Claude Certified Architect – Professional

Exam Guide

Version 1.0 · Effective July 2026 · Exam code: CCAR-P · This guide is subject to change without notice.

1. About This Certification

The Claude Certified Architect – Professional certification validates that an individual can design, build, and deliver production-grade AI solutions using Anthropic's Claude platform. It is intended for practitioners working in an architect role who select appropriate models, architectures, and API patterns; apply prompt and context engineering; integrate Claude into enterprise systems; and incorporate evaluation, security, compliance, and governance considerations into their designs.

This guide is the authoritative reference for candidates preparing to sit the exam. It describes the credential's purpose, the intended audience, the exam format and content blueprint, scoring, policies, and the rules of conduct that all candidates agree to. Read it in full before scheduling your exam.

2. Purpose and Value of the Credential

The primary purpose of the Claude Certified Architect – Professional certification is to provide an independent assessment of the knowledge, skills, and abilities required to architect Claude-based solutions competently in production environments. Earning the credential signals to employers, clients, and teams that the holder can own or significantly contribute to the full lifecycle of a Claude-powered system.

Candidates who earn the credential can demonstrate that they are able to:

- Design and prototype AI-driven solutions that address real business problems
- Select appropriate models, architectures, and API patterns for specific use cases
- Implement prompt engineering and context strategies to guide model behavior
- Integrate Claude into production systems using APIs, orchestration tools, and data pipelines
- Apply evaluation, monitoring, and observability practices to ensure solution quality
- Incorporate security, compliance, and governance considerations into system design
- Collaborate with cross-functional stakeholders and communicate architectural decisions effectively

3. Intended Audience

The certification is intended for mid- to senior-level technical professionals who design, build, and deliver production-grade AI solutions using large language models, particularly Claude. This audience primarily includes solution architects, AI/ML engineers, technical leads, and senior software engineers who operate at the intersection of business requirements and technical implementation.

These professionals translate business problems into scalable AI-driven solutions, including model selection, prompt engineering, orchestration of tools and agents, context management, and ensuring system safety, compliance, and governance. They are often involved in stakeholder engagement, advising clients or internal teams, and leading architectural decisions, including discussions of security, legal, and executive considerations. Candidates typically work across industries such as financial services, healthcare, retail, technology, education, and government.

This certification is not intended for entry-level developers, casual users of Claude-based applications, or individuals without experience designing end-to-end AI systems. It also excludes roles that are purely non-technical or limited to isolated tasks such as prompt writing without broader system design responsibility.

4. Minimally Qualified Candidate Profile

The exam is targeted at the minimally qualified candidate (MQC): an experienced practitioner who combines an engineering mindset with practical, real-world experience deploying AI solutions. The MQC designs, implements, and governs Claude-powered AI solutions within production environments, translating business requirements into scalable, secure, and reliable architectures.

The MQC demonstrates the ability to design end-to-end AI systems, including prompt and context engineering, retrieval-augmented generation (RAG), API integration, orchestration, and evaluation frameworks, and understand trade-offs related to cost, latency, performance, safety, and maintainability.

Recommended experience

- A foundation in software engineering best practices (modular design, separation of concerns, scalability)
- 3+ years of experience in systems architecture or platform engineering
- 6+ months of hands-on experience with Claude or comparable LLM-based systems in production
- Experience delivering end-to-end systems from discovery through deployment and operationalization

Prerequisites: There are no mandatory prerequisites or courses required to sit this exam. The experience above is recommended, not required. The credential is awarded based on exam performance alone.

5. Exam Details at a Glance

Credential	Claude Certified Architect – Professional
Exam code	CCAR-P
Number of items	63
Item format	Multiple-choice and multiple-response items; each item states how many responses to select
Time limit	120 minutes
Delivery	Proctored: online proctored and/or test center, per program policy
Passing score	Scaled score of 720 on a scale of 100–1,000
Exam fee	$175 USD
Validity period	12 months from the date the credential is awarded
Result reporting	Pass/fail with scaled score (100–1,000), plus percent-correct by domain on the score report

6. Exam Content Outline (Blueprint)

The exam blueprint defines the content domains measured and the approximate weight of each domain on the exam. Weights reflect the relative importance of each domain to competent performance as determined through the job task analysis. The percentages indicate the approximate proportion of scored items drawn from each domain.

Domain	Content Domain	Weight
1	Solution Design & Architecture	17%
2	Claude Models, Prompting & Context Engineering	13%
3	Integration	19%
4	Evaluation, Testing & Optimization	16%
5	Governance, Safety & Risk Management	14%
6	Stakeholder Communication & Lifecycle Management	14%
7	Developer Productivity & Operational Enablement	7%
| Total	100%	

Detailed objectives by domain

Each domain below lists the tasks a candidate is expected to perform. Exam items are written against these objectives.

Domain 1: Solution Design & Architecture (17%)

- Translate business problems into Claude-based AI solutions
- Design end-to-end architectures (input → processing → output → feedback loops)
- Select appropriate architectural patterns (workflow, agentic, augmented LLM)
- Design multi-agent systems and orchestration strategies
- Apply decomposition techniques for complex problem solving
- Align solutions to business value pillars (efficiency, transformation, productivity, cost, performance SLAs)

Domain 2: Claude Models, Prompting & Context Engineering (13%)

- Select appropriate Claude models based on trade-offs
- Design system prompts, templates, and guardrails
- Apply prompt engineering techniques (zero-shot, few-shot, chain-of-thought)
- Optimize context windows and manage token usage
- Implement prompt reuse strategies (caching, modular prompts, Skills)

Domain 3: Integration (19%)

- Evaluate tool/agent configuration for capability bloat
- Analyze authentication and authorization requirements to identify security gaps
- Evaluate accuracy-latency trade-offs and justify configuration decisions
- Analyze observability challenges and select monitoring strategies at scale
- Design a RAG pipeline with appropriate chunking and indexing strategies
- Apply retrieval strategies matched to data shape and query pattern
- Evaluate connection protocols and select the appropriate integration mechanism (MCP, API/CLI, agent-to-agent)
- Evaluate progressive discovery vs. monolithic context strategy

Domain 4: Evaluation, Testing & Optimization (16%)

- Define evaluation metrics (accuracy, latency, cost, safety, security)
- Design evaluation datasets and test frameworks using mixed methodologies
- Conduct A/B testing and iterative improvements
- Diagnose system issues (prompt failure, hallucinations, model mismatch)
- Optimize token usage, latency, and cost-performance trade-offs
- Monitor system performance using logging and observability tools

Domain 5: Governance, Safety & Risk Management (14%)

- Implement guardrails and safety controls
- Identify risks, limitations, and failure modes of LLM systems
- Apply human-in-the-loop validation strategies
- Ensure compliance with regulations (e.g., GDPR, HIPAA, FedRAMP)
- Address ethical AI considerations (bias, fairness, transparency)

Domain 6: Stakeholder Communication & Lifecycle Management (14%)

- Conduct structured discovery and requirement gathering
- Communicate architectural decisions and trade-offs
- Manage stakeholder feedback loops and expectation alignment (including SLAs)
- Document architectures and provide implementation guidance
- Support lifecycle phases (discovery, design, handoff, monitoring, iteration)

Domain 7: Developer Productivity & Operational Enablement (7%)

- Configure Claude tools and environments for teams (e.g., Claude Code)
- Improve developer workflows using AI-assisted tooling
- Support debugging and operational issue resolution

7. How to Prepare

There is no single required course. Anthropic does not guarantee that any particular resource ensures a passing result. Candidates are encouraged to combine hands-on experience with the resources below:

- Study the exam blueprint in Section 6 and self-assess against each objective
- Review official Anthropic documentation for the Claude API, models, prompt engineering, MCP, and Skills
- Build and operate at least one end-to-end Claude solution, including RAG, evaluation, and observability
- Practice architectural decision-making: model selection, integration protocols, and security trade-offs
- Complete the sample questions in Section 8 to familiarize yourself with item style

8. Sample Questions

These illustrative items show the style and cognitive level of the exam. They are not drawn from the live item bank. Correct answers and rationale appear after the questions.

Sample 1 · Domain 3 — Integration

A team exposes a customer-support agent that can read tickets, draft replies, issue refunds, and delete user accounts. Support staff only ever need to read tickets and draft replies. Applying least-privilege principles, which change best reduces risk?

A. Add logging to the refund and delete tools so misuse can be audited later.

B. Remove the refund and delete tools from the agent's configuration entirely.

C. Keep all tools but add a confirmation prompt before refunds and deletions.

D. Replace the agent with a larger model that follows instructions more reliably.

Sample 2 · Domain 2 — Models, Prompting & Context

An application sends the same 8,000-token system prompt and policy document on every request, followed by a short, varying user message. Latency and cost are both concerns. Which optimization most directly addresses both?

A. Truncate the policy document to the first 1,000 tokens.

B. Switch to the smallest available model regardless of task fit.

C. Place the static system prompt and policy before the dynamic content and enable prompt caching.

D. Move the policy document into a few-shot example block.

Sample 3 · Domain 4 — Evaluation & Optimization

A RAG system suddenly returns confident but incorrect answers after a document refresh, while latency and model version are unchanged. What is the most likely first place to investigate?

A. The model weights have silently changed.

B. The retrieval/indexing step is returning irrelevant or stale chunks.

C. The temperature setting is too low.

D. The context window has shrunk.

Answer key and rationale

Sample 1: B. Least privilege means removing capabilities the role does not require, eliminating the attack surface rather than monitoring or guarding it. Logging (A) and confirmations (C) are detective/compensating controls, not removal of unnecessary privilege; model size (D) is unrelated to authorization scope.

Sample 2: C. Ordering stable content first and enabling prompt caching lets repeated prefixes be reused, reducing both time-to-first-token and per-request cost without discarding required context. Truncation (A) loses needed policy; downsizing blindly (B) risks quality; relocating to few-shot (D) does not create a cacheable, reusable prefix.

Sample 3: B. Confident-but-wrong answers following a document refresh, with model and latency unchanged, point to retrieval feeding the model poor context, for example a broken re-index or mismatched embeddings. The other options would not be triggered specifically by a document refresh.

9. How the Exam Is Scored

The Claude Certified Architect – Professional exam is a criterion-referenced assessment: each candidate is measured against a fixed performance standard, not against other candidates. You pass by demonstrating the knowledge and skills defined in the blueprint, not by outperforming a percentage of peers.

Passing standard. The passing score was established through a formal standard-setting study in which trained subject matter experts judged the level of performance expected of a minimally qualified candidate. The score is reported on a scaled range of 100–1,000, and the cut score is 720.

Result reporting. Your result is reported as a pass or fail status with a scaled score from 100 to 1,000. Your score report also shows the percentage of items you answered correctly within each content domain. Section-level percentages are provided to help you understand your performance and are not used to determine your pass or fail result, which is based on your total scaled score.

10. Registration and Scheduling

Registration and scheduling are handled through the Anthropic Partner Academy and Pearson VUE:

1. Go to the certification page for your exam on the Anthropic Partner Academy and review the exam details.
2. Download the Exam Guide and review the Certification Terms and Conditions and the Certification Exam Policy before registering.
3. Register for the exam and complete checkout. The fee shown at checkout reflects any discount that applies to your partner tier.
4. Follow the confirmation instructions to create your Pearson VUE account, then sign in to schedule your exam session.
5. Choose an available date and select either online proctoring or a Pearson test center. The exam is delivered by Pearson VUE.
6. You may cancel or reschedule up to 24 hours before your appointment. Changes made within 24 hours forfeit the exam fee.

11. Exam Policies

Identification

On exam day you must present a valid, unexpired, government-issued photo identification. The name on your ID must match the name on your registration exactly. If you need to correct the name on your registration, contact certifications-support@anthropic.com before scheduling your exam.

Accommodations

Reasonable accommodations are available for candidates with documented disabilities or needs, in accordance with applicable law. Accommodations must be requested and approved by Pearson VUE before you schedule your exam. Do not schedule your appointment until your request has been approved. Request accommodations at pearsonvue.com/us/en/test-takers/accommodations.

Retake policy

Candidates who do not pass may retake the exam after a required waiting period. Waiting periods increase with each failed attempt: 14 days after the first, 30 days after the second, and 90 days after the third. You may take an exam up to four times within a rolling twelve-month period. Limits apply per exam, so not passing one exam does not prevent you from registering for a different one. The exam fee applies to each attempt.

No-show and late arrival

Candidates who fail to appear for a scheduled exam, or who arrive after the permitted late-arrival window, forfeit the exam fee and must re-register. Cancellation and rescheduling deadlines are described in Section 10.

12. Exam-Day Experience and Rules of Conduct

The exam is administered in a standardized, secure, proctored environment. Whether you test online or at a Pearson VUE test center, the following rules apply to protect the integrity of the credential for everyone who holds it.

During the exam you must:

- Remain within view of the proctor and webcam for the entire session, if testing online
- Keep your workspace clear of notes, books, phones, secondary monitors, and other materials

- Refrain from communicating with any other person during the exam
- Not capture, copy, photograph, or reproduce any exam content in any form

Prohibited items include mobile phones, smart watches, headphones, study materials, and any recording device. Permitted items, if any, such as scratch paper provided by the proctor, are specified by Pearson VUE.

Consequences of misconduct. Cheating, attempting to access prohibited resources, or disclosing exam content may result in invalidation of your result, revocation of your credential, and a ban from future exams.

13. Confidentiality and Non-Disclosure Agreement

Before the exam begins you must accept a confidentiality and non-disclosure agreement. By accepting, you agree that all exam content, including questions, answer options, and scenarios, is the confidential and proprietary property of Anthropic, and that you will not disclose, reproduce, or distribute any portion of it. If you do not accept the agreement, the exam session ends and no refund is issued.

14. Credential Maintenance and Recertification

The Claude Certified Architect – Professional credential is valid for 12 months from the date it is awarded. Because the underlying technology evolves rapidly, the credential is time-limited so that holders maintain current knowledge.

To renew on time, you review what has changed since you certified and complete a free, non-proctored assessment on the Anthropic Partner Academy. There is no fee for on-time renewal. If your credential lapses, you must retake the full exam at the full fee to regain certified status.

If exam content changes significantly, Anthropic may require holders to retake the full exam to recertify rather than complete the renewal assessment.

Holders remain subject to the rules of conduct described in Section 12.

15. Candidate Support, Appeals, and Privacy

Support

To correct the name on your registration, contact certifications-support@anthropic.com. For all other questions, including registration, scheduling, accommodations, and results, contact Pearson VUE support at pearsonvue.com/us/en/anthropic.html.

Appeals and complaints

You may appeal a decision within 14 days of the date you are notified of it, or, for a concern about your exam result, within 14 days of your exam date. Submit your appeal to Pearson VUE support at the link above. Appeals are reviewed under the program's appeals policy. The standard-setting outcome and the content of individual exam items are not subject to appeal.

Privacy

Personal data collected during registration and testing is handled in accordance with Anthropic's privacy policy, Pearson VUE's privacy policy, and applicable data-protection law.

16. Document Control

Version	Summary of change	Date
1.0	Initial publication	July 2026