



DIGITAL WORLD PULSE

CLAUDE HAIKU

5.5

FEATURES, PRICING & PRACTICAL PROMPTS
2026 GUIDE

18 copy-ready prompts with explained, practical examples

INDEPENDENT PRACTICAL GUIDE | RESEARCHED OCTOBER 21

Original Digital World Pulse practical guide

Research snapshot: October 21, 2026



DIGITAL WORLD PULSE

CONTENTS

From a precise brief to a verified practical result

03	What Haiku 5.5 Is For	
04	Access, Settings & Privacy	
05	Write A Brief That Works	Prompt 01
06	Summaries That Preserve Decisions	Prompt 02
07	Extract Fields Without Guessing	Prompt 03
08	Classify Messages Consistently	Prompt 04
09	Support Replies From A Policy	Prompt 05
10	Write Useful Copy From A Brief	Prompt 06
11	Localise Without Changing Facts	Prompt 07
12	Small Data, Correct Totals	Prompt 08
13	Fix A Small Bug, Then Check It	Prompt 09
14	A Small Api Request	Prompt 10
15	Choose Effort Deliberately	Prompt 11
16	Api Costs Without The Guesswork	Prompt 12
17	Know When To Use A Larger Model	Prompt 13
18	Repair The Specific Failure	Prompts 14-16
19	A Seven-Day Practical Trial	Prompts 17-18
20	Sources, Method & Disclaimer	

HOW TO USE THE EXAMPLES

Copy a numbered prompt, supply your own material where requested, then compare the response with the worked example and check. Fictional practice data is labelled. The example responses are editorial illustrations, not captured Claude sessions.



WHAT HAIKU 5.5 IS FOR

A small model for focused, repeatable work

Anthropic released Claude Haiku 5.5 on October 7, 2026. It is positioned for high-volume tasks where response time and cost matter: summaries, classification, extraction and routing. This guide turns those categories into concrete exercises whose results you can inspect.

THE MODEL AT A GLANCE

INPUT / OUTPUT	Text and images in; text out. Writing code is text generation, not proof that the code ran.
CAPACITY	The standard model specification lists a 1M-token context window and up to 128K output tokens. Limits are not a recommendation to send everything.
MODEL ID	claude-haiku-5-5 on the Claude API. Cloud-provider identifiers and product controls can differ.

A USEFUL FIRST WORKLOAD

Choose a task with clear inputs and a short answer: turn a meeting note into an action list, extract five fields from a brief, or classify messages into four defined categories. Set a rule for missing information and compare the output with the source. A small task makes mistakes easier to diagnose.

MODEL CAPABILITY IS NOT APP ACCESS

A model can interpret text, but browsing, account access and execution need the relevant product tools and permissions. Its listed knowledge cutoff is June 2026; current prices, policies and events need current sources. Anthropic's speed and capability statements are vendor claims, not independent measurements made for this guide.

Start with a narrow task and an observable definition of success.

ACCESS, SETTINGS & PRIVACY

Choose the route before you compare costs or features

Anthropic's Haiku product page lists selection in Claude.ai for Free, Pro, Max, Team and Enterprise users, across web and mobile. It also lists Claude Code and developer platforms. Actual limits, availability and organisation controls still apply to your account.

TWO DIFFERENT WAYS TO USE THE MODEL

CLAUDE APP	Use the model menu beside the send button; open More models if needed. Select Haiku 5.5 and check the displayed effort setting.
CLAUDE API	Developers send requests using the model ID, with usage billed under the applicable API arrangements. The token prices on page 16 are not an app subscription price.

BEFORE YOU PASTE REAL WORK

Remove unnecessary personal details, secrets and customer identifiers. For consumer accounts, review Model Improvement and the current privacy terms. Anthropic says consumer chats can be used for improvement when allowed, when flagged for safety review, or under other explicit opt-ins. Incognito chats are not used to improve Claude, even when Model Improvement is enabled.

COMMERCIAL DATA AND FEEDBACK

Anthropic says commercial inputs and outputs are not used for model training by default. Explicit feedback or other opt-ins can change how a submitted conversation is used. This is a training-policy distinction, not a blanket promise about retention, confidentiality or who can access every product. Check the rules that apply to your route.

A SIMPLE SETUP TEST

Select the model, paste a fictional sample from this guide and request a short structured result. Confirm that the displayed model is Haiku 5.5. If a model is missing in an enterprise account, ask the administrator about its availability rather than assuming that another model is equivalent.

Check the selected model, the billing route and the data rules separately.

WRITE A BRIEF THAT WORKS

State the task, input, format, limits and check

For a focused task, a compact brief often gives a cleaner result than a long list of generic instructions. Specify what the answer must contain, what it must leave out and how to handle absent information. The prompts here are original editorial exercises, ready to copy.

COPY PROMPT 01 - A COMPLETE FIRST TASK

Use only these fictional release notes. Added: CSV export. Fixed: duplicate reminder emails. Known issue: dark-mode chart labels are hard to read. No release date is supplied. Write a customer update with three headings: Added, Fixed, Still investigating. Keep it under 90 words. Preserve the known issue, add no promised fix date and do not claim a release date. End with one question the team must answer before publication.

WORKED EXAMPLE - EDITORIAL RESPONSE

Added
CSV export is now available.
Fixed
We corrected the issue that caused duplicate reminder emails.
Still investigating
Chart labels can be difficult to read in dark mode.
Before publication: What is the confirmed release date?

WHY THIS IS A STRONGER BRIEF

The source material is short, the audience is defined and the output has a fixed structure. The model is told to preserve an inconvenient limitation rather than polishing it away. The missing date remains a question. Check all three source items against the result; a nicer sentence does not justify adding a feature or promise.

REUSE THE PATTERN

Replace the notes and headings, then specify your own length limit and missing-data rule. Give a concrete correction if the answer fails. You can ask for concise reasons for a decision without requesting hidden internal reasoning.

A precise brief makes the result easier to review, not automatically correct.

SUMMARIES THAT PRESERVE DECISIONS

Separate agreement, proposals and unresolved questions

The important part of a meeting summary is often the difference between a decision and an idea. A proposed deadline is not an agreed deadline. Keep that distinction visible so a fluent summary does not create commitments that nobody made.

COPY PROMPT 02 - NOTES TO ACTIONS

Summarise these fictional notes using only the supplied facts. Maya: I can finish the guide by 28 October 2026 if the source pack arrives by 24 October. Leo: I will send the source pack on 24 October. Alex: Could we publish on 29 October? No one replies to Alex's proposal. Maya also needs a confirmed image brief; no owner is assigned to that brief. Return decisions, actions with owners and dates, and unresolved questions. Keep the publication date unconfirmed. Do not assign an owner that is not stated.

WORKED EXAMPLE - ACTION TABLE

Item	Owner	Date	Status
Send source pack	Leo	24 Oct 2026	Committed
Finish guide	Maya	28 Oct 2026	Conditional on sources
Publish guide	Not stated	29 Oct proposed	Not agreed
Confirm image brief	Not assigned	Not stated	Unresolved

EXPLAIN THE UNKNOWNNS

Leo's delivery is a commitment in the notes. Maya's finish date depends on that delivery. Alex's suggestion remains a proposal. The result should ask who confirms the image brief, when it is due, and whether the publication date is accepted. That is more useful than silently converting every sentence into a completed decision.

BEFORE YOU USE THE SUMMARY

Check each owner, deadline and conditional phrase against the source. If the source is a transcript, ask for a timestamp or quotation locator for disputed items. Do not treat a summary as proof that a message was sent, an event was created or a task was accepted.

A good summary keeps uncertainty visible and gives it a practical next step.

EXTRACT FIELDS WITHOUT GUESSING

Turn a short source into a consistent JSON record

Extraction works best when field names, types and missing values are defined in advance. The sample below is fictional and complete enough to practise. A request for JSON in a chat is not the same as using the API's schema-constrained structured-output feature.

COPY PROMPT 03 - A SMALL JSON EXTRACTION

Extract only from this fictional note: Order PO-204 is for North Studio: three desk lamps at USD 40 each. Delivery costs USD 12. Tax and the customer email are not stated. Required by 28 October 2026. Return JSON only, with keys order_id, customer, quantity, unit_price_usd, shipping_usd, subtotal_before_tax_usd, tax_usd, email, due_date. Use numbers for money and quantity, null for absent values, and YYYY-MM-DD for the date. Include delivery in subtotal_before_tax_usd. Show no invented tax-inclusive total.

WORKED EXAMPLE - VALID JSON, FICTIONAL ORDER

```
{
  "order_id": "PO-204", "customer": "North Studio",
  "quantity": 3, "unit_price_usd": 40,
  "shipping_usd": 12, "subtotal_before_tax_usd": 132,
  "tax_usd": null, "email": null,
  "due_date": "2026-10-28"
}
```

HOW TO CHECK IT

Three lamps cost USD 120; adding delivery gives USD 132 before any tax. The email and tax remain null. Check key names, numeric types and the date, then parse the JSON rather than judging it by appearance. For a production API workflow, use a supported schema and still validate the extracted facts against the source.

A valid structure does not prove that the values inside it are accurate.



CLASSIFY MESSAGES CONSISTENTLY

Define categories and an explicit route for ambiguity

Haiku's intended uses include classification and routing. The difficult part is usually the policy: what belongs in each class, which rule wins, and what happens when the message is unclear. Decide that before classifying a large inbox or dataset.

COPY PROMPT 04 - FOUR TICKETS, CLEAR RULES

Classify these fictional tickets. Labels: billing = charges or invoices; access = sign-in problems; bug = a feature fails after sign-in; review = unclear or outside those categories. For mixed tickets, use review. Return id, label and a reason of at most 12 words. A: I was charged twice this month. B: My password reset link has expired. C: I can sign in, but CSV export gives an error. D: My account looks wrong and I don't know why. Use only these labels; do not send replies or change accounts.

WORKED EXAMPLE - FICTIONAL TICKETS

ID	Label	Reason
A	billing	Reports a duplicate charge.
B	access	Password reset prevents access.
C	bug	Export fails after successful sign-in.
D	review	Problem is too vague for a reliable category.

WHY 'REVIEW' IS A USEFUL ANSWER

D could be about billing, access or a display issue. Forcing it into a confident category hides that ambiguity. Before scaling up, label a small reference set yourself, include mixed and vague examples, and compare the model's labels with those decisions. A self-reported confidence number is not a calibrated probability.

KEEP ROUTING SEPARATE FROM ACTION

A category can feed a queue, but a classifier should not automatically issue a refund or change a customer's account. Define the next step for review cases, and inspect errors by category so you can improve the rules rather than changing labels at random.

Consistent routing starts with a clear policy and a route for exceptions.

SUPPORT REPLIES FROM A POLICY

A helpful draft should stay inside the supplied rules

For customer support, give the model the policy it must use. A warm tone is useful, but it should not turn into an invented refund, exception or promise. This exercise uses a fictional store policy, not a statement about any customer's legal rights.

COPY PROMPT 05 - A POLICY-BOUND REPLY

Draft a reply under 100 words using only this fictional policy: unused items may be returned within 14 calendar days of delivery; the customer must provide an order ID; the support team reviews eligibility before authorising a return; refund timing is not specified. Customer: My unused item arrived on 10 October 2026. Today is 21 October. Can I return it? I cannot find my order number. Explain the next step without approving a return, promising a refund or inventing a processing time. Do not send the reply.

WORKED EXAMPLE - DRAFT, NOT A RETURN AUTHORISATION

Hi, thanks for getting in touch. Based on the dates you provided, 11 calendar days have passed since delivery, so your unused item appears to be within our 14-day return window. We need the order ID before our support team can review eligibility. Please check your order confirmation or account order history and share that ID when you find it. The team can then review the request and explain the next steps.

CHECK THE FACTS AND THE PROMISES

21 October minus 10 October is 11 days. The policy requires an order ID and review, so the reply asks for that information. It does not claim the return is approved, invent a refund date or turn an apparent match into a guarantee. An actual store must use its own current policy and applicable requirements.

WHEN A HUMAN SHOULD TAKE OVER

Ask the model to flag missing policy terms, conflicting dates or requests for an exception. Give support staff the source policy and draft together so they can review the decision, not only the wording.

Draft the helpful next step without inventing authority to grant it.

WRITE USEFUL COPY FROM A BRIEF

A concrete audience and a modest claim produce a better draft

A short promotional paragraph is a good bounded writing task. Supply the real offer and the intended audience, then prohibit claims the source material cannot support. The example below promotes a fictional workshop and includes no invented results or scarcity.

COPY PROMPT 06 - A COMPLETE WRITING CASE

Write a 70-90-word paragraph for a fictional beginner workshop called Clear AI Briefs. Facts: a 60-minute online session teaches task definition, useful input, output formats and checking results. Attendees practise on one sample task. Audience: independent writers and small-business owners. Tone: clear, welcoming, practical. No price or date is supplied, so do not invent them. Do not promise income, rankings or guaranteed time savings. Finish with a simple invitation to read the workshop details.

WORKED EXAMPLE - 74-WORD DRAFT

Give your next AI task a clearer starting point. Clear AI Briefs is a 60-minute online workshop for independent writers and small-business owners who want a practical way to ask for useful results. You will learn to define the task, supply relevant input, choose an output format and check the response. During the session, you will practise on one sample task and see how a better brief changes the work. Read the workshop details.

WHY THIS DRAFT MATCHES THE INPUT

The paragraph describes what is taught and what the attendee practises. It adds no price, date, guaranteed saving or unsupported testimonial. The invitation asks the reader to see the details rather than pretending registration is already open. Check the word count and every factual phrase against the brief.

A FOCUSED REVISION REQUEST

If the tone is too promotional, point to the sentence that causes it and ask for a plainer version with the same facts and length limit. Keep the factual source brief unchanged while revising style, so the model cannot solve a wording problem by inventing a stronger offer.

Useful marketing copy explains the actual offer and respects what is still unknown.



LOCALISE WITHOUT CHANGING FACTS

Keep names, numbers and conditions stable across languages

Translate for the audience while preserving the offer's exact facts. Specify language, tone and fixed details. For important public copy, ask a fluent reviewer to check the result.

COPY PROMPT 07 - ENGLISH TO ROMANIAN

Translate this fictional offer into natural Romanian for a general audience. Keep the brand name, price, quantity, deadline and time zone unchanged. Use a friendly, clear tone and correct Romanian diacritics. Add no claims or conditions. Source: Clear Notes includes 12 reusable planning templates for RON 99. The offer ends on 27 October 2026 at 18:00 Europe/Bucharest. The templates are digital files; no physical items are shipped. Return the translation, then a short fact-check table comparing the five fixed details.

WORKED EXAMPLE - ROMANIAN COPY

Clear Notes include 12 șabloane de planificare reutilizabile, la prețul de 99 RON. Oferta se încheie pe 27 octombrie 2026, la ora 18:00, Europe/Bucharest. Șabloanele sunt fișiere digitale; nu se expediază produse fizice.

WORKED EXAMPLE - DETAILS THAT MUST SURVIVE

Detail	Source	Translation check
Brand	Clear Notes	Unchanged
Quantity / price	12 / RON 99	12 / 99 RON
Deadline	27 October 2026	27 octombrie 2026
Time / zone	18:00 Europe/Bucharest	Unchanged
Delivery	Digital files only	No physical shipping

REVIEW MEANING, NOT ONLY FLUENCY

Moving the currency after the number is a wording choice. Changing 99 to 99.99 changes the offer. Recheck numbers and conditions after every rewrite.

Localise the language while preserving the offer's numbers and conditions.



SMALL DATA, CORRECT TOTALS

Use a weighted rate when the groups have different sizes

Use a small CSV to practise analysis that you can recalculate. Here, return rate means returned units divided by shipped units. The rows describe fictional shipments and should not be treated as business benchmarks, forecasts or evidence of a cause.

COPY PROMPT 08 - A CHECKABLE CSV TASK

Analyse this fictional CSV. Return each region's return rate and the overall rate, using returned/shipped. Show totals and formulas. Do not average the regional percentages to calculate the overall rate. Flag missing values, zero denominators or returned units greater than shipped units. Explain what these rows cannot tell us.

region,shipped,returned

North,100,4

South,300,3

East,100,0

WORKED EXAMPLE - VERIFIED ARITHMETIC

Region	Shipped	Returned	Return rate
North	100	4	4.00%
South	300	3	1.00%
East	100	0	0.00%
Total	500	7	1.40%

WHY THE OVERALL RATE IS 1.40%

$7 / 500 \times 100 = 1.40\%$. The simple mean of 4%, 1% and 0% is about 1.67%, which gives the smaller regions too much influence. North has the highest rate, but these rows do not establish a cause. Before using real data, define the date range and return cohort, then check the denominator and a sample row against the original file.

Count the events and the population before you trust a percentage.

FIX A SMALL BUG, THEN CHECK IT

The report should say what actually ran

Haiku is suited to focused coding tasks, but a proposed fix is not evidence of a passed test. Anthropic's model-specific guidance recommends asking coding agents for real verification. This toy function provides a small example with outcomes you can check independently.

COPY PROMPT 09 - A BOUNDED BUG FIX

Fix this Python function: `def discounted_total(price, qty, pct): return price * qty * pct / 100`. It should subtract the discount, return a value rounded to two decimals, reject negative prices, reject negative or non-integer quantities, and require a discount from 0 to 100. Test 50,3,10 -> 135.00; 50,3,0 -> 150.00; and a 120% discount -> `ValueError`. Return corrected code and the tests you actually ran. If execution is unavailable, say so. Do not expand this into an app.

WORKED EXAMPLE - ORIGINAL EDITORIAL FIX

```
def discounted_total(price, qty, pct):
    if price < 0 or qty < 0 or int(qty) != qty:
        raise ValueError("Invalid price or quantity")
    if not 0 <= pct <= 100:
        raise ValueError("Invalid discount")
    return round(price * qty * (1 - pct / 100), 2)
```

EXPLAIN THE BUG AND VERIFY THE FIX

The original returns the discount amount: $50 \times 3 \times 10\% = 15$. The corrected result is $150 - 15 = 135$. The printed checks for this original sample were run while preparing the guide; they are not a captured Claude session. This is a teaching example, not a full money-handling library. Test edge cases that matter to your real application.

Require observed results from a real check, or an honest statement that none ran.

A SMALL API REQUEST

Use the right model ID and keep the first call simple

Developers can use POST /v1/messages on the Claude API with their own authenticated account. The JSON body below is a concrete request example, not a full client or a request executed for this guide. Keep your API key out of prompts, public pages and source files.

COPY PROMPT 10 - THE REQUEST'S USER MESSAGE

Label 'I was charged twice.': billing, access, bug, or other. Label only.

REQUEST BODY - JSON SYNTAX CHECKED, NOT SENT

```
{
  "model": "claude-haiku-5-5",
  "max_tokens": 256,
  "thinking": {"type": "disabled"},
  "output_config": {"effort": "low"},
  "messages": [{
    "role": "user",
    "content":
      "Label 'I was charged twice.': billing, access, bug, or other. Label only."
  }]
}
```

EXPECTED TEXT AND AN IMPORTANT LIMIT

The desired text is billing. Inspect the returned content and stop reason; do not assume success because a request was accepted. Haiku 5.5 allows thinking to be disabled on the API at low through high effort. The Claude app's thinking controls differ. Omit temperature, top_p and top_k for this model. For a machine-consumed result, validate it against the four allowed labels.

Start with one request, an exact output contract and a validation step.



CHOOSE EFFORT DELIBERATELY

Use a setting for the task, then measure the result

Haiku 5.5 supports low, medium, high, xhigh and max effort; medium is the default. Effort affects token use, latency and quality. It is a product or API setting, not something guaranteed by writing 'think harder' in a normal prompt.

A PRACTICAL STARTING POINT

Effort	Where to start testing
low	Short, routine tasks with an exact result contract.
medium	The default starting point for most work.
high	Longer work or tasks with strict instructions.
xhigh / max	Use only if your evaluations justify the extra effort; compare larger models too.

COPY PROMPT 11 - SCORE QUALITY, NOT LENGTH

Evaluate these summaries using the source only. Source: Guide due 28 Oct if sources arrive 24 Oct; publication date proposed for 29 Oct but not agreed. A: 'The guide will be published on 29 Oct.' B: 'The guide targets 28 Oct, conditional on sources arriving 24 Oct. Publication on 29 Oct remains unconfirmed.' Choose the more faithful summary. Explain the decisive difference in two sentences and list any unsupported claim.

WORKED EXAMPLE - A SMALL EVALUATION

B is more faithful because it preserves the dependency and the unconfirmed publication proposal. A turns an unapproved proposal into a definite publication commitment. The error matters more than the difference in length.

CHANGE ONE VARIABLE AT A TIME

Run the same representative task at two settings, keeping prompt and input unchanged. Compare correctness and editing effort, not only response length. Anthropic warns that low effort in long agent prompts can lead to skipped searches, checks or unfinished work. Higher settings are not a substitute for an explicit output contract.

Keep the least expensive setting that meets your measured quality standard.

API COSTS WITHOUT THE GUESSWORK

Prompt length changes the rate for the entire request

These are Claude Platform token prices, not Claude app subscription charges. Haiku 5.5 uses one price tier for prompts up to 100,000 tokens and a higher tier above that threshold. The higher rates apply to the request, not only to the excess input.

USD PER 1 MILLION TOKENS - STANDARD RATES

Token category	Prompt <=100K	Prompt >100K
Uncached input	USD 0.10	USD 0.50
Output	USD 0.50	USD 2.50
Cache reads	USD 0.01	USD 0.05

COPY PROMPT 12 - CALCULATE A SMALL API BUDGET

Calculate the token-only cost of 100 fictional requests. Each has 8,000 uncached input tokens and 600 billable output tokens. Rates: USD 0.10 per million input and USD 0.50 per million output. Show cost per request and total. Assume every prompt stays below 100,000 tokens and there are no tools, cache writes, retries or other charges. Do not treat the output allowance as actual usage.

WORKED EXAMPLE - VERIFIED COST CALCULATION

Per request: $8,000 / 1,000,000 \times \text{USD } 0.10 = \text{USD } 0.0008$ input; $600 / 1,000,000 \times \text{USD } 0.50 = \text{USD } 0.0003$ output. Total = USD 0.0011 per request, or USD 0.11 for 100 requests. At the higher tier, 200,000 input plus 1,000 output tokens would cost USD 0.1025 before other charges.

CHECK ACTUAL USAGE

Prompt length includes cache reads and writes. Cache writes, tools and special routing can add costs; check the full price list. The newer tokenizer changes counts versus Haiku 4.5, so count with the target model and inspect actual usage. This calculation excludes those extras and assumes the stated billable counts.

Budget from measured usage, the correct tier and the charges your workflow actually incurs.



KNOW WHEN TO USE A LARGER MODEL

Route by task difficulty and the cost of an error

A small model's low price is useful when it completes the task reliably. Anthropic positions Haiku for narrow, repeatable tasks and recommends larger models for more complex agentic coding. Treat model selection as a decision to evaluate on your work, not a universal ranking.

COPY PROMPT 13 - ROUTE A SMALL WORKLOAD

Propose an evaluation route for these tasks. A: label support messages using four clear categories. B: interpret conflicting clauses in a high-value contract. C: fix one discount calculation with supplied tests. D: refactor several services without a complete test suite. Return task, first approach and required review. Do not claim any model is infallible. Distinguish a small-model trial, a larger-model comparison and qualified human review.

WORKED EXAMPLE - AN EDITORIAL ROUTING PLAN

Task	First approach	Required check
A	Trial Haiku on a labelled sample	Compare labels and ambiguous cases
B	Qualified human leads; AI organises text	Legal interpretation needs expert review
C	Try a focused fix	Run the supplied checks and edge cases
D	Plan scope; compare a larger model	Human oversight and a verification strategy

DO NOT OPTIMISE PRICE IN ISOLATION

Compare the total cost of a correct, reviewed outcome. Cheap output that requires repeated repairs can cost more time than a stronger model's first draft. Preserve the same input and acceptance criteria when comparing models. Avoid giving a broad task to an unattended workflow until its permissions, tests and stopping conditions are clear.

KEEP THE ROLES CLEAR

A model can extract a clause, organise evidence or propose code. That does not make it the decision-maker for every consequence. Match the level of human review to the task and the risk of an incorrect result.

Choose the workflow that reaches a reliable reviewed result, not just the lowest token price.

REPAIR THE SPECIFIC FAILURE

Three follow-ups that make the next attempt more useful

When a response fails, identify whether the problem is a false fact, a broken format or an unverified action. Use a focused follow-up with the original source still available. A new polished answer needs its own check; a correction is not proof that everything is fixed.

COPY PROMPT 14 - REMOVE THE UNSUPPORTED CLAIM

Your summary says publication is confirmed for 29 October, but the notes say it was only proposed. Correct that claim using the original notes. Preserve the source-dependent 28 October draft target. Show the corrected sentence and briefly identify what changed. Add no new dates or commitments.

COPY PROMPT 15 - REPAIR JSON WITHOUT INVENTING VALUES

Return the extraction again as valid JSON only. Use the exact required keys and numeric types from the original request. Keep `tax_usd` and `email` null because the source does not supply them. Do not use markdown fences, comments or guessed values. Recalculate the subtotal from the source before answering.

COPY PROMPT 16 - ASK FOR EVIDENCE OF A CHECK

Before calling the coding task complete, report the exact check command, whether it ran, the observed result and any failure. If no execution tool is available, say which checks remain unrun. Distinguish a suggested fix from a verified result. Do not invent a test log or add unrelated features.

WORKED EXAMPLE - AN HONEST LIMIT

An agent without execution access should say that the function was proposed but its tests were not run. In a connected environment, inspect the actual tool result or test output. A schema can constrain output shape, but it does not certify that a date, price or extracted value is true. Use source checks and application validation together.

Name the failure, preserve the evidence and verify the corrected result.

A SEVEN-DAY PRACTICAL TRIAL

Measure finished work, review time and real usage

Pick one repeatable task and compare it with your usual process. Log the prompt, model, effort, errors, corrections, time and actual usage. A fictional worked calculation can teach the method; it cannot tell you how much Haiku will save on your own work.

A SMALL TRIAL YOU CAN FINISH

DAYS 1-2	Record a manual baseline and run the same task with a complete brief.
DAYS 3-4	Review errors; change one part of the prompt and test again.
DAYS 5-6	Try another comparable input and one alternative effort setting.
DAY 7	Compare reviewed quality, total effort and actual cost; keep, revise or drop.

COPY PROMPT 17 - INCLUDE REVIEW AND REPAIRS

Review this fictional trial: manual task 40 minutes; AI setup 5, generation 1, review 10, repairs 8. Calculate total AI-assisted time, the time difference and percentage change. Recommend another trial only if the final reviewed quality is acceptable. Explain why one case does not establish typical savings.

WORKED EXAMPLE - VERIFIED TIME CALCULATION

$5 + 1 + 10 + 8 = 24$ minutes. Difference: $40 - 24 = 16$ minutes, or 40% of the manual baseline. This is one fictional case, including review and repair time. Repeat on comparable tasks before making a performance claim.

COPY PROMPT 18 - PLAN THE NEXT SMALL TEST

Use my trial log to suggest one next test. State the task, smallest required input, exact output and pass/fail check. Prefer a frequent, narrow task with an answer I can verify. Propose only one trial; do not connect accounts, change settings or schedule recurring work.

Keep the workflow because its reviewed output is useful, not because its demo was impressive.



DIGITAL WORLD PULSE

SOURCES & DISCLAIMER

Primary references, method and limits

- [1] Anthropic - Haiku 5.5 launch, October 7
- [2] Claude Platform - Haiku 5.5 model specifications
- [3] Anthropic - Haiku access and use cases
- [4] Claude Platform - Effort settings and model guidance
- [5] Claude Platform - Prompting Haiku 5.5
- [6] Claude Platform - Pricing and the 100K threshold
- [7] Anthropic Privacy - Consumer data and model training
- [8] Anthropic Privacy - Commercial data and model training
- [9] Claude API - Create a Message
- [10] Claude Platform - Structured outputs
- [11] Claude Platform - Count tokens before a request
- [12] Claude Help - Change model, effort and thinking settings

Source titles and numbered page references are clickable. All sources checked October 21, 2026.

METHOD & EDITORIAL DISCLAIMER

Independent Digital World Pulse guide researched October 21, 2026; not affiliated with or endorsed by Anthropic. Product facts come from the linked primary sources. Prices, access, limits and settings can change. No live Haiku session, API call or independent model benchmark was performed for this guide. The 18 prompts and worked responses are original editorial examples, not captured Claude outputs or guaranteed results. Fictional notes, orders, policies, offers and datasets are labelled as practice material. The printed Python sample, JSON syntax, arithmetic and word limits were checked locally. AI-generated artwork is conceptual, not a product screenshot. Verify the facts and the resulting actions before using real work.

Use this guide as a practical starting point; verify the result in your own account.