

Claude Certified Architect - Professional

CCA-P
 Official exam code CCAR-P

Mid-senior solution architects / AI-ML engineers (3+ years systems arch, 6+ months Claude in prod).

63	720	7	120 min	\$175	12 mo
SCORED ITEMS	TO PASS, OF 1000	WEIGHTED DOMAINS	TIME LIMIT	LIST PRICE	VALIDITY

This is a ledger, not a brochure

Most exam guides are read once. This one is meant to be marked up: you record where you actually are against each domain, work the two that cost you most, log two timed mocks, and end on a line that says book or do not book. Print it, or annotate it on screen.

Every *exam* figure — item count, cut score, domain weights, fee, validity — is transcribed from Anthropic's published blueprint, version 1.0-2026-07. The study-pace and mock-timing figures are this site's own, not Anthropic's. It does not reproduce exam content, and it is not the official guide.

Take the free CCA-P diagnostic

Read the full guide online



What the CCA-P weighs

Weights are the share of scored items drawn from each domain. The item counts are those weights applied to a 63-question exam, which is the number worth budgeting study hours against.

D1	Solution Design & Architecture		17%	~11 items
D2	Claude Models, Prompting & Context Engineering		13%	~8 items
D3	Integration		19%	~12 items
D4	Evaluation, Testing & Optimization		16%	~10 items
D5	Governance, Safety & Risk		14%	~9 items
D6	Stakeholder Communication & Lifecycle		14%	~9 items
D7	Developer Productivity & Operational Enablement		7%	~4 items

Where the exam actually is

Integration alone carries 19% — about 12 of the 63 items. The three heaviest domains together are 52% of the paper.

How the result comes back

Pass or fail with a scaled score from 100 to 1000, plus the percentage of items you answered correctly in each domain. Those per-domain percentages are reported to help you read your performance; they do not decide the pass.

Tick what you know. Cross what you do not.

That is the whole system. Two marks, and you add them up.

- ☒ **Tick.** You could explain this to someone else without looking it up.
- ☐ **Cross.** You could not. This is what you study.

Then add up your ticks

Each domain page has a score box. Count your ticks, write the number in, and compare it with how many questions that domain is worth. A domain worth 16 questions where you ticked 2 of 6 topics is where your exam is won or lost.

Do it in this order

1. **Take the free diagnostic first, before you revise.** Write your score on page 5. Revising first only tells you what looks familiar.
2. **Go through pages 6 to 10 and mark every topic** tick or cross. Be hard on yourself here. A generous tick costs you an exam fee later.
3. **Study the two domains with the most crosses and the highest question count.**
4. **Sit a full mock, then mark the same pages again** in the second column. Anything that went from tick to cross is what the real exam would have caught.

Take the free diagnostic

Open the 63-question mock

How the questions are written

Every question follows the same shape. Knowing the shape is worth a few marks on its own, because you stop misreading what is being asked.

1. THE SITUATION	A short description of a real setup and what is going wrong. Two or three sentences. Every number in it matters.
2. THE CATCH	A phrase that rules most answers out on its own — "as a first step", "without adding infrastructure", "easiest to maintain". Read this before you read the answers.
3. THE QUESTION	"Which is most effective" and "what would you check first" are different questions. Answer the one asked.
4. HOW MANY TO PICK	Some questions want one answer, some want two or more. The question always tells you. Check this first.
5. THE ANSWERS	Usually one right, one that does too much, one that fixes a different problem, and one that is simply wrong.



When you are not sure of the answer

Check how many answers to pick	It changes the question. The wording always tells you.
Use the catch to rule answers out	"As a first step" or "easiest to maintain" usually removes two of them straight away.
Prefer the answer that changes the design	Where one answer changes how the thing is built and another just changes the model or the wording, the build answer is usually the one being tested.
Watch for the answer to a different problem	It will be correct in general and wrong here. Check it against the symptom described.
Flag it and come back	You have about 1.9 minutes a question. A flagged question you return to beats a guess you never revisit.

Write down where you are before you revise

A number you wrote two weeks ago is the only proof the fortnight worked. Take the diagnostic without revising first, and fill the top row in.

Your diagnostic

DATE	QUESTIONS RIGHT	OUT OF	WEAKEST DOMAIN
		10	

How you did in each domain

The middle column is how many questions that domain is worth on the real exam. Write in how many you got right, and the difference is your gap — the number of marks currently sitting on the table.

DOMAIN	QUESTIONS ON THE EXAM	YOU GOT RIGHT	GAP
D1 Solution Design & Architecture	about 11		
D2 Claude Models, Prompting & Context Engineering	about 8		
D3 Integration	about 12		
D4 Evaluation, Testing & Optimization	about 10		
D5 Governance, Safety & Risk	about 9		
D6 Stakeholder Communication & Lifecycle	about 9		
D7 Developer Productivity & Operational Enablement	about 4		
Total	63		

Your mocks

A mock score is a guide to how ready you are, not a prediction of your official result. It uses the same 720 pass mark so you can compare week to week, but only Anthropic scores the real exam.

DATE	QUESTIONS RIGHT	OUT OF	PASSED?	WHAT WENT WRONG
		63		
		63		

Sit a full 63-question mock

Solution Design & Architecture

D1

17% of the exam · about 11 of 63 items

Translate business problems into Claude-based AI solutions and design end-to-end architectures (input -> processing -> output -> feedback loops).

YOUR SCORE FOR THIS DOMAIN

After the diagnosticof 11

After the mockof 11

TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK
Translating Business Requirements into Claude Solutions Translating a business problem into a Claude solution means producing a build-ready requirements spec, not picking an architecture pattern.	☐	☐
End-to-End Architecture Design (Input to Feedback Loop)	☐	☐
Workflow vs Agent: Choosing the Right Architecture Pick the cheapest tier that actually satisfies the task: an augmented LLM call for a single unaided prompt, a workflow when the steps are known in advance, an agent only when the model must dynamically decide its own path.	☐	☐
Aligning Solutions to Business Value Pillars Every architectural choice - pattern, model tier, guardrail overhead - has to trace back to a named business-value pillar an executive sponsor recognizes: efficiency, transformation, productivity, cost, or performance SLAs.	☐	☐

NOTES – WHAT YOU GOT WRONG, AND WHY

Two domains

D2

Claude Models, Prompting & Context Engineering

13% of the exam · about 8 of 63 items

YOUR SCORE FOR THIS DOMAIN	After the diagnostic _____ of 8	After the mock _____ of 8
TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK
	☐	☐
	☐	☐
	☐	☐

D7

Developer Productivity & Operational Enablement

7% of the exam · about 4 of 63 items

YOUR SCORE FOR THIS DOMAIN	After the diagnostic _____ of 4	After the mock _____ of 4
TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK
	☐	☐
	☐	☐
	☐	☐

NOTES – WHAT YOU GOT WRONG, AND WHY

Integration

D3

19% of the exam · about 12 of 63 items

YOUR SCORE FOR THIS DOMAIN

After the diagnostic of 12

After the mock of 12

TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK
Authentication & Access Control for Agent Systems Agent authentication and authorization is a stack of controls, not one setting: which identity the process runs as, what credentials each tool call carries, and which actions need a human in the loop.	☐	☐
Accuracy-Latency-Cost Tradeoff Analysis	☐	☐
Observability & Monitoring for Agents at Scale	☐	☐
Progressive Discovery vs Monolithic Context Strategy	☐	☐

NOTES – WHAT YOU GOT WRONG, AND WHY

Evaluation, Testing & Optimization

D4

16% of the exam · about 10 of 63 items

Defining evaluation metrics across accuracy, latency, cost, and safety; designing evaluation datasets and mixed-methodology test frameworks; running A/B tests and iterative improvements; diagnosing system issues (prompt failure, hallucinations, model mismatch); and monitoring production performance with logging and observability tooling.

YOUR SCORE FOR THIS DOMAIN

After the diagnostic of 10

After the mock of 10

TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK

NOTES – WHAT YOU GOT WRONG, AND WHY

Two domains

D5

Governance, Safety & Risk

14% of the exam · about 9 of 63 items

YOUR SCORE FOR THIS DOMAIN	After the diagnostic _____ of 9	After the mock _____ of 9
TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK
Guardrails & Safe Deployment Practices	☐	☐
LLM System Risks & Failure Modes	☐	☐
Regulatory Compliance for AI Systems (GDPR, HIPAA, FedRAMP)	☐	☐

D6

Stakeholder Communication & Lifecycle

14% of the exam · about 9 of 63 items

YOUR SCORE FOR THIS DOMAIN	After the diagnostic _____ of 9	After the mock _____ of 9
TOPICS IN THIS DOMAIN	KNOW IT	AFTER MOCK
Structured Stakeholder Discovery	☐	☐
Communicating AI Decisions & Value to Stakeholders	☐	☐
Stakeholder Feedback Loops & SLA Alignment	☐	☐

NOTES – WHAT YOU GOT WRONG, AND WHY

Work every concept, free

Two weeks, in the order that pays

People already working with Claude day to day report 2-4 weeks at 60-90 minutes a day.
From a cold start, 6-8 weeks. The order matters more than the hours.

WHEN	WHAT	RECORDED ON
Day 1	Diagnostic, cold. No revision first. Write the per-domain percentages down.	Page 5
Day 1	Mark every domain in the first-pass column, honestly.	Pages 6–10
Days 2–6	Integration first — 19% of the paper, and the single largest block of marks available.	Page 8
Day 7	Full 63-question mock , timed, no notes. The mock here runs tighter than the real 120 minutes on purpose — train faster than race pace.	Page 5
Days 8–12	Solution Design & Architecture , plus whatever the mock exposed. Re-mark in the second column.	Page 6
Day 13	Second timed mock. Different questions, same conditions.	Page 5
Day 14	The decision on page 12. Book, or take another week.	Page 12

Fourteen days, at a glance

1
MEASURE
☐

2
WORK
☐

3
WORK
☐

4
WORK
☐

5
WORK
☐

6
WORK
☐

7
MOCK
☐

8
WORK
☐

9
WORK
☐

10
WORK
☐

11
WORK
☐

12
WORK
☐

13
MOCK
☐

14
DECIDE
☐

Tick each square as the day is done. Two squares are timed mocks; the last one is the decision on page 12.

Day 1 starts here

Day 7 mock



Booking, sitting, and what happens if it goes wrong

Registering

Registration and scheduling run through the Anthropic Partner Academy and Pearson VUE: open your exam's page on the Partner Academy, read the terms and exam policy, register and check out (the fee shown reflects any partner-tier discount), create your Pearson VUE account from the confirmation, then choose a date and either online proctoring or a test centre. You can cancel or reschedule up to 24 hours ahead; inside 24 hours the fee is forfeit.

If you do not pass

Retakes wait 14 days after the first attempt, 30 after the second, 90 after the third, up to four attempts in a rolling twelve months. The fee applies to every attempt. Limits are per exam, so a fail on one does not block registering for another.

Pacing, on the day

120 minutes for 63 scored items is 1.9 minutes each. Check yourself at the quarter marks.



The decision

Second mock cleared 720 with room

☐ yes☐ no

No domain still carrying empty boxes above 15% weight

☐ yes☐ no

Sat both mocks under time, without notes

☐ yes☐ no

Three yes: book it. Any no: one more week. Booked for

Sit the mock that decides it

Registration walkthrough

On the day

Government-issued photo ID, and the name on it must match your registration exactly. The workspace has to be clear of notes, books, phones, secondary monitors and any recording device. You accept a confidentiality agreement before the exam begins; declining ends the session with no refund.

Afterwards

The credential is valid 12 months. On-time renewal is a free, non-proctored assessment on the Partner Academy; let it lapse and you retake the full exam at the full fee.

