GSDC
Global Skill Development Council

PATTERNS REFERENCE · PRINTABLE CHEAT-SHEET

The Agentic AI Patterns Reference

A printable cheat-sheet of all six core patterns — one-page diagrams, when to use each, and the common pitfalls to avoid.

INSIDE THIS REFERENCE

- One-page diagram for each of the six patterns
- When to use each — and common pitfalls
- A pattern-selection decision guide
- Governance & agentic AI security best-practices checklist
- Career roadmap, governance roles & salary guide
- 2026 hiring trends in agentic AI

Global Skill Development Council · Vendor-neutral, globally recognised certification

How to use this reference

Most agentic systems are combinations of a few well-understood patterns. Learn these six, know when each applies, and you can design almost anything — then compose them.

Built for

- Engineers designing agentic systems
- Architects standardising patterns across teams
- Anyone learning agentic AI who wants a quick reference
- Interview and whiteboard prep

How to read each card

- The diagram shows the shape of the pattern
- "What it is" gives the one-line idea
- "When to use" tells you where it fits
- "Common pitfalls" is what to avoid

Print pages 4–9 as a standalone cheat-sheet, or use the decision guide on page 10 to pick a pattern fast. Start the smallest pattern that solves the problem, then compose upward.

The six patterns at a glance

Demand in 2026 is for engineers who can design systems of agents, not single prototypes — and these are the building blocks.

1 · ReAct
Reason and act in a loop with tool use.

2 · Tool-Use / Router
Pick the right tool for each request.

3 · Orchestration
Supervisor delegates to worker agents.

4 · Reflection / Critic
Generate, critique, improve.

5 · Memory-Augmented
Read/write short-term & vector memory.

6 · Autonomous Loop
Perceive, plan, act, reflect.

The next six pages give each pattern its own one-page card with a diagram.

Pattern 1 · ReAct (Reason + Act)

The agent interleaves reasoning with tool actions, observing results and iterating until it can answer.

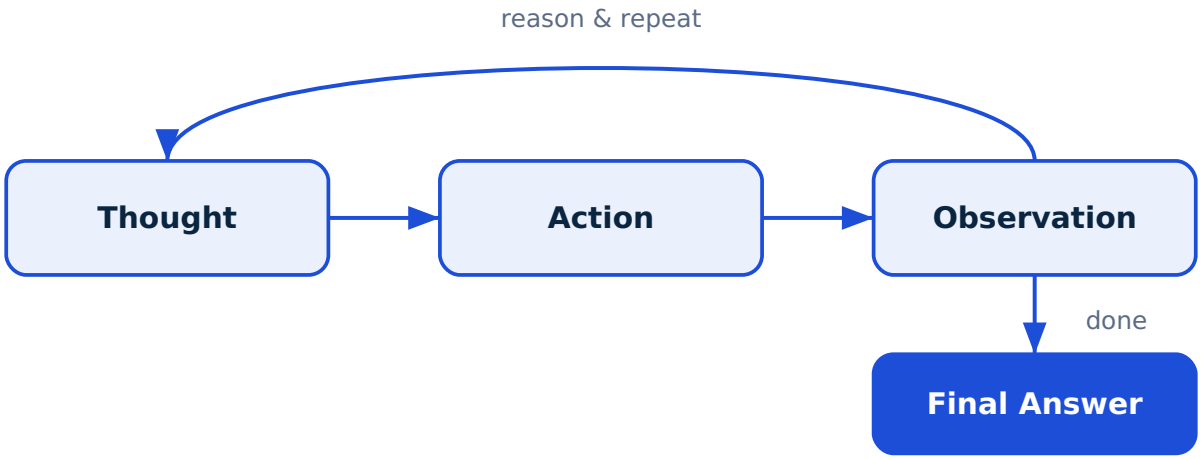

When to use

- Tasks needing tool use + step-by-step reasoning
- Question answering over external data
- When you want transparent, traceable steps

Common pitfalls

- Infinite loops with no stop condition
- Letting the agent act before it reasons
- No cap on iterations or token budget

GSDC
Global Skill Development Council

Pattern 2 · Tool-Use / Router

A routing agent picks the right tool or API for each request and integrates the result.

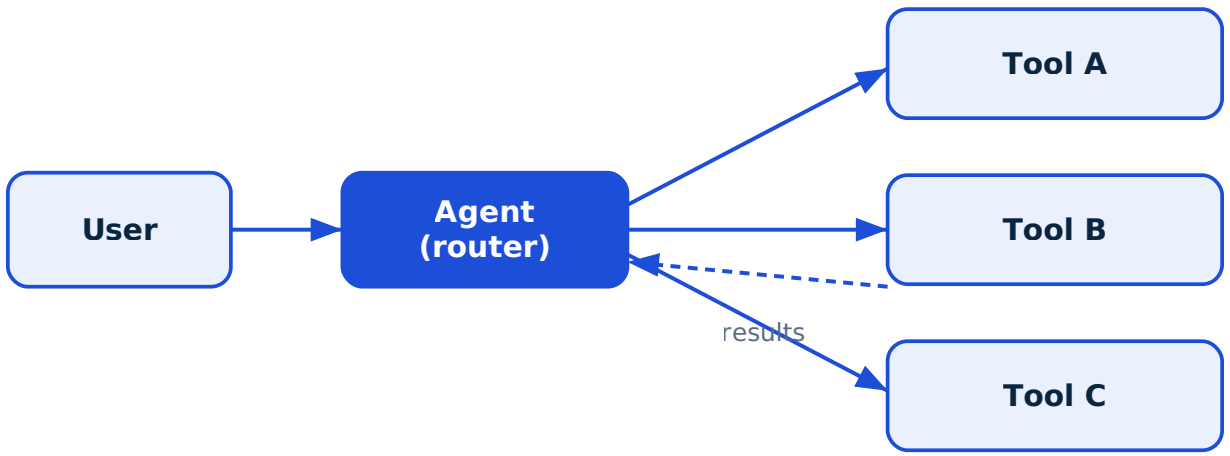

When to use

- Many specialised tools or data sources
- Clear, separable capabilities
- When intent maps cleanly to a tool

Common pitfalls

- Overlapping tools that confuse routing
- Vague tool descriptions
- No fallback when no tool fits

Pattern 3 · Multi-Agent Orchestration

A supervisor decomposes work and delegates to specialised worker agents, then aggregates results.

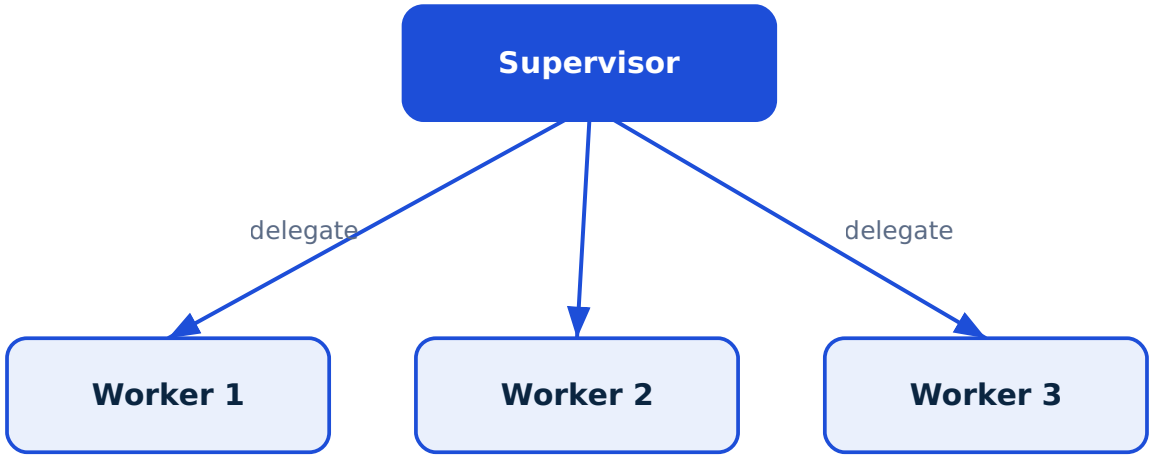

When to use

- Complex tasks with distinct sub-skills
- Parallelisable workloads
- When one agent's context gets overloaded

Common pitfalls

- Over-engineering simple tasks
- Unclear ownership / duplicated work
- No conflict-resolution protocol

Pattern 4 · Reflection / Critic

A generator produces output; a critic reviews and feeds back until quality criteria are met.

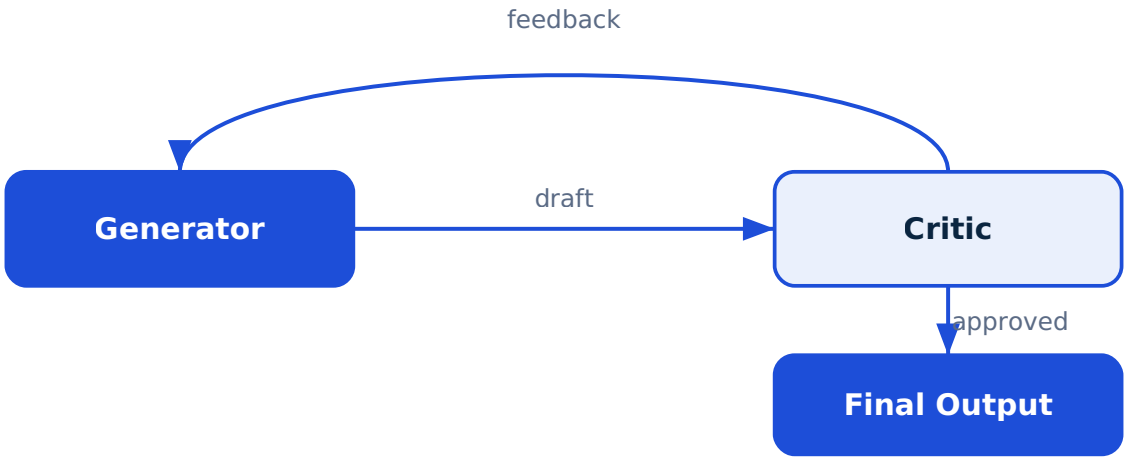

When to use

- Quality-critical outputs (code, analysis)
- Tasks with checkable criteria
- When first drafts are often wrong

Common pitfalls

- Critic with no clear rubric
- Endless revision without a stop rule
- Cost blow-up from too many rounds

Pattern 5 · Memory-Augmented Agent

The agent reads and writes short-term, episodic and semantic (vector) memory to stay coherent over time.

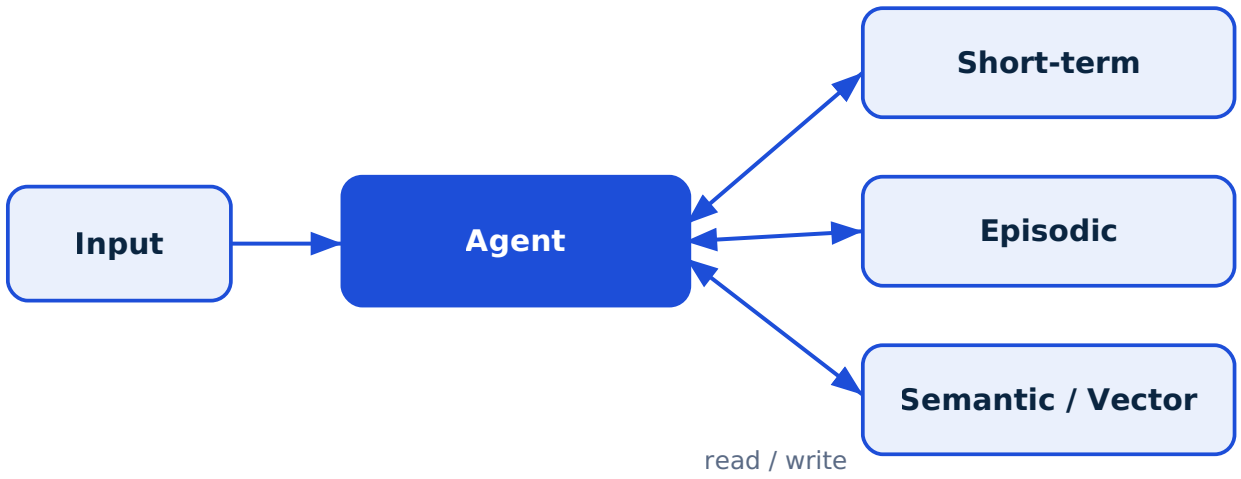

When to use

- Long or multi-session interactions
- Personalisation and recall
- When context windows aren't enough

Common pitfalls

- Storing everything (noisy recall)
- No retention / eviction policy
- Leaking PII into long-term stores

Pattern 6 · Autonomous Plan-Act Loop

A continuous perceive-plan-act-reflect cycle lets the agent operate with minimal supervision.

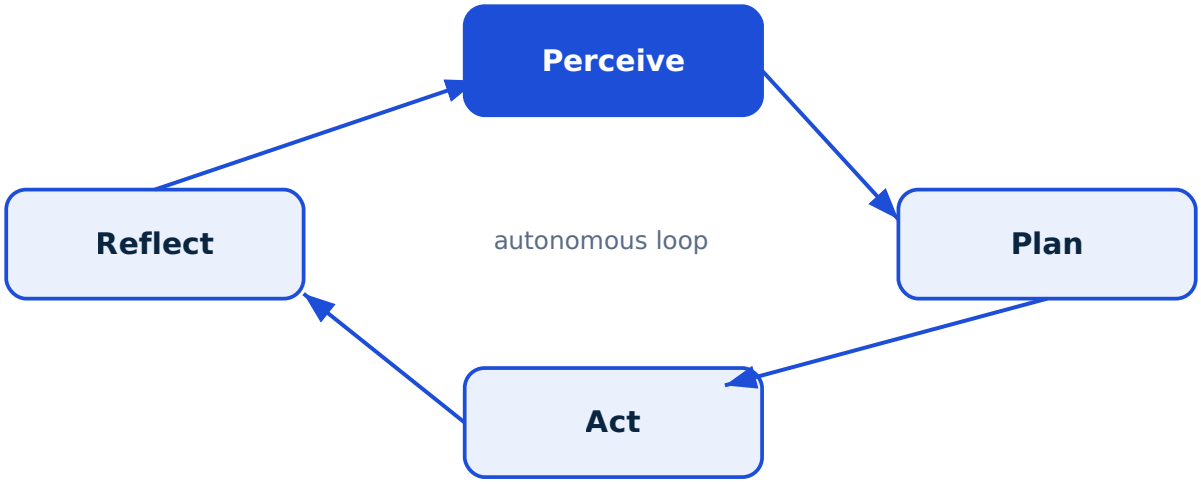

When to use

- Monitoring & autonomous operations
- Goal-driven, long-running tasks
- When the environment keeps changing

Common pitfalls

- Acting without human-in-the-loop on risk
- No reflection step (repeats mistakes)
- Missing guardrails / spend caps

Pattern-selection decision guide

A fast way to pick a starting pattern. Match the question to the pattern, then compose as needed.

If your task…	Start with
Needs reasoning plus tool calls, step by step	ReAct
Maps cleanly to one of several tools/APIs	Tool-Use / Router
Has distinct sub-skills or is parallelisable	Multi-Agent Orchestration
Demands high output quality with checkable criteria	Reflection / Critic
Spans long or multiple sessions and needs recall	Memory-Augmented Agent
Runs continuously toward a goal with little supervision	Autonomous Plan-Act Loop

> Rule of thumb: pick the simplest pattern that works, add memory when context isn't enough, and only reach for orchestration when one agent's scope gets overloaded.

Governance & security best-practices checklist

Every pattern above needs guardrails before it goes near production.

- ✓ Least-privilege access for every tool and data source
- ✓ Input/output guardrails and prompt-injection defences
- ✓ PII detection, redaction and data-handling controls
- ✓ Full audit trail of agent decisions and actions
- ✓ Rate limits and spend caps on autonomous actions
- ✓ Human-in-the-loop checkpoints for high-risk steps
- ✓ Bias detection, monitoring and regular red-teaming

Cross-pattern pitfalls & anti-patterns

Mistakes that show up no matter which pattern you choose.

Design anti-patterns

- Reaching for multi-agent when one agent would do
- No stop conditions — loops that never terminate
- Vague tool/agent descriptions that confuse routing
- Treating the LLM as deterministic

Operational anti-patterns

- Shipping without monitoring or rollback
- No evaluation harness for agent behaviour
- Unbounded cost from excess iterations
- Autonomy on risky actions with no human gate

The certification covers each pattern's failure modes and how to engineer around them.

Where pattern mastery leads

Knowing these patterns cold is what separates builders from architects on the ladder.

1. **Agentic AI Developer**
   Builds and wires individual agents.
2. **Agentic AI Expert**
   Designs, governs and evaluates systems of agents.
3. **AI Agent Architect**
   Owns the end-to-end autonomous architecture.
4. **Lead / Principal AI Engineer**
   Leads teams and sets technical direction.
5. **AI Strategy Lead**
   Owns the enterprise roadmap and ROI.

AI governance roles

Patterns plus governance is a powerful, well-paid combination — here's the parallel track.

Roles in demand

- AI Governance Lead
- Responsible AI / AI Risk Manager
- AI Compliance Specialist
- VP of AI Governance

What they own

- Governance frameworks & approval workflows
- Bias detection and mitigation
- Audit trails and incident response
- Security best practices for autonomous agents

Salary guide for these roles

US base-pay benchmarks for the senior roles that design agentic systems.

Role	Avg base (US)	Typical range
Principal AI Engineer	**$323K**	**$253K–$421K**
Lead AI Engineer	**$197K**	**$156K–$252K**
AI Agent Architect (AI Architect)	**$189K**	**$142K–$257K**
AI Solution Architect	**$185K**	**$150K–$230K**
Senior AI Engineer	**$145K–$310K**	**$400K+ total comp**

Sources: Glassdoor, 2026 (Principal/Lead AI Engineer, AI Architect, AI Solution Architect); senior total-comp from KORE1 / market data, 2026.

GSDC
Global Skill Development Council

From reference to mastery

A cheat-sheet helps you recognise patterns; the certification makes you fluent in building, combining and governing them — with proof.

36+ hrs
Expert-led video, lifetime access

32 labs
Hands-on, AI-powered builds

Daily
Live sessions with global experts

What you build

- Every pattern here, hands-on, plus how to combine them
- A capstone and recognised credential as proof

Reassurance

- 4.4★ Trustpilot · 4.8★ Reviews.io
- 7-day money-back guarantee

RISK-FREE

Backed by a 7-day money-back guarantee

Enrol with confidence — full refund within 7 days if it isn't right for you.

Enroll Now →

GSDC
Global Skill Development Council

Exam details & what's included

The exam

- 40 multiple-choice questions · 65% to pass
- 90-minute duration · closed book · English
- Valid 5 years · complimentary retake

Inside the programme

- 36+ hours expert-led video, lifetime access
- Daily live sessions · 3 one-on-one SME connects
- 2 practice exams · capstone + job support

Delivered as one programme — no separate fees, lifetime access.

Frequently asked questions

Are these all the patterns there are?

They're the six core building blocks. Most real systems combine several — master these and you can read and design almost any agentic architecture.

Can I print this as a cheat-sheet?

Yes — pages 4–9 are one pattern per page by design, plus the decision guide on page 10.

How do I choose between patterns?

Use the decision guide on page 10, and start with the simplest pattern that solves the problem.

Do the diagrams map to real frameworks?

Yes — they correspond to constructs in LangChain, LangGraph, AutoGen and RAG-based memory, all covered in the programme.

Frequently asked questions

How do governance and security fit in?

They wrap every pattern. Use the checklist on page 11 before any pattern goes to production.

Will this help in interviews?

Very much — being able to name, diagram and compare these patterns is exactly the architect-level signal hiring managers look for.

Is there a guarantee on the programme?

Yes — a 7-day, 100% money-back guarantee.

Put the patterns to work

You've got the diagrams, the decision guide and the guardrails. The next step is building them for real — combining patterns into systems you can ship and defend.

- Keep pages 4–9 as your cheat-sheet
- Use the decision guide to pick a starting pattern
- Apply the governance & security checklist before launch
- Get certified to master and combine all six

Tap the footer below or any button in this reference to enrol in the Agentic AI Expert Certification.