

THE ENTERPRISE FIELD GUIDE

The Enterprise GenAI Field Guide

The reference manual for AI architects and team leads. Twenty-four detailed use cases across six industries, a full model decision matrix, the deployment checklist and RAG blueprint that make systems production-ready, ready-to-adapt governance and security policy templates, and a 90-day enterprise implementation roadmap.

WHAT'S INSIDE

01 · 24 USE CASES

Four detailed use cases across each of six industries.

02 · DECISION MATRIX

ChatGPT vs Claude vs Gemini, full comparison.

03 · DEPLOYMENT

Readiness checklist + an enterprise RAG blueprint.

04 · POLICY TEMPLATES

Governance & security policies to adapt.

05 · 90-DAY ROADMAP

From discovery to operationalised at scale.

06 · TEAM & SKILLS

The roles and capability you'll need.

6

INDUSTRIES

24

USE CASES

5

POLICY TEMPLATES

90

DAY ROADMAP

A companion to GSDC's generative-AI use-case resources. Tool and model capabilities reflect early-2026 positioning and change quickly — verify current specifics with each vendor and on the certification page. Policy templates are starting points, not legal advice.

START HERE

Why most GenAI pilots never reach production

Enterprises rarely fail at the demo. They stall on the gap between a convincing prototype and a system that's secure, governed, reliable and cost-controlled at scale. This guide is built around closing that gap.

TRAP 1**Use case without value**

Building because the tech is exciting, not because it solves a costed business problem. Start from value and risk, not novelty.

TRAP 2**Demo-grade retrieval**

Naive RAG impresses in a demo and disappoints in production. Reliable retrieval is an engineering discipline, not a quick win.

TRAP 3**Governance as an afterthought**

Security, privacy and policy bolted on late become the blocker. Bake them in from the first pilot.

TRAP 4**No evaluation**

Without measurement you can't tell improvement from regression — or prove the system is safe to scale.

HOW TO USE THIS GUIDE

Find use cases for your industry (pages 3–8). Choose your models (9–10). Apply the deployment checklist and RAG blueprint (11–12). Adapt the policy and security templates (13–15). Then run the 90-day roadmap (16–17) with the right team (18).

01 · USE CASES

1

Financial Services & Banking

High value, high regulation — grounding and human review are non-negotiable.

USE CASE 1**Research & filing summarisation**

Condense filings, earnings calls and analyst notes into briefs with citations.

RAG + Summarisation

USE CASE 2**Customer-service copilot**

Draft agent replies grounded in current product and policy documents.

RAG copilot

USE CASE 3**Risk & fraud narratives**

Turn alerts and transaction data into readable case narratives for analysts.

Summarisation

USE CASE 4**Compliance drafting**

Draft disclosures and regulatory responses from approved source policy.

RAG + Templating

SECTOR CAUTION

No automated decisions on credit, fraud or advice without human review. Mask sensitive data before it reaches a model, cite every figure, and keep a full audit trail for regulators.

01 · USE CASES

2 Healthcare & Life Sciences

Powerful for documentation and synthesis — never for unsupervised diagnosis.

USE CASE 5

Clinical note summarisation

Summarise visit notes and discharge summaries for clinician review.

Summarisation

USE CASE 6

Patient information assistant

Answer admin and general health-info questions, with strict scope guardrails.

RAG

USE CASE 7

Literature synthesis

Synthesise research papers into reviews with traceable citations.

RAG + Summarisation

USE CASE 8

Prior-authorisation drafting

Draft authorisation and claims letters from records for staff review.

RAG + Templating

SECTOR CAUTION

Protected health information requires strict handling and masking. Clinician verification is mandatory; assistants inform, they do not diagnose. Build escalation paths into every patient-facing flow.

50% OFF

Turn use cases into production systems

The GSDC Generative AI Certification teaches the patterns behind every use case here — hands-on, in the AI Studio. The natural next step from this reference. Enrol now with 50% off.

[Explore the certification →](#)

Recognised in 90+ countries

01 · USE CASES

3 Retail & E-commerce

Scale and personalisation — watch for fabricated product details.

USE CASE 9

Product content generation

Generate descriptions, attributes and SEO copy at catalogue scale.

Generation

USE CASE 10

Conversational search

Natural-language product discovery and recommendations.

RAG + Retrieval

USE CASE 11

Support automation

Resolve common queries; route the rest to humans with context.

RAG agent

USE CASE 12

Review & VoC summarisation

Turn thousands of reviews into themes and prioritised actions.

Summarisation + Classification

SECTOR CAUTION

Generated product claims need accuracy and brand-voice review before publishing. Conversational search must retrieve real specifications — never invent them. Guard any tool access to order and customer data.

01 · USE CASES

4 Manufacturing & Supply Chain

Knowledge at the point of need — keep sources version-correct.

USE CASE 13

Maintenance knowledge assistant

Q&A over manuals and SOPs for technicians on the floor.

RAG

USE CASE 14

Quality & defect analysis

Summarise inspection and defect reports into trends.

Summarisation + Classification

USE CASE 15

Supplier contract review

Extract terms and flag risks across many supplier agreements.

RAG + Extraction

USE CASE 16

SOP & training generation

Draft work instructions and training from approved sources.

Generation

SECTOR CAUTION

Maintenance and safety content must reference the correct equipment version — wrong-version answers are a safety risk. Keep humans on root-cause analysis and require SME sign-off on generated procedures.

01 · USE CASES

5 Technology & Software

Productivity across the SDLC — with security and licence review.

USE CASE 17

Code assistance & review

Draft, explain and review code with a human in the loop.

Generation + Analysis

USE CASE 18

Internal docs assistant

RAG over wikis, runbooks and tickets, scoped by team access.

RAG

USE CASE 19

Incident & postmortem summaries

Turn logs and threads into clear incident summaries.

Summarisation

USE CASE 20

Test generation

Draft test cases from specifications for engineer validation.

Generation

SECTOR CAUTION

Review generated code for security flaws and licence issues before merging. Scope documentation retrieval by team permissions so the assistant can't surface content a user shouldn't see.

LIMITED TIME

Equip your team to ship these, not just demo them

For a limited time, the GSDC Generative AI Certification is open with a special offer — the fastest way to build production-grade GenAI capability across your team.

[Enrol while it lasts →](#)

Offer ends soon

01 · USE CASES

6 Professional Services & Legal

Knowledge-heavy work — grounding, citations and confidentiality first.

USE CASE 21

Contract review & redlining

Identify clauses, risks and deviations across documents.

RAG + Extraction

USE CASE 22

Proposal & RFP drafting

Draft responses from a curated knowledge base.

RAG + Templating

USE CASE 23

Research memo synthesis

Synthesise sources into structured, cited memos.

RAG + Summarisation

USE CASE 24

Knowledge-management assistant

Firm-wide Q&A over documents with confidentiality controls.

RAG

SECTOR CAUTION

Professional review remains accountable for every output. Preserve confidentiality and privilege, cite sources, and segregate client data so retrieval never crosses engagement boundaries.

02 · DECISION MATRIX

ChatGPT vs Claude vs Gemini

General positioning of the three leading enterprise assistants. Capabilities converge and change fast, so treat this as orientation and confirm current specifics with each vendor before committing.

| DIMENSION | CHATGPT · OPENAI | CLAUDE · ANTHROPIC | GEMINI · GOOGLE |
|---------------------------|-------------------------------|-----------------------------------|-------------------------------|
| General strength | Broad, versatile all-rounder | Careful reasoning, strong writing | Multimodal, search-grounded |
| Long documents | Strong | A noted strength | Strong |
| Multimodal | Yes | Yes | Yes (native focus) |
| Ecosystem | Largest third-party ecosystem | API & enterprise partners | Google Cloud & Workspace |
| Enterprise / admin | Enterprise tier available | Enterprise tier available | Via Google Cloud |
| Best-fit | Broad use, widest tooling | Document-heavy, careful drafting | Google-stack orgs, multimodal |

Positioning only — not an endorsement, and not specific to any version. Context limits, pricing, model names and features differ and update frequently; verify on each provider's official site. Open-weight models (e.g. Llama, Mistral) are an additional option where data control or self-hosting matters.

50% OFF TODAY

Learn to choose and combine models — at half price

The GSDC certification covers hands-on tool selection across the leading assistants on real challenges. Enrol today and claim a flat 50% saving on your seat.

Claim 50% off →

Applied at enrolment

02 · DECISION MATRIX

When to choose which — and why not to pick just one

REACH FOR

ChatGPT

Broad general-purpose work and the widest selection of third-party integrations and tooling.

REACH FOR

Claude

Long-document analysis, careful reasoning and high-quality drafting and editing.

REACH FOR

Gemini

Organisations already on Google Cloud and Workspace, and multimodal workloads.

THE ENTERPRISE PATTERN

Don't marry one provider

Mature teams route tasks to the best-fit model behind an internal gateway, so they can switch or mix providers without rewriting applications. This controls cost, avoids lock-in, and lets you adopt new models as they ship.

DECISION FACTOR

Data & deployment

If data residency or self-hosting is essential, weigh open-weight models alongside the hosted options.

DECISION FACTOR

Total cost & latency

Right-size the model to the task; not every job needs the largest, most expensive option.

03 · DEPLOYMENT

Enterprise deployment readiness checklist

Before any GenAI system goes to production, walk this checklist. A demo proves it can work once; these items prove it works reliably, securely and affordably.

DATA & SOURCES

- Sources inventoried, owned and refreshable
- Sensible chunking & metadata in place
- Retrieval scoped to user permissions

IDENTITY & ACCESS

- SSO and role-based access control
- Per-team data segregation enforced
- Tool/action permissions least-privilege

SECURITY & PRIVACY

- PII/PHI masking before model calls
- Prompt-injection & data-egress defences
- Vendor data-use terms reviewed (no training on your data)

COST & PERFORMANCE

- Token budgets & caching configured
- Right-sized models per task
- Fallback on provider outage

EVALUATION & MONITORING

- Fixed evaluation set & metrics
- Tracing & regression alerts
- User-feedback capture loop


GOVERNANCE & CHANGE

- Policies live (use, data, risk tiers)
- Audit logging & incident plan
- User training & rollout plan

03 · DEPLOYMENT

The enterprise RAG blueprint

The production pattern that lifts retrieval past the naive plateau and makes answers trustworthy and auditable.



ENTERPRISE ADDITIONS

Wrap the pipeline

Access-scoped retrieval, PII redaction, audit logging of every query and source, and guardrails on inputs and outputs — on top of the core flow.

DIAGNOSE FAILURES

It breaks in two places

Bad retrieval (wrong chunks → fix data & search) or **bad generation** (good chunks, bad answer → fix prompting & evaluation). Ask “were the right chunks retrieved?” first.

PRODUCTION BASELINE

Naive retrieval (embed, top-k, stuff the prompt) plateaus around 70–80% precision. Hybrid search, metadata filtering, reranking and citations are the baseline for systems you'd trust with customers.

48 HOURS ONLY

Build a production RAG pipeline — enrol within 48 hours

This blueprint is core to the GSDC certification's hands-on build track. Act within 48 hours to lock in the current enrolment offer.

[Secure my seat →](#)

Window closes in 48h

04 · POLICY TEMPLATES

Governance policy templates (1 of 2)

Starting points to adapt with your legal, security and compliance teams. Fill the fields for your organisation — these are scaffolding, not legal advice.

Template A · Acceptable Use Policy

Permitted uses [list approved tasks & use cases]

Prohibited uses [e.g. no unsupervised decisions, no confidential data in unapproved tools]

Approved tools [list sanctioned assistants & platforms]

Disclosure [when AI assistance must be disclosed]

Template B · Data Handling & Classification

Data tiers [public / internal / confidential / restricted]

Tier-to-tool rules [which data tiers may be used with which tools]

PII / PHI handling [masking, minimisation, consent requirements]

Retention & logging [how long prompts/outputs are kept, where]

04 · POLICY TEMPLATES

Governance policy templates (2 of 2)

Template C · Model & Vendor Governance

Approved vendors [list, with data-use terms reviewed]

Model inventory [models in use, owners, purpose]

Review cadence [how often vendors/models are re-assessed]

Contracts & DPAs [data-processing agreements in place]

Template D · Human-in-the-Loop & Risk Tiers

Risk tiering [low / medium / high by impact]

Required review [human sign-off rules per tier]

Prohibited automation [decisions never made without a human]

Escalation [when & how the system hands off to a person]

RISK-FREE · 50% OFF

Build real governance capability — risk-free

The GSDC certification's responsible-AI module turns templates like these into practiced skill, backed by a 7-day money-back guarantee. Enrol with nothing to lose and 50% off.

Start risk-free →

7-day money-back guarantee

04 · SECURITY

Security checklist & incident template

Generative AI adds new attack surfaces — prompt injection, data leakage through outputs, over-permissive tools. Treat it as a security-critical system.

SECURITY CHECKLIST

- Least-privilege access to data & tools
- Prompt-injection defences & input validation
- Output filtering for data leakage
- Secrets never placed in prompts
- Full audit logging of prompts & outputs
- Red-teaming before launch & periodically

Template E · AI Incident Response

| | |
|--------------------|--|
| What counts | [harmful output, leak, misuse, outage] |
| Notify | [owners, security, legal, comms] |
| Contain | [disable feature/tool, revoke access] |
| Review | [root cause, fix, policy update] |

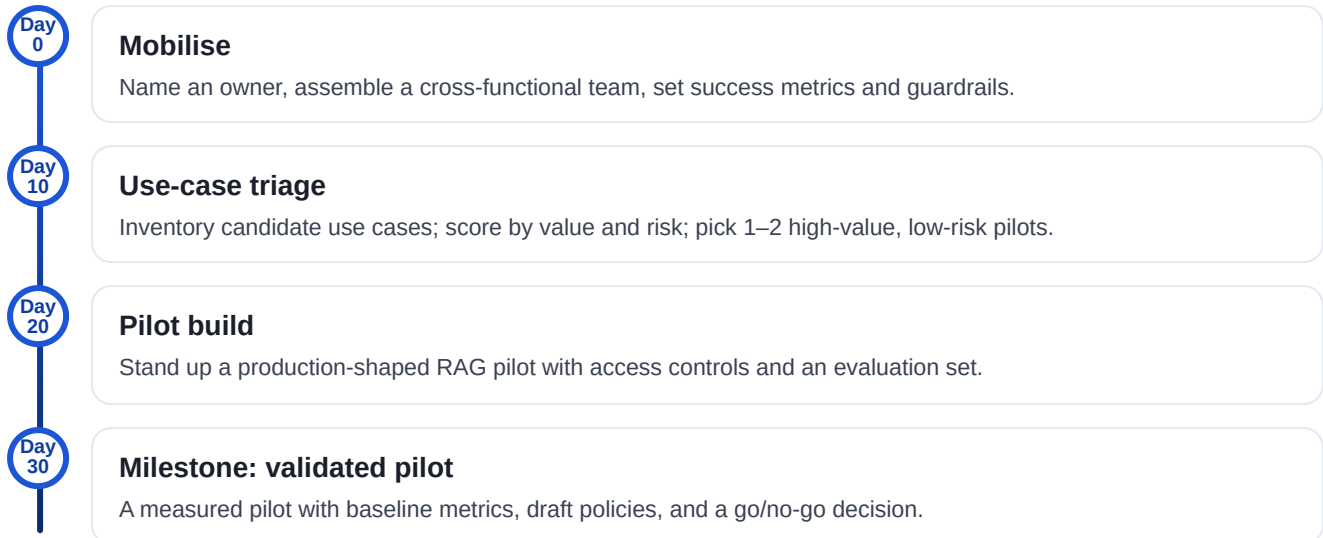
MINDSET

Assume any text reaching the model could be adversarial, and any output could be seen by the wrong person unless controlled. Design for both from day one.

05 · IMPLEMENTATION ROADMAP

The 90-day implementation roadmap

From scattered experiments to a governed, scaling capability. Phase 1 proves value safely on a focused pilot.



PHASE 1 PRINCIPLE

Start narrow and safe. One genuinely useful, well-governed pilot builds more momentum than five half-built experiments.

OFFER INSIDE

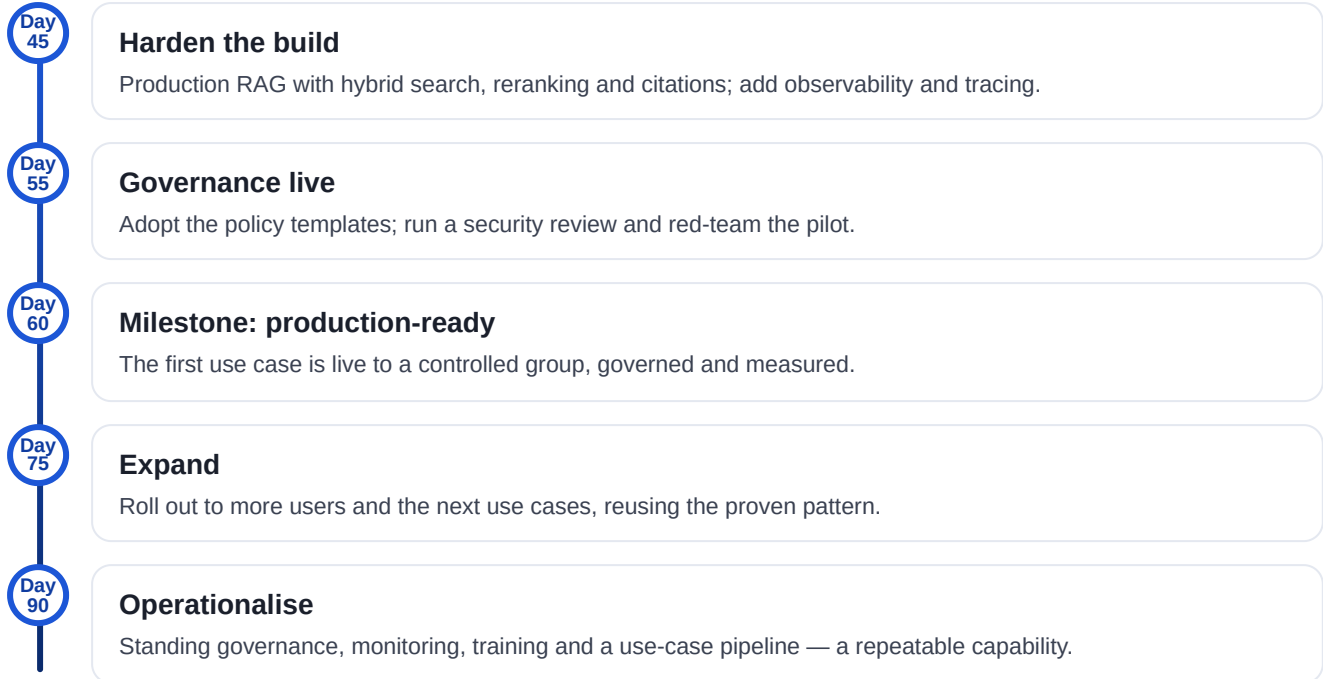
Build the team behind the roadmap — credential recognised in 90+ countries

A skilled, certified team is what carries a pilot to production. The GSDC certification builds that capability and proves it. A special enrolment offer is on the page.

Get certified →

Globally recognised credential

05 · IMPLEMENTATION ROADMAP

Days 31–90 · build, govern & scale**90-DAY OUTCOME**

One use case in production, a reusable deployment pattern, live governance, and a team that can repeat the process — not a one-off demo.

06 · TEAM & SKILLS

The team behind production GenAI

Enterprise GenAI is a team sport. These are the roles a deployment needs — and the capability the certification builds across them.

| ROLE | OWNS | KEY SKILLS (SYLLABUS) |
|---------------------|---------------------------------|-----------------------|
| AI Architect | The end-to-end stack & patterns | M2, M5, M6, M7 |
| GenAI Engineer | Building the pipelines & agents | M3, M5, M6 |
| Product / Team Lead | Use-case value & adoption | M1, M4, M9 |
| Governance Lead | Policy, risk & compliance | M7, M8 |
| Domain SME | Accuracy & sign-off | M1, M5 |

WHY CERTIFY THE TEAM

A shared certification gives everyone a common vocabulary and a verified baseline of skill — so architects, engineers, leads and governance speak the same language and move faster together. The nine-module syllabus maps directly to the roles above.

HALF PRICE

Build production-ready capability across your team — at half price

The GSDC certification gives your people the hands-on skills and shared language to ship enterprise GenAI. Enrol now at 50% off and start building the team.

[Start the pathway →](#)

Hands-on from day one

RECAP

Enterprise readiness on one page

Your quick-reference before the next steering meeting.

PICK USE CASES BY

Business value × manageable risk. Start with one or two; reuse the pattern.

CHOOSE MODELS BY

Task fit, data needs, cost — and route across providers behind a gateway.

DEPLOY WITH

Production RAG, access control, evaluation, observability — the full checklist.

GOVERN WITH

Use, data, vendor, risk-tier and incident policies — live from the pilot.

Hybrid search

Vector + keyword retrieval; the production baseline.

Reranking

A second pass ordering chunks by true relevance.

Grounding

Forcing answers from retrieved context, with citations.

Risk tiering

Matching human-review requirements to impact.

THE ONE-LINE TAKEAWAY

Enterprise GenAI succeeds when value, reliable retrieval, evaluation and governance are designed in from the first pilot — and carried by a skilled team.

WHERE TO GO NEXT

From pilot to production capability

You have the use cases, the model guidance, the deployment blueprint, the policy templates and a 90-day roadmap. The differentiator now is a team that can execute — and prove it.

CHOOSE

A high-value, low-risk first use case.

BUILD & GOVERN

With the blueprint, checklist and policies.

SKILL THE TEAM

With a shared, verified certification.

The bottom line

Most organisations already know *what* they want generative AI to do. The advantage goes to those who can deploy it reliably, securely and responsibly — at scale. This guide maps that path; a skilled, certified team walks it.

LAST CHANCE · 50% OFF

Make enterprise GenAI production-ready — at 50% off

You have the reference; the last step is the capability to execute it. Equip your team with the GSDC certification now at half price — your last prompt in this guide, so make it count.

Enrol & get certified →

Offer applied at checkout

© 2026 Global Skill Development Council (GSDC). Educational companion to GSDC's generative-AI use-case resources. Model and tool positioning reflects early-2026 information and changes quickly; verify with each vendor. Policy templates are starting points, not legal advice — adapt them with your legal, security and compliance teams. Verify current program and pricing details on the official certification page.