherent action.

Content consolidation at scale becomes achievable because FishtankOS can analyze performance data across hundreds of pages, identify which content is cannibalizing which, determine the right consolidation strategy, author a new definitive piece drawing from the original sources, and publish it with the correct metadata and attribution — all in one workflow a single team member can initiate and oversee in a fraction of the time a traditional production process would require. Landing page creation becomes a brief-to-publish workflow rather than a multi-week production effort involving designers, developers, copywriters, and CMS specialists.

When every system in an organization's technology ecosystem is connected and orchestrated through a single intelligence layer, the constraint on what a team can accomplish shifts entirely to the quality and ambition of the outcomes they choose to pursue. The limit is no longer the tooling, the headcount, or the number of specialists a task has to pass through. That is a fundamental shift in what enterprise digital teams are capable of, expanding the ceiling on the work itself, and it is exactly what FishtankOS is built to deliver.

Figure 13 — FishtankOS skills map across the ecosystem, from foundational to organizational intelligence



FishtankOS is the agentic layer orchestrating across all systems simultaneously.

Agentic workflows • Context aggregation • Decision automation • Powered by Anthropic



04

Act four: The benefits

↗ Scale at full speed

The system carries the institutional knowledge, so the people don't have to.

↗ Access from anywhere

On-the-go intelligence delivered to where your teams already live.

↗ Speed and efficiency

Requests do not queue up. Campaigns do not wait on bottlenecks.

↗ Built-in expertise

Every connected system is usable by everyone. No specialists required.

↗ Unlimited possibilities

Connected systems unlock a new level of what's possible.

The most significant organizational benefit of a connected intelligence layer is not that it makes existing work faster. It is that it makes work possible that previously was not, because the execution cost of ambition has fundamentally changed.

The five benefits in this section are the direct organizational consequences of connecting intelligence across your full technology ecosystem. Each one reinforces the others, and together they compound to create a meaningful shift in what enterprise digital teams are capable of.



Benefit 01

Scaling

Every new hire performs to standard from day one.

For most teams, capability lives in people — the CMS architect, the metadata owner, the one strategist who knows the brand voice cold. It walks out the door every time someone leaves. The reflex is to train more, but the data says that doesn't work: 84% of companies still haven't redesigned jobs around AI, choosing to educate the workforce instead of rearchitecting how work gets done. Fluency rises. The structure never changes.

FishtankOS moves that knowledge out of people and into the operating model. The CMS architecture, the metadata schema, the brand voice — all of it lives in the organizational intelligence layer. A new contributor doesn't ramp for weeks; they describe what they need, and the system applies every standard on their behalf, correctly, the first time.

The training cycle breaks. The quality floor rises and stays risen. The organization scales its contributor base without the governance overhead that's always made distributed authoring fragile — because the knowledge no longer depends solely on who's in the room.

*The old, people-dependent model (white)
versus a new system that actually scales (yellow).*

People-dependent


Train everyone on the platform.
Consolidate when it gets messy.
Someone leaves, knowledge dies.
Start over, increase investment.
Same broken cycle perpetuates.

Performance at scale

New people ramp up and excel by describing what they need.
Knowledge stays in the system, instead of with the person.
Intelligence embedded directly in the operating model.
People describe what they want, the system handles the how.
Systems scale without repeating the cycle.

Figure 14 — Redesigning workflows beats retraining new people

Workflow redesign is one of the strongest contributors to business impact of all factors tested.



Note. Figure 14 from QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (n = 1,993). Copyright 2025 by McKinsey & Company.

¹ QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation* (McKinsey & Company, November 2025)

² Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge* (Deloitte, January 2026)





Intelligence delivered to where work already happens

With FishtankOS your team keeps working exactly as they do today, with intelligence delivered directly into the tools they already use.

Benefit 02

Access

Working from anywhere just became the new normal.

Your team shouldn't have to chase around intelligence, it should come to them. Worker access to sanctioned AI tools jumped 50% in a single year, from under 40% to almost 60% of the workforce. Yet among the people who have it, fewer than 60% actually use it day to day, a number that's barely moved.² Access went up. Usage didn't. The tools were handed out, but they stayed where they were installed, behind one more login, on one more platform, in one more place you had to be sitting down to reach.

FishtankOS changes where the intelligence lives. It comes to you through the channels your team is already in (Microsoft Teams, Slack, email), so the full capability of the platform travels with the person instead of staying at the desk. Any device. Any channel. Any time you're ready to act. A content update, an analytics query, a workflow trigger, a campaign adjustment — any of it starts with a single plain-language message, from any device, wherever you happen to be. That collapses the gap between noticing something and acting on it to near zero. A fix spotted on a phone between meetings gets made before it's forgotten. The window in which your organization's intelligence is available stops matching office hours and starts matching real life. Now, intelligence travels with the person, and value builds on the go.

Access from anywhere

On the commute.
Between meetings.
No desktop required.

FishtankOS connects everywhere you work

Reads the request from Teams.
Automatically applies brand standards.
Works across all existing systems.

² Deloitte AI Institute, *State of AI in the Enterprise: The Untapped Edge* (Deloitte, January 2026)




Benefit 03


Speed

Do more faster, and get to the work that matters.

Before FishtankOS



After FishtankOS — do more work that matters




The real opportunity with speed isn't moving faster on the same tasks just so you can do more tasks. It's clearing the shallow layer entirely so your team's best hours go toward the wabout ork that actually moves the business. Execution, assembly, formatting, data entry, context switching, coordination: the shallow work that fills the day is the hidden tax on getting to the work that actually matters. It's why efficiency is the near-universal AI goal (80% of companies list it as a goal) while only about a third are using AI to genuinely transform how the business operates. What you do with the capacity you've freed up is where the real value compounds.

FishtankOS clears the shallow layer. It processes requests across systems at once, collapsing the sequential handoffs and endless context switching that turn a simple task into a multi-hour chain, so execution stops eating the hours your best people should spend thinking.

The real shift is a team that spends its best hours on strategy, creativity, and growth, the work it was actually hired to do.

Figure 15 — Growth and innovation are the new performance benchmarks

Efficiency is table stakes. Objectives in growth and innovation are where high performers are pulling away.



Note. Figure 15 from QuantumBlack AI by McKinsey, *The State of AI in 2025: Agents, Innovation, and Transformation*. (n = 1,993). Copyright 2025 by McKinsey & Company.



Benefit 04

Expertise

Every system running at max, no special expert required.

Every platform in your enterprise stack is now within reach of everyone on your team. Platform expertise is encoded directly into every connected system. The knowledge that used to live with one or two specialists is now available to everyone who can describe what they need. It's why leaders name insufficient skills the single biggest barrier to integrating AI, and why fewer than half of companies have done much, if anything, about it.

FishtankOS closes that gap: anyone on the team can operate any system at an expert level. The only skill that matters is knowing what to ask. Describe the desired outcome and FishtankOS supplies the system knowledge. Pipeline data without a CRM admin, analytics without touching the tool, publishing without learning the CMS.

Expertise stops being about credentials and becomes about clarity. Prompt well, and you can operate any system at an expert level, putting the intelligence locked inside complex platforms within reach of everyone.

Figure 16 — The only skill you need is knowing what to ask

Prompt well, with clarity and intention, and the entire marketing ecosystem operates at a maxed-out level of expertise.



Benefit 05

Unlimited possibilities

**The sky is the limit.
Your only limitations are self-imposed.**

For the first time, the limiting constraint on what a team can achieve is no longer execution. Scaling, access, speed, and expertise (the four benefits before this one) remove what used to cap output; headcount, specialist availability, the slow grind of multi-system work. What's left is imagination.

Under the old model, the most ambitious initiatives stayed on the whiteboard. The execution cost was too high, the specialist coordination too slow, and the multi-system sequencing too fragile for anyone to commit to at scale. A cross-system initiative meant weeks of planning, specialist coordination, and sequential execution no one had capacity for, so teams aimed at what was achievable instead of what was possible. FishtankOS collapses the cost of that complexity, and the initiatives that once felt out of reach become routine.

It changes the question a team asks itself. The question, 'can we execute this?' becomes, 'what's worth doing?' Ambition pays — the organizations seeing the most value are the ones thinking biggest, with high performers 3.6x more likely to be pursuing transformative change. Aim higher, and the system keeps up. The ceiling is ambition, and ambition is yours to set.

What's the next great thing you are ready to create?



➤ AI IN YOUR DXP

**Connect what you have,
transform how you work.**



F I S H T A N K
A Verndale Company

The #1 Sitecore Platinum Partner, dedicated exclusively to the platform.

Fishtank is a Sitecore Platinum Partner operating at the intersection of engineering excellence, wild creativity and extreme attention.

The largest scale, down to the smallest detail.



12
2026 Sitecore MVPs

27
Years in the business

800+
Digital transformations

60+
Sites on SitecoreAI

100%
Sitecore focused

2025 Strategic Excellence Award

2024 Experience Award Finalist

2023 Rising Partner Award



Professional
services



Financial
services



Education



Non-profit



Healthcare



Energy



Manufacturing



Utilities