

FOR SENIOR LEADERSHIP · INSURANCE · TECHNOLOGY · FINANCIAL SERVICES

# The Business-Driven AI Era

A paradigm shift in how enterprises build, test, and operate software — illustrated through a real-world transformation.

**Inside this paper:** Why business-driven AI is the third great shift in enterprise software delivery, how the 2026 AI tool landscape actually breaks down, why "vibe coding" is the wrong answer for enterprise apps, and how Arcline Technologies used InsightStudio and InsightTestBench to reform business app delivery and QA in twelve months — without abandoning GitHub Copilot.

# Contents

<b>01</b>	Executive summary	3
<b>02</b>	Three waves of how software gets built	4
<b>03</b>	Why this is happening now	5
<b>04</b>	The 2026 enterprise AI tool landscape	6
<b>05</b>	The vibe-coding trap, in detail	9
<b>06</b>	What an agentic AI fabric actually is	10
<b>07</b>	Industry context — insurance, tech, finance	11
<b>08</b>	The new role of IT	12
<b>09</b>	Case study — Arcline Technologies	13
<b>10</b>	A pragmatic path to get there	18
<b>11</b>	Four questions to ask your team this quarter	19
<b>12</b>	Why VerticalServe	20
<b>13</b>	Appendices · Decision frameworks · Glossary	21

# 01 Executive summary

Every enterprise has spent the last twenty years optimizing *how* IT delivers software — from waterfall to agile to DevOps. Each generation shortened the cycle but kept the same fundamental shape: **the business asks for an app, IT builds it.**

A third shift is now underway. The combination of capable agentic AI, production-grade no-code fabrics, and the accumulated weight of shadow IT has created the conditions for a profound inversion: **the business side composes the apps themselves, and IT operates the fabric beneath them.**

**The thesis of this paper.** Business-driven AI delivery is the next operating model for enterprise software. It is *not* vibe coding — which is a useful tool for engineers and a dangerous one for business users. It is governed agentic platforms that let business teams compose apps that inherit identity, audit, security, and operations by default.

This shift is real, and it is measurable. The early adopters across insurance, technology, and financial services are reporting 5–10× more apps in production per quarter, weeks-to-minutes time-to-value for departmental workflows, and — counterintuitively — *stronger* IT control, not weaker, because every app inherits the platform's guardrails by default.

**Inside this paper.** We describe the three waves of how enterprise software gets built; survey the 2026 AI tool landscape (GitHub Copilot, Claude Code, Microsoft 365 Copilot, Lovable, v0, ServiceNow, Salesforce, and the VerticalServe stack); explain what an agentic fabric is and why it differs from vibe coding; and walk through how **Arline Technologies**, a mid-size B2B SaaS company, used InsightStudio and InsightTestBench to transform its business app delivery and QA function in twelve months — while keeping GitHub Copilot productively in the IDE for engineering.

## SHADOW IT, TODAY

**35–60%**

of employees use AI tools their org has not sanctioned

## IT DELIVERY BACKLOG

**14 months**

typical wait for a new internal business app at a mid-size enterprise

The companies that recognize this shift early — and pick the production fabric over the vibe-coding shortcut for business apps — will compound their advantage quarter after quarter. The companies that don't will spend the next decade cleaning up another generation of shadow IT.

## 02 Three waves of how software gets built

Every enterprise has lived through two prior shifts in how software gets delivered. The third is happening now.

<p><b>WAVE 1</b></p> <p><b>Waterfall</b></p> <p>≈ 1990 – 2010</p> <p>Business writes requirements. IT scopes a project. Six to eighteen months later, software ships. Predictable on paper, brittle in practice — requirements drift, and the software that ships isn't what was needed.</p>	<p><b>WAVE 2</b></p> <p><b>Agile + DevOps</b></p> <p>≈ 2010 – 2024</p> <p>Iterative delivery. Business and IT share a backlog. Sprint cadence shortens the loop. Better outcomes, but the bottleneck is still <i>developer capacity</i>. Apps that can't get on the roadmap fall to shadow IT.</p>	<p><b>WAVE 3</b></p> <p><b>Business-driven AI</b></p> <p>2025 +</p> <p>Business teams compose apps themselves on a curated fabric of pre-validated AI / data / workflow primitives. IT shifts from delivery to enablement. The bottleneck stops being engineering hours.</p>
--	--	--

The first two waves were *engineering* shifts. The third is an *organizational* shift. It changes who builds, what IT does, and how the company gets work done.

	Waterfall	Agile / DevOps	Business-driven AI
<b>Who builds</b>	IT only	IT, with business in standups	Business, on a curated fabric
<b>Time to first version</b>	6–18 months	4–12 weeks	Hours to days
<b>Bottleneck</b>	Requirements drift	Developer capacity	Capability library curation
<b>IT's role</b>	Implementer	Iterative implementer	Platform / governance owner
<b>App lifecycle</b>	Personal investment	Sprint-tracked	Marketplace + versioned + auditable
<b>Shadow-IT pattern</b>	Spreadsheets	Personal automations	Unmanaged AI subscriptions

## 03 Why this is happening now

---

Three trends converge in 2026 that did not all exist together before:

Trend	What it enables
<b>Capable agentic LLMs</b> Claude Opus 4.7, GPT-5, Gemini 2.5	A natural-language prompt now describes intent accurately enough to generate working workflows, queries, and integrations — without a human translating to code.
<b>Production-grade no-code fabrics</b>	Validated connectors, agents, identity, and workflow primitives that compose into enterprise-grade apps — not demos.
<b>Shadow-IT exhaustion</b>	A decade of unmanaged spreadsheets, personal Power Automate flows, browser-tab ChatGPT use, and unsanctioned SaaS has created risk that leadership is actively trying to consolidate.

---

*A business analyst in 2026 has the same productive capacity, with the right platform, that a junior developer had in 2020 — but with governance baked in. That's the unlock.*

## 04 The 2026 enterprise AI tool landscape

The phrase "AI tooling" has been stretched so far that it has lost useful meaning. Let us be precise. There are six distinct categories of AI-assisted software tooling in enterprise use today — each with a different audience and a different fit.

CAT 1

**Chat assistants.** Claude, ChatGPT, Gemini, Perplexity Pro. General-purpose conversational reasoning for everyone. Stateless; no persistent workflows; no system connectors.

CAT 2

**Productivity copilots.** Microsoft 365 Copilot, Workspace Duet, Salesforce Einstein, Notion AI. Embedded assistance inside Office, CRM, etc. Cannot compose new apps; confined to the vendor surface.

CAT 3

**IDE coding assistants.** GitHub Copilot, Cursor, Tabnine, Codeium. Real productivity gain for developers (35–55% faster on routine code). An *engineering* tool, not a business-enablement tool.

CAT 4

**Agentic coding assistants.** Claude Code, Devin, GitHub Copilot Workspace, Aider. Autonomous multi-file refactors and PR drafting. Senior engineers' tool. Operates on source code; output requires engineering review.

CAT 5

**RISKY FOR ENTERPRISE** **Vibe-coding / prompt-to-app.** Lovable, v0, Bolt, Replit Ghostwriter. Generates working prototypes from prompts. No SSO, no audit, no connectors. The next generation of shadow IT if deployed organizationally.

CAT 6

**Agentic AI fabrics for the enterprise.** InsightStudio, ServiceNow App Engine, Salesforce Agentforce, Microsoft Power Platform with Copilot Studio. Business teams compose governed apps; the fabric inherits identity, audit, security, and operations.

### Where each category fits

Category 6 – Agentic AI fabric (InsightStudio)  
Audience: Business teams  
Output: Production-governed apps in an internal marketplace

Categories 3+4 – IDE & agentic coding (Copilot, Claude Code)  
Audience: Software engineers  
Output: Faster engineering, refactored code, PRs

Category 2 – Productivity copilots (M365, Workspace)  
Audience: Knowledge workers in existing apps

Category 1 – Chat assistants (Claude, ChatGPT)  
Audience: Everyone, for ad-hoc reasoning

— Avoid for the enterprise app stack —

Category 5 – Vibe coding (Lovable, v0, Bolt)  
Risk: The next generation of shadow IT

**The mistake leadership most often makes** is substituting one category for another — using Copilot to "enable business users" (it doesn't; it accelerates developers), using Lovable to ship production apps (it can't; the governance isn't there), or expecting a Category 6 fabric to replace IDE assistants (it won't; engineers still need their tools). The right strategy is *all of the above, in the right places*.

## 05 The vibe-coding trap, in detail

Vibe-coding tools — Lovable, v0, Bolt, Replit Ghostwriter — let anyone describe an app and get working code back. The demos are impressive. The pattern in the enterprise is recognizable: it's the same pattern as personal spreadsheets in the 2000s and personal Power Automate flows in the 2010s. Individual productivity wins; organizational debt accumulates.

Concern	Vibe coding Lovable · v0 · Bolt · Replit	Production AI fabric InsightStudio
Who builds it	An individual, often off-hours	A business team with explicit ownership
What gets produced	Working code in a personal cloud	A governed app in the org's marketplace
Security review	None — or after the fact	Built into the fabric: SSO, secrets, encryption, audit
Compliance	Unknown until audit time	SOC 2 / GDPR / HIPAA controls inherited per-app
Identity & access	Whoever has the URL	Org SSO, group grants, request flows
Data egress	Often public APIs with company data	Stays inside your VPC; egress policies enforced
Maintenance	Orphaned when builder leaves	Platform team owns runtime; apps are org assets
Discoverability	Word-of-mouth; Slack thread; lost	Internal marketplace; access requests
Integration	Copy-paste, brittle	First-class connectors to ERP, CRM, warehouse, identity
TCO over time	Low at start; explodes when IT cleans up	Predictable; managed infrastructure
Risk profile	Uncontrolled, unaudited, unobservable	Bounded — every action logged, every app governed

The vibe-coding path delivers the first 10 apps fast. It also creates the next decade's tech debt. **The point of a production AI fabric is to deliver the first 10 apps fast *and* the next thousand sustainably.**

## 06 What an agentic AI fabric actually is

Most "no-code" platforms in the market are workflow tools. They each do one thing — process automation, or chatbots, or report generation — and become another silo for IT to integrate.

An **agentic AI fabric** is the integrated layer beneath every app a business team would want to build. It contains:

**01 A capability library.** Pre-validated agents (extractors, lookups, summarizers), connectors (Salesforce, SAP, Snowflake, SharePoint, Outlook, custom APIs), and tools (databases, filesystems, OCR, vision). Sandboxed and audited.

**02 A composition surface.** A visual + prompt-driven builder where a business user describes an app and the fabric assembles the workflow, the form, the output panels, the integrations.

**03 A runtime.** Every app runs through the same engine. Identity, secrets, rate limiting, audit, observability are runtime concerns — not app concerns.

**04 A marketplace.** The catalog of apps the org has produced. Teammates discover apps like SaaS subscriptions, but inside the org's governance perimeter.

**05 A governance plane.** Admin controls for data reach, model allowlists, app approval, SIEM export.

Apps composed on a fabric inherit all of this. The business team focuses on *what the app should do*; the fabric handles *how it does it safely*.

**InsightStudio** is one such fabric — for business app composition. **InsightTestBench** is the parallel fabric for the QA function. Both run on the open **InsightWorker** engine and the same \$15 bundle format, so an enterprise that deploys one gets the operational shape of the other essentially for free.

## 07 Industry context

---

The business-driven AI shift will be felt differently across industries. Three of the largest enterprise verticals warrant specific mention.

### **Insurance**

Insurance enterprises run on long-tail department-specific workflows — broker submission triage, loss-run analysis, policy comparison, claims intake routing, reinsurance reporting — that historically required either a custom-built tool or a manual spreadsheet process. The combination is exactly what an agentic fabric is designed for: each workflow becomes an app composed by the line of business that owns it, integrated with the policy admin system, claims system, and SharePoint, governed centrally.

Carriers that have moved early typically retire 30–50% of their shadow spreadsheets in the first year, and bring 40–80 governed line-of-business apps into their internal marketplace by month 18.

### **Technology / B2B SaaS**

Tech companies have the engineering capacity to build internal tools themselves — and yet typically still have a long backlog of unbuilt ones, because engineering capacity is rightly reserved for the customer-facing product. The result is a familiar pattern: customer success running spreadsheets, finance using personal Power Automate, product ops with ChatGPT browser tabs over private customer data.

Tech companies are also the most likely to be early adopters of GitHub Copilot and Claude Code inside engineering — which is the *right place* for those tools. The question is how to provide an equally sophisticated platform for the *non-engineering* side. The case study in section 9 is a representative example.

### **Financial services**

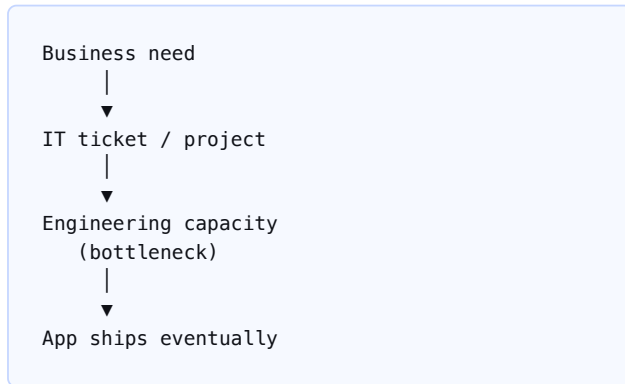
Financial firms — banks, asset managers, broker-dealers — have the highest governance bar. Every internal app needs to pass audit, model risk management, and regulatory review. The vibe-coding path is essentially closed; a business-side employee cannot legitimately deploy a vibe-coded app to a trading floor.

The agentic fabric path is the only path that maps cleanly to a financial firm's governance posture: every app inherits the same audit log; every model call is captured for MRM review; every data access is authenticated; the runtime is fully on-prem or in the firm's own VPC.

## 08 The new role of IT

Senior leaders often interpret "business builds apps" as "IT is sidelined." The opposite is true. IT's role becomes *more strategic*, not less.

### Before — IT as builder



IT's strategic energy gets consumed by the long tail of departmental apps. There is rarely time for the platform thinking that drives organization-wide leverage.

### After — IT as enabler



The long tail of departmental apps moves off IT's roadmap. IT's calendar gets liberated to do the work it was always supposed to do.

### IT's new focus areas

- **Platform reliability** — uptime, observability, capacity planning
- **Capability curation** — vetting new connectors, integrations, models before they enter the library
- **Policy + governance** — who can see what data, what crosses organizational boundaries, audit + SIEM integration
- **Risk + compliance partnership** — collaborating with InfoSec, Legal, and Audit on the *one* runtime, not on a thousand individual apps
- **Strategic data initiatives** — the platform projects that move the needle for the whole company

## 09 Case study — Arcline Technologies

*The following case study is a composite, drawn from real customer engagements. Identifying details have been changed; the operational patterns and outcomes are accurate.*

**Arcline Technologies** is a mid-size B2B SaaS company. ~1,800 employees, ~210 engineers, ~28 QA, multiple customer-facing dashboards plus a long tail of internal apps for sales operations, customer success, finance, and people ops. Public-cloud-native; AWS-based; Okta-managed identity.

### Starting state — Q4 2025

Arcline's engineering team had standardized on GitHub Copilot two years earlier and had begun a pilot of Claude Code for senior staff engineers. Coding productivity was real and measurable. But outside engineering, the company was running on the familiar mix:

- **Business teams in queue.** The IT delivery team had a 14-month backlog of internal-app requests from sales ops, customer success, finance, and people ops. The vast majority would never be built.
- **Shadow IT everywhere.** ~40 active Power Automate flows owned by individuals. ~20 personal Zapier accounts. Customer success was running churn scoring in a maintained-by-one-analyst Google Sheet. Finance was using ChatGPT browser tabs to draft variance commentary against quarterly numbers — pasting in raw P&L data.
- **QA bottleneck on every release.** 28 QA engineers writing Playwright tests by hand. Regression: two weeks of QA effort per release. Flaky tests had reached the point where 30% of CI failures were dismissed as "probably the test."
- **Vibe-coding seepage.** A handful of product managers had started using Lovable and v0 to prototype features for customer demos. Two of those prototypes had been shown to customers as if they were product — creating commitments engineering didn't know about.

*The CIO recognized the pattern. The CTO recognized the pattern. The CFO was looking at AI spend across the company — fragmented across nine vendors — and asking the obvious question: what is our consolidated AI strategy?*

### PHASE 1 · Q1 2026

#### InsightStudio pilot — sales ops & customer success

Sales ops composed five apps in the first six weeks: lead triage assistant, account research briefer, pipeline anomaly digest, quote-to-contract draft, renewal risk scorer. Customer success composed four: ticket triage + escalation router, health score aggregator, QBR prep pack, churn signal alerter. Each app inherited Okta SSO, audit logging, and the same observability the IT team used for everything else.

### PHASE 2 · Q2 2026

#### InsightTestBench reform of QA

Each of seven major customer-facing surfaces was bootstrapped by its product owner. The bench logged in via Okta service account, crawled every page, generated 70–100 cases per surface. Product owners refined cases by chat. Scheduled regression ran nightly with RCA-classified failures pinged only on real regressions. Once Arcline wired its GitHub repo into settings, the bug-fix agent began opening PRs against confirmed product bugs with the failing test, the patch, and a Jira ticket linked.

### PHASE 3 · Q2 2026

#### Coexistence with Copilot and Claude Code

GitHub Copilot stayed in the IDE. Claude Code stayed for senior engineers' repo-wide refactors. Neither needed to be displaced. The stack acquired clear layers — Cat 3+4 for engineers, Cat 6 (InsightStudio / InsightTestBench) for business and QA. Lovable, v0, and Bolt were explicitly removed from the sanctioned list. The two product managers who had shown v0 prototypes to customers ported them to InsightStudio within a week.

### PHASE 4 · Q3 2026

#### Marketplace + shadow-IT consolidation

The InsightStudio marketplace reached 63 governed apps. Cross-team reuse began — customer success's ticket triage was forked by support engineering for incident triage; finance's variance commentary drafter was adopted by FP&A. 22 of 40 Power Automate flows migrated; 16 sunsetted. 18 of 20 personal Zapier accounts closed. Browser-tab ChatGPT use for finance work retired.

## Results after 12 months

Metric	Before	After 12 months
IT departmental-app delivery backlog	14 months	4 months
Business-led apps in production	0	63
QA effort per release cycle	2 weeks × 28 engineers	~2 days × 25 engineers
Regression flake rate	~30% of CI failures dismissed	< 8%
Critical regressions caught pre-release by the bench	n/a (manual catches only)	11 in the 12 months
PRs opened by the bug-fix agent + merged	0	47 of 71 opened (66% merge rate)
Documented shadow-IT workflows	60 +	4 (audit exceptions)
Annual fragmented AI spend	\$3.2M across 9 vendors	\$1.4M (Studio + Bench + retained Copilot)

The CFO's "consolidated AI strategy" question had an answer. The CTO's engineering org kept Copilot and Claude Code. The CIO's IT organization transitioned from delivery to platform — and was growing the *platform* team rather than the delivery team. Most importantly: **the business teams could ship.**

# 10 A pragmatic path to get there

---

The Arcline pattern is repeatable. The recommended sequence:

## PHASE 1 · MONTHS 1–3

### Pilot

Pick 1–2 business teams with a clear, painful, non-mission-critical workflow. Stand up InsightStudio in a non-production environment. Train 2–3 builders. Success metric: a working, governed app shipped in under two weeks, without writing code.

## PHASE 2 · MONTHS 2–5 (PARALLEL)

### QA reform

Pick 1–2 product surfaces. Stand up InsightTestBench. Bootstrap. Move scheduled regression off manual Playwright; wire RCA + webhook. Wire the bug-fix-agent flow to a sandboxed branch. Success metric: product owners drive the test suite directly.

## PHASE 3 · MONTHS 4–9

### Marketplace + shadow-IT consolidation

Move both products to production-grade footprint in your VPC. Onboard 3–5 more business teams. Audit existing shadow IT (Power Automate, Zapier, personal AI subscriptions); categorize: migrate, sunset, exception. Establish IT's platform-operations rhythm.

## PHASE 4 · YEAR 2+

### New default

Business-driven app delivery becomes default for non-mission-critical workflows. Marketplace catalog reaches 100+ apps. Cross-team app reuse becomes routine. New business teams onboard in days, not the months traditional IT delivery required.

The risk profile is bounded — Phase 1 can be done with a single VM in a non-production network. The decisions are reversible at every stage.

# 11 Four questions to ask your team this quarter

---

## 1 · How many requests for new internal apps did our IT team decline, defer, or deprioritize last year?

That number is the size of your shadow-IT surface and the opportunity for business-led delivery.

## 2 · What percentage of our employees are using AI tools we haven't sanctioned?

That number is your inverse measure of platform adoption — and your direct measure of unmanaged risk. Anonymous surveys typically reveal 35–60%.

## 3 · Among the AI tools we *have* sanctioned, who is each one for?

GitHub Copilot is for engineers. Claude Code is for senior engineers. M365 Copilot is for productivity-inside-Office. None of these replaces a Category 6 fabric for the business app problem. If leadership's mental model is "we have AI covered because we deployed Copilot," that's an alarm bell.

## 4 · What would change if our business teams could build production-grade, governed apps without an IT ticket?

This is the question that gets to the strategic value of the shift, not just the cost savings.

# 12 Why VerticalServe

The technology to do business-driven AI delivery at enterprise grade has only existed for about a year. Most enterprises will adopt it in the next 18–24 months regardless of vendor. The first movers will be 12 months ahead of the second movers on app catalog size, internal capability, and IT operational maturity.

The **VerticalServe stack** is built for the enterprise reality:

## INSIGHTSTUDIO

The agentic fabric for **business app composition**. Visual builder + prompt-to-app + marketplace + governance.

## INSIGHTTESTBENCH

The agentic fabric for **QA**. Bootstrap suites from a brief, vision-grounded cases, AI RCA, scheduled regression watchdog, bug-fix agent that opens PRs.

## INSIGHTWORKER

The open **engine** beneath both. Available standalone if you want to compose your own surfaces. Every \$15 bundle is plain YAML you can read, fork, extend.

## All three share the enterprise posture

- **Self-hosted** in your VPC. Data never leaves your environment.
- **Open at the bundle level**. Every primitive is a \$15 YAML bundle. No proprietary lock-in.
- **Identity-aware** from day one. Okta, Azure AD, Google Workspace via OIDC.
- **Audit-first** runtime. Every app run, tool call, and model invocation captured. SIEM webhook export.
- **Per-org governance**. Model allowlists, regional pinning, per-team quotas, per-app reviewers.
- **Operationally simple**. Docker-compose to start; Kubernetes/Helm for production. Standard MySQL. Standard observability.

The companies that move first will define what "business-driven AI" looks like inside their industry. They will set the expectation that internal teams ship apps in days. They will retire shadow IT systematically rather than fighting it indefinitely. Their IT organizations will be in the position they have always wanted to be in — **enabling business outcomes, not delivering departmental tickets**.

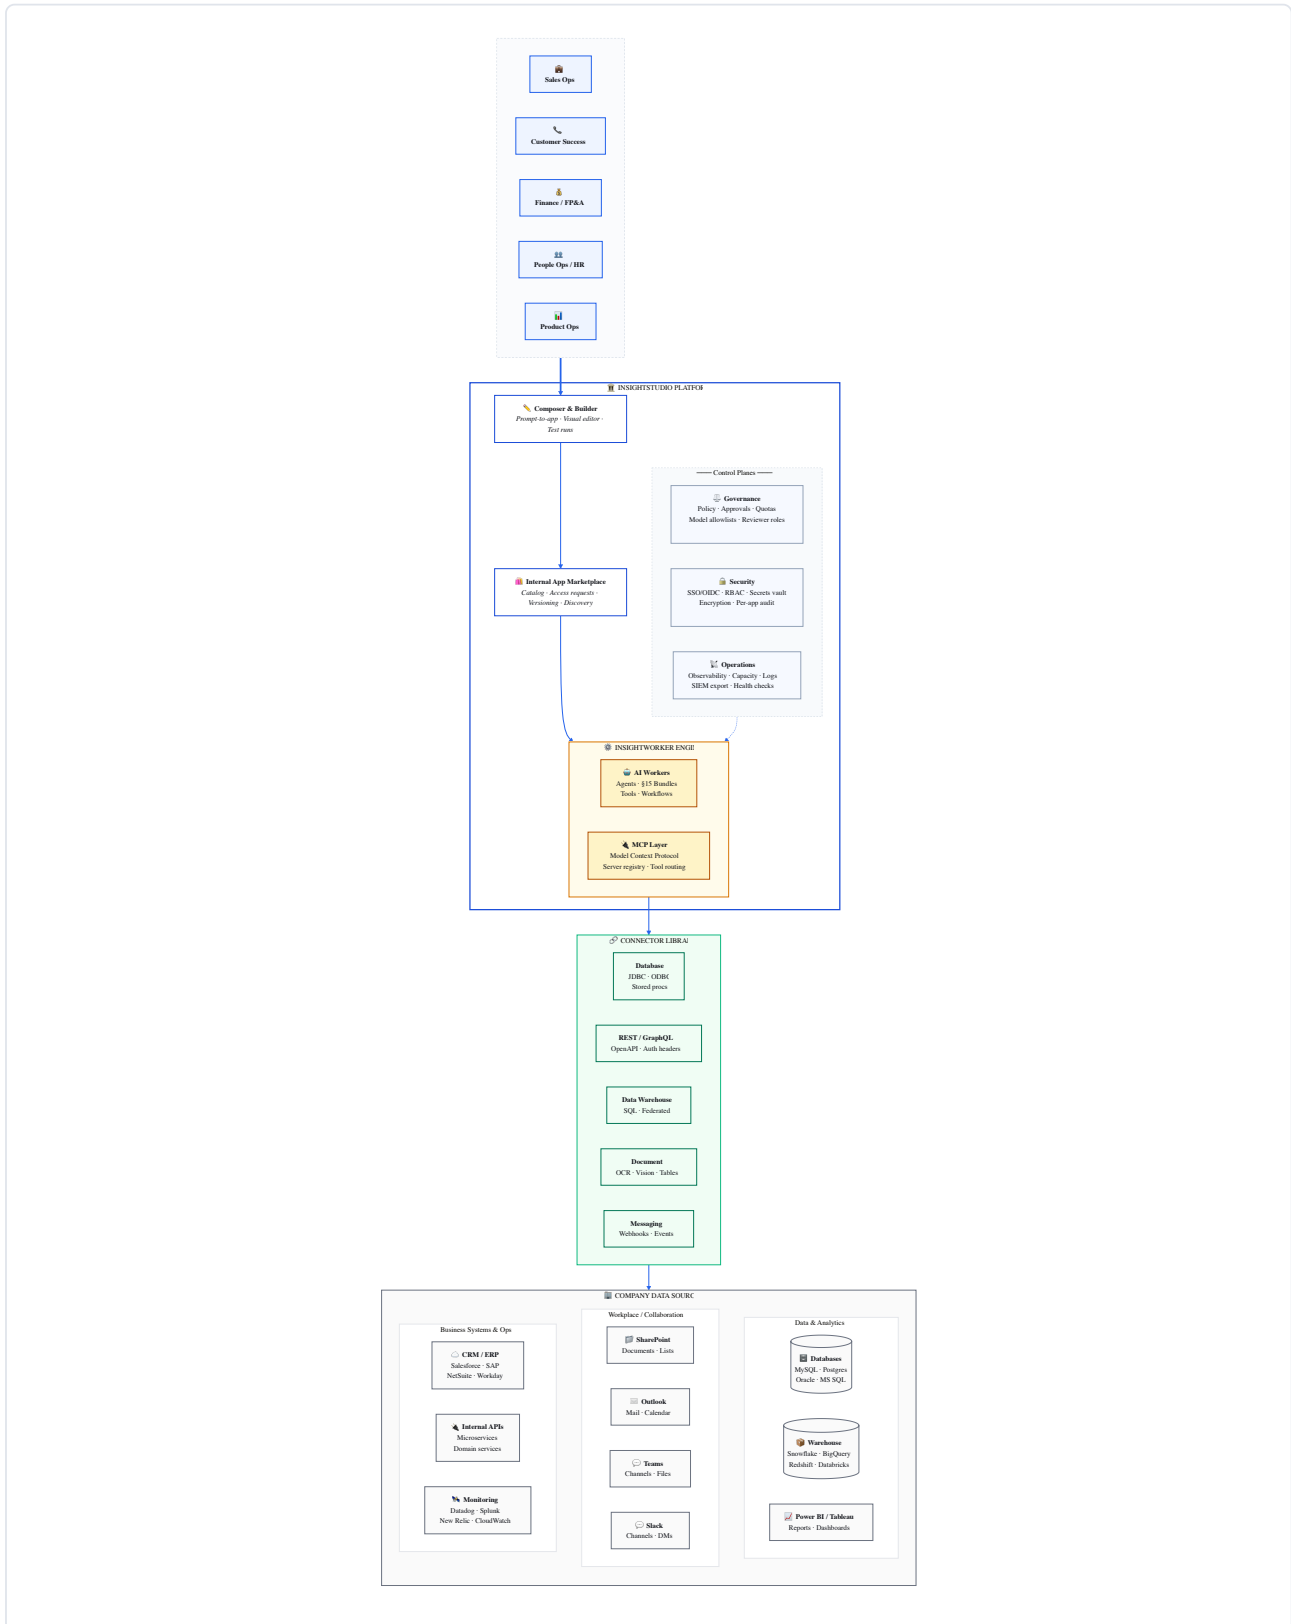
# A Appendix 2026 enterprise AI tool decision tree

Match the right category of tool to the right audience and problem.

If your need is...	The right category is...	Specific tools
Faster code completion for engineers	IDE coding (Cat 3)	GitHub Copilot, Cursor, Codeium
Multi-file refactors / agentic code changes	Agentic coding (Cat 4)	Claude Code, Devin, Copilot Workspace
Help inside Outlook / Word / Excel	Productivity copilots (Cat 2)	M365 Copilot, Workspace Duet
Ad-hoc reasoning, drafting, research	Chat assistants (Cat 1)	Claude, ChatGPT, Gemini
Throwaway prototype for a design review	Vibe coding (Cat 5, caveats apply)	Lovable, v0, Bolt
<b>A business team to ship a governed, persistent, integrated app</b>	<b>Agentic AI fabric (Cat 6)</b>	<b>InsightStudio</b> , ServiceNow App Engine, Salesforce Agentforce, Power Platform
<b>Your QA team transformed into AI-driven test operators</b>	<b>Specialized agentic fabric (Cat 6)</b>	<b>InsightTestBench</b>
You want the open engine under all of the above	—	<b>InsightWorker</b>

# B Appendix Reference architecture

The end-to-end architecture: business teams compose apps at the top; InsightStudio's control planes enforce governance, security, and operations; the InsightWorker engine routes agents and tools through MCP; a connector library bridges to every category of company data source.



IT operates the platform, control planes, engine, and connector library (the bottom four layers). Business teams operate the composer + marketplace (the top two). The marketplace is the meeting point where governed apps surface for the rest of the organization.

## C Appendix Glossary

---

<b>Agentic AI</b>	AI systems that plan and execute multi-step actions autonomously, calling tools and reasoning over their results, rather than producing one response per prompt.
<b>\$15 app bundle</b>	InsightWorker's open format for a packaged AI application: a YAML manifest, a workflow spec, an optional widget UI description. Portable across InsightStudio, InsightTestBench, and the CLI.
<b>Capability library</b>	The curated set of pre-validated connectors, agents, and tools available to apps composed on a fabric.
<b>Marketplace</b>	The internal catalog of apps an organization has composed; enables discovery, access requests, governance.
<b>Production AI fabric</b>	A Category 6 platform that lets business teams compose governed apps with inherited identity, audit, security, and operational guardrails.
<b>RCA agent</b>	InsightTestBench's failure-classification agent. Reads case spec + execution log + page state and verdicts each failure as test design, product bug, or environmental.
<b>Shadow IT</b>	Software, automations, or AI usage in the organization that has not been formally sanctioned, governed, or audited.
<b>Vibe coding</b>	The practice of generating working code from a prompt using consumer-tier tools (Lovable, v0, Bolt, Replit); produces prototypes, not production apps.

NEXT STEP

# Let's walk this through against your environment.

The Arcline arc — pilot, QA reform, marketplace consolidation — is repeatable for an insurance carrier, a B2B SaaS company, or a financial firm. The decisions are reversible at every phase, and the risk profile is bounded.

We would be glad to share a 90-day pilot plan tailored to your industry, your current AI tooling, and your existing shadow IT footprint.

[Request a 90-day pilot conversation](#)

**VerticalServe Inc.** [contact@verticalserve.com](mailto:contact@verticalserve.com)

---

**InsightStudio** [insightaistudio.com](https://insightaistudio.com)

---

**InsightTestBench** [insighttestbench.com](https://insighttestbench.com)

---

**InsightWorker** [insightworker.com](https://insightworker.com)

---

© 2026 VerticalServe Inc. All rights reserved.