

SOFTBLUES · BUYER'S GUIDE · MAY 2026

Claude Implementation Buyer's Guide.

Platform choice, governance, costs, and vendor selection for UK and Ireland businesses.

Anthropic Partner Network applicant

Google Cloud Partner since 2014

London

16 slides

Six chapters. Sixteen slides. Forty minutes.

Written for the people who actually have to choose a platform, budget the work, and sign the contract.

1. The platform choice

When Claude is the right answer, when Copilot is, when Gemini is, and how to weigh the soft factors that rarely make it into RFPs.

2. Governance and compliance

SSO, SCIM, retention, audit logs. Mapped against UK GDPR, ICO AI guidance, and FCA conduct rules.

3. Total cost of ownership

Licence costs, API costs, integration costs, internal time, ongoing partnership. Sample budgets at 100, 500, and 1,500 staff.

4. Realistic rollout timing

Where most rollouts slip. The 90-day model. Department sequencing. How to avoid the common stalls.

5. Build versus buy

Custom MCP servers, Vertex AI agents, custom Skills. When custom is worth the cost, when it is not.

6. Questions for any vendor

The 12 questions that surface real expertise versus PowerPoint expertise. Reference checks, work samples, contract terms.

FOREWORD

Most vendor guides are **thinly disguised sales collateral.**

This one is written differently.

Most AI vendor guides overstate the case for the vendor that published them. They avoid honest comparisons. They dodge specific cost information. The buyer reads them and ends up no closer to a decision.

This guide is written differently by design. Where Claude is the wrong answer for your situation, the guide says so. Where the costs are higher than the marketing suggests, the guide gives you real numbers. Where vendors typically obscure deployment risk, the guide names the failure modes we have seen.

The point is to be useful. If after reading this guide you decide Claude is not for you, that is a win for us too. We did not waste your time, and you have a clearer view of your options.

Written by Ivan Pylypchuk and Niall Marshall. Reviewed every 90 days. Last revision May 2026.

Claude, Copilot, or Gemini? Honest answer.

The headline strengths of each, and the buyer types that fit each one best.

Claude Enterprise

Strongest at long-form reasoning, complex analysis, document work, and coding. Best enterprise security controls (SCIM, custom retention, audit logging). Anthropic's safety-first model fits regulated industries. Works across AWS, Google Cloud, and Azure.

Best for: knowledge-intensive teams, regulated industries, engineering organisations.

Microsoft Copilot

Tightly woven into Microsoft 365 and Teams. Best if your daily work happens in Word, Excel, PowerPoint, Outlook, and SharePoint. Heavy reliance on M365 metadata for context. Less flexible outside the Microsoft stack.

Best for: M365-native companies whose AI use case sits inside Office workflows.

Gemini Enterprise

Tightly coupled to Google Workspace. Strong for internal document search across Drive and Gmail. Larger context window. Google ecosystem integration is the differentiator.

Best for: Google Workspace shops that want everything under one Google umbrella.

Three signals that Claude fits.

If two or three of these are true for your company, Claude is the most likely first call.

Knowledge-intensive work

Your team reads contracts, drafts reports, summarises research, reviews documents at scale. Claude consistently outperforms on long-form reasoning and complex analysis. The productivity gain is largest where the work is most cognitive.

Governance and compliance matter

You operate under UK GDPR, FCA conduct rules, SRA standards, or an internal compliance framework that takes audit seriously. Anthropic's enterprise security controls (SSO, SCIM, audit, custom retention) and safety-first model fit the posture regulated buyers need.

Multi-cloud or cloud-agnostic stack

You run on AWS, Google Cloud, Azure, or a mix. You do not want to be locked into one vendor's productivity ecosystem. Claude Enterprise is available across all three major clouds and does not require you to be a Microsoft 365 or Google Workspace house.

Three use cases where another tool fits better.

Honest about platform fit. Where the work sits clearly outside Claude Enterprise's strengths, we say so.

Heavy M365 use with happy Copilot adoption

If your team lives in Outlook, Word, Excel, Teams, and SharePoint, and is already getting real value from Copilot, switching introduces friction with limited upside. Copilot's M365 metadata access is genuinely useful. Stay with what works.

Image generation as the primary need

Claude does not generate images. If your dominant AI use case is creative imagery, marketing assets, or visual content production, Gemini, OpenAI's tools, or specialist image platforms are the right pick. Claude can sit alongside them, not replace them.

Public-facing customer support chatbot

Claude Enterprise is built for internal teams using Claude across their workflow. If your primary need is a customer-facing chatbot or voice agent on your website or phone line, our separate AI Voice Agents programme is the right offer. Different architecture, different deployment.

What governed AI actually means.

Four pillars your IT and compliance teams will want to verify before any rollout.

Identity

Single sign-on through your identity provider (Okta, Microsoft Entra, Google Workspace). SCIM for automated user provisioning and de-provisioning. Role-based access at platform, project, and tool level.

Data retention

Granular control over how long prompts, outputs, and conversation history are stored. From immediate deletion to indefinite retention, per the policy your compliance team sets. Anthropic does not train on enterprise data by default.

Audit logging

Every prompt and response logged with user, timestamp, model, and tool calls. Filterable by user, department, document, or date. The audit trail your regulators and internal compliance officers actually want.

Action governance

Connectors and agents need permission scoping. Who can invoke which tool. What data flows where. DLP gateways and policy enforcement before content leaves your perimeter.

Certifications worth verifying with Anthropic directly: SOC 2 Type II, SOC 3, ISO 27001, ISO 42001, HIPAA-ready under BAA. Available across AWS, Google Cloud, and Azure.

How Claude Enterprise lines up with UK and Irish regulation.

Not a substitute for your legal review, but a starting point for your compliance officer.

UK GDPR and ICO

Anthropic operates as a processor under UK GDPR. Standard Contractual Clauses are available for EU and UK data transfers. ICO AI guidance is well-served by Claude Enterprise's audit logging, retention controls, and Data Protection Impact Assessment material that Anthropic publishes.

FCA conduct rules

For FCA-regulated firms, audit trails on AI-assisted client work satisfy record-keeping obligations. Role-based controls let you enforce Conduct Rule Two (acting with due care) and SMCR accountability. Custom compliance Skills can encode firm-specific suitability checks.

EU AI Act and Irish regulation

EU AI Act high-risk classifications mostly fall outside Claude Enterprise's general-purpose use, but the act's transparency and governance requirements map cleanly to the controls described on the previous slide. For Irish firms, the same regime applies in full.

Four cost components. Most buyers only see two.

Plan for all four upfront. The internal time and partnership costs are where budgets quietly slip.

Licence costs

Anthropic charges per seat for Claude Enterprise, typically discounted at volume. Team plan starts around \$25 per seat per month. Enterprise pricing is negotiated based on seat count and term length.

API costs

Custom agents, MCP servers, and high-volume workflows hit the Anthropic API at usage rates. Budget separately. Prompt caching, model routing, and the Batch API typically cut this by 40 to 80 percent when designed well.

Integration costs

Custom MCP servers, identity provider configuration, connector authorisation, and audit dashboards. Most companies need 2 to 5 custom integrations in year one. Each is a discrete piece of engineering.

Internal time and partnership

Your team's time during rollout (often underestimated). Training time. Ongoing administration. Implementation partnership costs. Together these often exceed the platform licence in year one.

Indicative year-one budgets across three engagement shapes.

Real ranges from comparable engagements. Smaller engagements scale down. Larger engagements scale up.

Small team or focused department

Anthropic licences: 10 to 50 seats year one.

API and custom builds: modest, mostly Skills and 1 custom MCP server.

Implementation: shorter focused engagement, often a compact strategy phase plus targeted deployment.

Ongoing partnership: light retainer.

Order of magnitude: mid-five figures GBP for year one all-in.

Whole company, 500 staff

Anthropic licences: 150 to 300 seats year one.

API and custom builds: moderate. 3 to 5 MCP servers, 1 or 2 Vertex AI agents, cost optimisation pass.

Implementation: full three-phase programme, multi-department rollout.

Ongoing partnership: active retainer with regular custom build work.

Order of magnitude: low to mid-six figures GBP for year one.

Whole company, 1,500 staff

Anthropic licences: 500 plus seats, multi-year contracting.

API and custom builds: significant, including compliance tooling and multiple Vertex AI agents.

Implementation: phased rollout across operations, 9 to 12 months.

Ongoing partnership: substantial retainer with dedicated team.

Order of magnitude: mid-six figures GBP year one, larger year two.

Plan for 90 days. Then years of expansion.

The three phases that work in practice, with the failure modes we see most often called out.

Phase 1, Strategy

2 to 3 weeks. Discovery, departmental interviews, technical and compliance review, shadow AI audit. Deliverables: executive presentation, strategy document, 12-month migration map, training curriculum, build catalogue.

Common stall: skipping discovery and going straight to deployment without leadership alignment.

Phase 2, Deployment

6 to 10 weeks. Five parallel tracks: platform foundation (SSO, audit), Projects and Skills per department, native and custom connectors, role-based training, adoption tracking.

Common stall: treating deployment as IT setup rather than change management. The platform launches; the team does not adopt.

Phase 3, Continuous partnership

Ongoing. Custom MCP servers, Vertex AI agents, compliance tooling, cost optimisation, voice extensions, adoption engineering. Quarterly roadmap reviews.

Common stall: treating the partnership as ad-hoc consulting rather than continuous capability building.

When to use native, when to customise.

Three categories of capability, each with a clear decision rule.

Use native Anthropic capabilities when they exist

Anthropic ships new native connectors, Skills, and features every month. Slack, Microsoft 365, Google Workspace, Salesforce, HubSpot, Jira, Confluence, GitHub are all native. Use them. Native means Anthropic maintains it, secures it, and updates it with Claude's evolving capabilities.

Build custom MCP servers for non-native systems

Where Anthropic does not have a native connector and the integration is critical (your internal CRM, a sector-specific platform, a legacy database), a custom MCP server is the right pattern. Typical effort: 1 to 4 weeks per server. Most companies need 3 to 5 in year one.

Build Vertex AI agents for specialised workloads

Where you need fine-tuned models, multi-agent orchestration, or long-running stateful processes, Google Cloud Vertex AI is the right destination. Agents are exposed to Claude as MCP tools or via Google's A2A protocol. Typical effort: 3 to 8 weeks per agent.

Twelve questions for any AI implementation vendor. Questions 1 to 6.

Real expertise versus PowerPoint expertise. These six surface most of the difference.

1. Show me production work

Ask for live deployments you can reference. A live system answering real user queries beats a case study PDF every time. If they cannot point to one, ask why.

2. How do you use Claude internally?

Vendors who do not run on Claude themselves are theorising. Ask which Projects, which Skills, which Connectors. If the answer is vague, the deployment will be too.

3. How many Certified Claude Architects?

Anthropic's certification programme is real. Vendors with certified architects on staff have invested in the platform. Numbers under five for a vendor pitching this work is a yellow flag.

4. Walk me through one timeline

Phased work with named workstreams beats hand-waving about timelines. If they cannot walk through a real engagement's weekly cadence, they have not run one.

5. What governance do you set up by default?

Good vendors have a default governance posture: SSO patterns, retention recommendations, audit configurations. Bad vendors say "we configure to your needs" without offering opinion.

6. Show me a custom MCP server you built

Asking for the source code is unreasonable. Asking for an architecture diagram, an explanation of auth, retries, and audit, is reasonable. Vagueness here means inexperience.

Questions 7 to 12. Cost, risk, and what comes after.

Half of these test how the vendor handles things going wrong. That is when you find out who they really are.

7. Tell me about a deployment that did not go to plan

Every vendor has had one. The honest answer covers what happened, what they learned, what they changed. Vendors who claim a perfect track record are either lying or new.

8. How do you price?

Fixed-price phases, day-rate consulting, or hybrid. Each has tradeoffs. What you want to avoid is open-ended scope at day rates. Ask for the maximum exposure on a typical engagement.

9. What happens when Anthropic obsoletes your work?

Anthropic ships fast. Good vendors expect this and structure engagements to migrate gracefully. Bad vendors resist replacement because they are protecting custom work.

10. Can I keep the strategy if I do not continue?

The Phase 1 strategy should be yours regardless. If the vendor makes it conditional or proprietary, that is a control move, not a service. Walk away from those.

11. Who is on my account day to day?

Names of the actual people. Real engineers and architects, not just an account manager. The person who designs the system should be reachable directly during the engagement.

12. What do you not do?

The most useful question. Vendors who answer "we can do anything" are not partners. Honest vendors have a clear scope and will refer you elsewhere outside it.

If you would like to continue the conversation.

Three ways to engage with Softblues. None of them commit you to anything.

1. Score your readiness

Take the Claude Readiness Score, a 10-question diagnostic that scores your company across people, process, technology, and culture. Five minutes, free, personalised PDF report sent by email.

softblues.io/claude-enterprise/readiness-score

2. Book a 20-minute discovery call

A no-pitch conversation with Ivan (UK) or Niall (Ireland) about your situation, your stack, your priorities, and whether Claude makes sense as the starting point. Honest answer either way.

softblues.io/contact

3. Read more

The Softblues Claude Enterprise hub covers the three-phase programme, the operations we work in, the connectors we deploy, and the case studies we can share publicly.

softblues.io/claude-enterprise

Softblues at a glance: London-based AI consultancy. 50 plus engineers and specialists. Google Cloud Partner since 2014. In the final stage of joining the Anthropic Claude Partner Network with 10 Certified Claude Architects on the team. Our company runs on Claude every day across sales, marketing, engineering, finance, and operations.

Thank you.

Any questions on this guide, or want to talk about what it would look like for your company,
just say. We read every reply.

Ivan Pylypchuk, CEO

ivan@softblues.io · softblues.io