BairesDev Areas of Expertise | Anthropic

Where our Claude practice is deepest, and where we would recommend to leverage BairesDev

Claude in 9 of our 13 most recent AI case studies · Claude Code, API and Bedrock all in production

	WHERE WE SUPPORT	WHERE WE'LL POINT YOU ELSEWHERE
Agentic Software Delivery	Claude Code agent fleets — coordinator and workers, one live agent per ticket — taking config-only integration tickets from days to QA-ready in hours, with quality gates cleared before QA and lessons promoted into shared memory.	Autonomous merge with no human reviewer. An engineer reviews, owns and merges every change we ship.
Context & Cost Engineering	Distilled memory banks that cut tokens per repo-reasoning task by ~70%, one agent scoped per directory instead of per monorepo, and MCP servers exposing only a whitelisted tool set per agent.	Model training or fine-tuning. We shape context, retrieval and tooling around Claude; we don't compete with frontier labs.
Agentic Ops & SRE	Investigation harnesses that pull data across systems, run the checks, update the ticket and hand back a structured incident summary — a full day reduced to well under an hour.	Managed 24x7 watch floors and IR retainers. We build and harden the harness; MSSPs run the monitoring.
Knowledge & Document Agents	Agents over Confluence, Jira and Slack that choose which tools to call and derive real semantics rather than returning raw text, plus document extraction and enterprise retrieval workflows.	Certified ITSM and knowledge-suite implementation — vendor PS or a specialized SI.
Customer Operations at Volume	Claude classifying, routing, resolving and drafting responses across 200,000+ support tickets a month in production, with automations handling the downstream steps and cost per ticket tracked.	BPO and contact-center staffing. We build the automation; the operation stays with the customer or their BPO.
Adoption, Enablement & Governance	Claude Code and Claude Max rollouts, engineer enablement and certification tracks, and review and approval gates for MCP servers, skills and plugins before they reach a repository.	Seat procurement and licensing negotiation — that stays between Anthropic and the customer.
How We Engage	Defined Claude workstreams with scoped deliverables, and senior engineers embedded in a customer's team delivering with Claude on the customer's own entitlement.	Deals where Claude consumption sits on our entitlement rather than the customer's. We confirm that first and tell you either way before registering.

BairesDev

Agentic AI Success Stories | Anthropic

Nine recent engagements shipping with Claude — the outcome and the surface it runs on

An engineer reviews and owns every change · prototype to production in weeks, not quarters · client names withheld

THE WORKLOAD	WHAT CHANGED	CLAUDE SURFACE & STACK
Agentic ticket delivery Freight brokerage	Config-only integration tickets went from days to QA-ready in hours using a coordinator-and-worker fleet of Claude Code agents, one live agent per Jira ticket — ~10x faster in the initial pilot.	**Claude Code** · Cursor · GitHub Spec Kit
Context & token engineering Collectibles grading	Tokens per repo-reasoning task cut ~70% (45K to 13K) with distilled memory banks; large-repo task completion rose from ~40% to ~92%; time to first triage insight 35 to 6 minutes.	**Claude** · **MCP** · .NET · Kubernetes · Datadog
AI help desk at volume Drive-through coffee chain	200,000+ tickets a month now classified, routed and resolved autonomously. Resolution 4 hours to 39 minutes, cost per ticket $0.72 to $0.58, roughly $258K in annual savings.	**Claude on Amazon Bedrock**
Agent-assisted SRE Enterprise no-code platform	Customer investigations cut from a full day to 30 minutes–2 hours across 5 systems and 90+ clusters, a 40–50% throughput increase, with the engineer approving every infrastructure action.	**Claude** · AWS · Kubernetes · Terraform · Helm
Release intelligence Satellite and online radio	A 10-minute Confluence and Jira lookup became a Slack question answered in about a minute — roughly 90% faster — with real release semantics. Prototype to production in ~3 weeks.	**Claude on Amazon Bedrock** · ECS Fargate · Terraform
Fraud investigation Residential internet and TV	A routine low-priority ticket became a confirmed 2-year phone-spoofing fraud in ~4 working days: 2 years of call history correlated across 8 systems, hundreds of recordings classified. NPS 10/10.	**Claude**
Investigation data platform Autonomous vehicles	15–30 analyst hours freed per week, with 4 disconnected sources unified into one pipeline carrying anomaly detection and ML root-cause labeling. Proposal to production in ~6 weeks.	**Claude** · OpenAI · Databricks · PySpark
Agentic QA authoring Legal case management	An agent that drafts test scenarios and locates existing cases cut hundreds of hours of non-testing QA work per cycle down to double digits, and halved the time to find an existing test case.	**Claude** · GitHub Copilot · .NET · Azure
Test automation from zero Retail e-commerce	Zero to 95%+ automation coverage of critical journeys in about 3 months across 3 storefronts, cutting manual testing time ~60% and manual reporting effort ~80%.	**Claude** · OpenAI · Playwright · CI/CD

BairesDev