THE AI ARCHITECT'S GUIDE TO 2026 AGENTS

Where They Live. What They Do.
What You Can Leverage Today.

A systems-level field guide for small business owners
written from 20+ years of enterprise architecture practice

BrotherILA.com • brandt.delany@gmail.com • 346.589.9888

PREFACE

Why an Architect Wrote This

Most AI content right now is written by marketers or enthusiasts. They know the tools. They know the demos. What they rarely know is how systems actually connect — where data lives, what breaks at scale, and which decisions made in month one create expensive problems in month six.

I spent over twenty years building enterprise systems across oil and gas, petrochemical, and industrial technology. Not writing about them — deploying them. Watching what works under real operating conditions, real data loads, and real organizational constraints. That background is exactly what small business AI conversations are missing.

This guide is not about any single tool. It is about the architecture — the map of agent types that exist right now, where they live in the real world, who the major players are, and what a small business owner can actually pick up and use today versus what belongs in an enterprise budget.

The reason this taxonomy matters: you cannot make good technology decisions without knowing what category of problem you are solving. A business owner who says "I need an AI agent" is like a contractor saying "I need a tool." Which tool depends entirely on what you are building.

This guide gives you the vocabulary and the map. Once you understand where each agent type lives and what it actually does, you stop being a buyer at a vendor's mercy and start being an architect making deliberate decisions.

The architect doesn't wire the electrical or pour the concrete. The architect knows why the load-bearing walls go where they go.

EXECUTIVE SUMMARY

The Eleven Agent Types at a Glance

There are eleven distinct types of AI agents operating in business systems today. Each has a different architecture, a different role, and a different risk profile. Most businesses need a combination — not a single agent, and not all eleven.

#	Agent Type	Summary
1.	Reactive / Autonomous Worker	Trigger → action. No memory, no judgment. Zapier, Make, Shopify Flow.
2.	Planned Agent	Breaks goals into steps and executes. CrewAI, LangGraph, AutoGen.
3.	Policy-Based Agent	Acts inside approved rules and guardrails. ServiceNow, IBM Watson, Rippling.
4.	Goal-Based Agent	Given an outcome, chooses its own path. Google DeepMind, OpenAI, Cohere.
5.	Event-Driven Agent	Sleeps until triggered. Stripe webhooks, PandaDoc, Shopify Flow.
6.	Optimizer Agent	Math and scoring select the best move. Pricefx, Google Smart Bidding, Amazon Repricer.
7.	Utility Agent	Ranks by preference and satisfaction, not just pass/fail. MadKudu, Clearbit/HubSpot.
8.	Human-in-the-Loop Agent	Recommends; a human approves high-risk actions. UiPath, Scale AI, Workato.
9.	Memory / Continuity Agent	Remembers context over time. Intercom, Pinecone, Notion AI, Fireflies.ai.
10.	Swarm / Multi-Agent System	Many agents cooperate. CrewAI, AutoGen, Amazon Bedrock Multi-Agent.
11.	Control-System Agent	Inputs → constraints → decision → action → feedback → repeat. Siemens, Google DeepMind.

For most small businesses, the right starting stack is three or four agent types working together — not eleven. Knowing the full map lets you choose deliberately.

THE GUIDE

The Eleven Agent Types

01 **Reactive / Autonomous Worker** **Small Business → Mid-Market**

Watch this. Do that. Report back.

The simplest agent class. No strategy layer, no memory, no judgment. It receives a trigger and executes a predefined action. Think of it as a very reliable drone — it does exactly what it is told, every time, without ever asking why. The backbone of most small business automation today.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
Zapier, Make.com, Microsoft Power Automate, Shopify Flow	A Zapier zap watches a Google Form. When someone submits, it sends a confirmation email, logs the entry to a Google Sheet, and notifies you via Slack. No memory. No decision. Trigger fires, action runs.

02 **Planned Agent** **Mid-Market → Enterprise**

Break the goal down. Build the roadmap. Execute.

A planner takes a high-level objective and decomposes it into ordered steps, assigns subtasks, monitors progress, and adapts when something breaks. More intelligent than a reactive agent — but more moving parts means more places for things to go wrong. Worth the complexity for multi-step workflows.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
CrewAI, LangGraph, AutoGen (Microsoft)	A CrewAI pipeline receives a sales lead. One agent researches the company, another drafts a personalized outreach email, a third schedules the follow-up task in HubSpot. Each step is assigned, monitored, and handed off.

Goal-Based and Planned agents are often two halves of the same brain — one defines success, the other builds the roadmap to get there.

03 **Policy-Based Agent** **Mid-Market → Enterprise**

Rules drive behavior. The agent acts inside approved lanes.

This agent does not invent or improvise. It operates strictly within predefined rules, constraints, and guardrails. Strong for regulated industries, brand-sensitive businesses, and any system where consistency and compliance matter more than creativity. The Governor of your AI stack.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
Salesforce Einstein, ServiceNow, IBM Watson, Rippling	A ServiceNow workflow routes IT tickets strictly by category, priority, and SLA rules. The agent never

	improvises. A P1 incident always goes to the on-call engineer. A billing question always goes to finance. No exceptions.

04 **Goal-Based Agent** **Enterprise / Research**

Given an outcome, not just a task. It finds its own path.

Instead of following steps, this agent is handed a target state and chooses its own actions to reach it. It evaluates options, selects moves, and adjusts based on results. This is where AI starts to look less like automation and more like decision-making. Powerful — and requires careful guardrails.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
Google DeepMind, OpenAI (Operator-style agents), Cohere	Google's AlphaFold was given one goal: predict protein structures. It was not given a step-by-step method. It chose its own computational path to achieve an outcome that had eluded biochemists for fifty years.

Goal-Based and Planned agents are often deployed together — one sets the target, the other builds the execution plan.

05 **Event-Driven Agent** **Small Business → Enterprise**

Asleep until something happens. Then it moves fast.

Event-driven agents are dormant until a specific trigger fires — inventory drops below a threshold, a form is submitted, a payment clears, a contract is signed, a customer complains. The moment the event fires, the agent wakes up and acts. Extremely efficient because it only consumes resources when needed.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
Zapier, Make.com, AWS Lambda, Shopify Flow, Stripe Webhooks, PandaDoc, Linear	Shopify Flow watches inventory. When a product drops below 10 units, it automatically creates a reorder request and notifies the buyer. Or: when a PandaDoc contract is signed, it fires a webhook that creates a project in Linear, updates HubSpot to Closed Won, and generates the invoice.

06 **Optimizer Agent** **Mid-Market → Enterprise**

Scoring and math find the best move among options.

These agents use mathematical models, scoring functions, or algorithms to evaluate multiple possible actions and select the highest-value option. Excellent for problems with clear metrics: price, margin, efficiency, click-through rate, delivery time. The engine behind most AI-powered pricing, bidding, and scheduling systems.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
Prisync, Pricefx, Google Ads Smart Bidding, Amazon Repricer	Amazon's repricer adjusts seller prices hundreds of times per day based on competitor pricing, inventory levels, and

	margin rules — automatically finding the price that maximizes revenue without dropping below the floor.

07 **Utility Agent** **Mid-Market → Enterprise**

Not just pass/fail. It asks: is this worth acting on?

The Utility Agent goes beyond binary success/failure to evaluate choices by preference, satisfaction, and weighted criteria. It ranks options using a utility function — a mathematical representation of "how good is this outcome given all factors?" This is the agent that filters out technically correct answers that are practically useless.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
MadKudu, Clearbit (now HubSpot), Phenom People	MadKudu scores inbound leads not just on fit, but on intent signals, timing, and likelihood to convert. A lead that matches your ICP but showed up at 11pm on a Sunday with a free email address scores low utility. It never hits a rep's queue.

Classic example: a garage sale has the vintage camera you want (goal met) — but it's 20 miles away and raining (low utility). The agent does not alert you.

08 **Human-in-the-Loop Agent** **All Levels**

The agent recommends. A human approves before it executes.

This architecture inserts a human checkpoint before high-risk actions are taken. The agent handles research, drafting, and analysis — the human approves the final action. Essential for money movement, legal documents, medical decisions, and brand-sensitive communications. Not a limitation — a feature. The right design for decisions where being wrong is expensive.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
UiPath, Scale AI, Labelbox, Workato	UiPath's automation processes insurance claims end-to-end until it hits an exception it cannot confidently resolve. It flags the exception, routes it to a human reviewer with full context, waits for approval, then resumes. The human only sees what requires judgment.

09 **Memory / Continuity Agent** **Small Business → Enterprise**

It remembers. Every conversation. Every decision. Every preference.

These agents maintain persistent context over time — customer history, business goals, prior decisions, brand voice, and stated preferences. Without memory, every interaction starts from zero. With it, the agent accumulates knowledge and becomes more useful the longer it operates. The foundation of genuine AI-powered customer relationships.

MAJOR PLAYERS	REAL-WORLD EXAMPLE

Intercom, Salesforce Einstein, HubSpot, Zendesk, Pinecone, Notion AI, Fireflies.ai, Otter.ai	Intercom's AI assistant remembers every prior conversation with a customer — their previous issues, their product tier, their stated preferences — and uses that history to personalize every response and route escalations correctly without asking the customer to repeat themselves.

10 **Swarm / Multi-Agent System** **Enterprise / Research**

Many specialized agents. One coordinated outcome.

Multiple agents with distinct roles cooperate or compete to solve complex problems. A researcher agent gathers data. A writer agent drafts. A reviewer agent checks quality. An orchestrator coordinates the whole sequence. Powerful for tasks too complex for a single agent — but adds coordination complexity that requires careful architecture before deployment.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
CrewAI, AutoGen (Microsoft), Amazon Bedrock Multi-Agent	Amazon Bedrock Multi-Agent allows enterprises to deploy fleets of specialized agents that hand tasks to each other — one searches product catalogs, one checks inventory, one processes the order — coordinated by an orchestrator that manages the full customer interaction.

Powerful for complex research or simulation. High coordination risk for small business — deploy tightly scoped or not at all.

11 **Control-System Agent** **Enterprise / Industrial**

Inputs → constraints → decision → action → feedback → adjustment.

The control-system agent is modeled on industrial process control. It continuously reads inputs, applies constraints, makes a decision, takes an action, measures the result, and feeds that result back as the next input. It never stops. It never drifts. It self-corrects. Less chatbot, more operating system. The strongest business model for autonomous operation.

MAJOR PLAYERS	REAL-WORLD EXAMPLE
Siemens (industrial), Honeywell, Google DeepMind, Tesla Autopilot, n8n, LangGraph	Google's DeepMind AI manages data center cooling. It continuously reads temperature sensors, adjusts cooling systems, measures energy consumption, and recalibrates — reducing cooling energy use by 40% with no human in the loop. The loop never opens.

For small business: n8n or LangGraph with a feedback loop built in. Output from step 4 feeds back into step 1. This is the architecture that makes autonomous operation real.

ARCHITECT'S RECOMMENDATION

Don't Build the Pyramid. Build the Right Stack.

The eleven agent types above are not a checklist. They are a vocabulary. Most small businesses need three or four agent types working together cleanly — not an enterprise-grade swarm of eleven.

The most common mistake: deploying the most sophisticated agent type available when a simpler one solves the problem with less risk. A control-system agent running your pricing model is overkill if a policy-based optimizer handles the same decisions with half the complexity.

USE CASE	RECOMMENDED STACK
Dynamic pricing	Event-Driven + Policy-Based + Optimizer + Human Approval
Lead qualification	Event-Driven + Utility Agent + Memory Agent
Client onboarding	Event-Driven + Planned + Policy-Based
Content production	Planned + Swarm (scoped) + Human-in-the-Loop
Inbound call handling	Reactive + Policy-Based + Human-in-the-Loop
Autonomous reporting	Memory + Planned + Control-System

Knowing the map is step one. Knowing where your business sits on it — and which agents you can deploy this week versus which ones require infrastructure you haven't built yet — is step two.

NEXT STEP

Is Your Business AI-Ready?

Reading this guide gives you the map. The AI Surface Audit tells you exactly where your business sits on it.

Most small business owners have never had an enterprise architect look at their operation. They have vendors who sell software and consultants who sell more software. They rarely have someone who maps the entire system — the tools, the data, the workflows, the security posture — and gives them a clear, prioritized 90-day plan built on 20+ years of enterprise solutions architecture experience.

That is exactly what the AI Surface Audit delivers. And it starts at $250.

TIER	WHAT YOU GET	INVESTMENT
Triage	AI Score, Stack Snapshot, Top 3 Priorities, 30-Day Quick Wins	From $250
Standard	Full 10-Tab Audit, Workflow Atlas, Security Assessment, 90-Day Roadmap, ROI Model	Contact for Pricing
Transformation	Full Audit + Implementation Planning + Agent Architecture + Ongoing Advisory	Contact for Pricing

Delivered in approximately 7 days. Enterprise rigor. Small business price.

Start your AI Surface Audit today.

BrotherILA.com

brandt.delany@gmail.com • Text 346.589.9888

linkedin.com/in/ila-brandt-delany-a43a613