

AI BUSINESS OPERATING SYSTEM

27 Ready-to-Use AI Prompts

Copy-ready prompts from AI Business Operating System. Replace bracketed placeholders with your own business context.

How to use this pack

Each prompt is reproduced from the book. Copy the prompt, replace the bracketed placeholders, keep the stated constraints intact, and review every output against the human-gate rules in the relevant chapter.

Prompt 1.1 — Surfacing hidden recurring tasks

You are helping a business owner complete a task-capture exercise for automation planning. I will describe one role in my business in plain language. Break that role into every distinct recurring task you can identify, even ones that might seem too small or too obvious to mention. For each task, state: (1) the task name in five words or fewer, (2) your best estimate of how often it recurs — daily, weekly, monthly, or irregular, (3) one clarifying question I would need to answer to confirm this task belongs on my list. List tasks separately even if they look related; do not group them. Do not recommend which tasks to automate — that decision uses a separate scoring method I will apply myself, after this list is complete. Here is the role: [describe the role, its main responsibilities, and a typical week, in 150–300 words].

Prompt 1.2 — Turning a task list into audit-ready entries

I will paste a rough, informally worded list of tasks captured during a one-week audit. Rewrite each item as a single, specific task description that could be scored against a fixed four-part test — meaning each rewritten task must have a clear start point and end point, must not bundle more than one distinct activity, and must not use a department or category name in place of an actual task. Where an item is too vague to rewrite as a single task, do not guess — instead, list it under a separate heading titled “Needs clarification” with the specific question I would need to answer to split or clarify it. Do not score, rank, or recommend anything about these tasks. Here is the list: [paste the raw captured list, one task per line].

Prompt 2.1 — Drafting a Connect-layer specification

I need a written specification for connecting two systems in my business, to hand to a developer or IT contractor. I will describe both systems and what needs to move between them. Produce a specification with these exact sections: (1) Source system and the specific data fields to be sent, (2) Destination system and the specific data fields it expects to receive, (3) Trigger — what event causes data to move (a new record, a schedule, a manual action), (4) Failure handling — what should happen if the destination system is unreachable when the trigger fires, (5) A one-paragraph plain-language summary suitable for a non-technical manager to approve before work begins. Do not recommend specific software products or platforms —

describe the specification in terms general enough to be built on whatever tools I already use. Here are the two systems and what needs to move between them: [describe source system, destination system, and the specific data involved].

Prompt 3.1 — Drafting a first BRIEF-format prompt

Help me draft a complete prompt using a six-part structure: Background, Role, Instructions, Examples, Format, and RULES. I will describe the task I want an AI system to perform repeatedly in my business. Produce a draft prompt with all six sections clearly labeled. In the Background section, include only context I actually provide — do not invent details about my business. In the RULES section, propose at least three specific constraints appropriate to this exact task, stated as “never” or “always” instructions, not general cautions. Leave the Examples section as a labeled placeholder with the instruction “insert one real, worked example of the desired output here” only if I have not provided one — otherwise use the example I give you. Here is the task: [describe the task, who it’s for, and what a correct result looks like, in 100–200 words].

Prompt 3.2 — Diagnosing a failing prompt

I am going to describe a prompt I’m using and a specific case where its output was wrong. Work through this sequence, in order, and stop at the first question that applies: (1) Did the output do something the prompt’s rules explicitly forbade? (2) Did it rely on facts or context not present in the prompt’s background section? (3) Did it adopt a tone or level of authority inconsistent with the role I defined? (4) Did it misunderstand the task itself, doing something other than what the instructions asked? (5) Did it complete the task correctly but in the wrong format? If none of the above, state that the most likely cause is a missing worked example. For whichever question applies, name the specific line in my prompt that should be added or changed, and write that line for me directly — do not describe the fix in the abstract. Here is my current prompt and the case where it failed: [paste prompt, then describe the input and the incorrect output].

Prompt 4.1 — Drafting category-specific auto-reply templates

I will give you a category of customer email and three to five real examples of messages in that category, along with the correct reply for each. Using the BRIEF method, draft a reusable reply template for this category: Background (what this category of email is and who it’s from), Role (a helpful, accurate representative of the business, never overpromising), Instructions (how to extract the specific details needed — order number, date, item — from an incoming message and insert them into the reply), Examples (use the ones I provide), Format (a short, direct email reply, no more than 120 words), and RULES (state explicitly: never confirm a delivery date not present in the order data; never apologize for a fault without a manager’s review; never mention pricing not already stated in the order record). Here is the category and examples: [paste category name and 3–5 example message/reply pairs].

Prompt 4.2 — Drafting the unclear-case refusal message

Write a short, honest message to be sent automatically whenever an inbound email cannot be confidently categorized or answered. The message must: (1) never pretend to have answered the question, (2) state plainly that the message has been forwarded to a person, (3) name the general area of the business handling it

(not a specific individual's name), (4) give a realistic response-time expectation. Keep it under 40 words, and do not use any apologetic language beyond a single brief acknowledgment — the goal is clarity, not a display of concern.

Prompt 5.1 — Drafting an ARS scoring note

I will describe one reference document my business uses. Score it against these eight factors, answering each with 1 (clearly meets it) or 0 (does not), and a one-sentence justification for each score: Authority — (1) is the source named and identifiable, (2) does the source have direct authority over this subject, (3) is it verifiable against a primary document. Recency — (4) is a publication or issue date known, (5) is it still within any stated validity window. Specificity — (6) does it directly address the exact question it would be used for, (7) does it contain specific actionable detail rather than general background, (8) does it avoid contradicting anything else in our reference library. Total the eight scores and state the resulting status: 7–8 Approved, 5–6 Approved with review date, 3–4 Provisional, 0–2 Rejected. If you lack information to score any factor confidently, say so explicitly rather than guessing, and tell me exactly what information you'd need. Here is the document: [describe the document, its source, its date if known, and its intended use].

Prompt 5.2 — Flagging likely recency risk

I will paste a list of documents with their publication dates and subject areas. For each one, flag whether it is likely to have been superseded, based only on the subject area and how frequently that category of document typically changes (for example: tax thresholds, regulatory notifications, and vendor pricing change often; internal process documents and historical records typically do not). Do not claim to know whether a specific document has actually been superseded — you cannot verify that. State only a likelihood level (High / Medium / Low) and the reasoning behind it, and recommend a review-date interval for anything you mark Medium or High. Here is the list: [paste document name, date, and subject area, one per line].

Prompt 6.1 — Tracing a finding to its primary source

I will give you a research finding and the source I currently have for it. Assess whether this source is likely Tier 1 (a primary, original document), Tier 2 (a credible secondary account that itself cites a primary source), or Tier 3 (anything else, including unverified or uncited claims). If Tier 2 or 3, tell me specifically what kind of primary source would need to exist to upgrade this finding, and suggest concrete search terms I could use to look for it myself — do not claim to have found or verified that primary source unless I explicitly provide it to you. Here is the finding and its current source: [paste finding and source].

Prompt 7.1 — Explaining a KPI movement

I will give you this period's value for a KPI, its definition and formula, the prior period's value, and the underlying data that fed the calculation — all from our own approved systems only. Write a plain-language explanation, no more than 100 words, of what changed and the most likely contributing factor visible in the data I've given you. Do not speculate about causes not supported by the data provided — for example, do not attribute a change to external events (weather, competitor action, economic conditions) unless that information is explicitly included in what I've given you. If the data doesn't clearly explain the movement, say so directly rather than guessing.

Here is the KPI, its formula, both periods' values, and the underlying data: [paste KPI definition, formula, current and prior values, and supporting data].

Prompt 8.1 — Drafting content against pillars and voice

Using the BRIEF method, draft a piece of content for the pillar and format I specify. Background: our content pillars are [list your pillars] and our voice attributes are [list your voice attributes]. Role: an in-house content writer who never overstates a claim. Instructions: write a piece mapped explicitly to the pillar I name, in the format I specify. Examples: [paste one approved example of this format in our voice]. Format: [specify length and structure]. RULES: never use any phrase or claim pattern from this prohibited list: [paste your prohibited-language list]; never make a comparison to a named or unnamed competitor; never state a claim about product performance without citing an approved source. Here is the pillar, format, and topic: [specify].

Prompt 9.1 — Planning derivatives from a Knowledge Asset

I will describe one approved Knowledge Asset — its topic, its main points, and its source material. Propose ten distinct derivative formats this asset could be repurposed into, each suited to a different channel or moment of use (for example: a short social caption, an email snippet, a printable reference card, a spoken-style script). For each, state the format, the likely channel, and one sentence on which specific part of the parent asset it draws from. Do not draft the derivatives themselves yet — only plan them. Here is the asset: [describe the topic, main points, and source].

Prompt 9.2 — Drafting one derivative from an approved asset

Using the BRIEF method, draft the following single derivative from an already-approved Knowledge Asset. Background: the parent asset's approved content is [paste the relevant excerpt]. Role: a writer producing this one specific format only. Instructions: produce exactly the derivative I specify below, using only information present in the parent asset. Examples: [paste one example of this format done well, if available]. Format: [specify the exact format — word count, structure, channel]. RULES: never introduce a claim, number, or recommendation not present in the parent asset; never contradict the parent asset's stated guidance. Here is the specific derivative to produce: [specify format and channel].

Prompt 10.1 — Drafting a platform-adapted post

I will give you one piece of approved source content and a target platform. Using the BRIEF method, adapt it for that platform. Background: the source content is [paste it]. Role: a social media writer for our brand, matching the voice attributes [list them]. Instructions: adapt the source content for the specified platform's typical format and length, without adding any claim not present in the source. Examples: [paste one approved example post on this platform]. Format: [specify platform-appropriate length and structure]. RULES: never add a claim not present in the source content; never use any phrase from this prohibited list: [paste prohibited-language list from Chapter 8]. Here is the platform: [specify].

Prompt 10.2 — Classifying an incoming message

I will paste an incoming social media message. Classify it as one of exactly four types: General Inquiry, Sales Interest, Complaint, or Spam/Irrelevant. State your classification and a one-sentence justification. If the message shows any sign of dissatisfaction, frustration, or a negative experience — even mild or indirect — classify it as Complaint rather than General Inquiry; when uncertain between the two, always choose Complaint. Do not draft a reply. Here is the message: [paste message].

Prompt 11.1 — Reviewing a transcript for factual accuracy

I will paste a transcript of a recording. Read it and flag, as a numbered list, any statement that appears to be a specific factual claim — a number, a date, a name, a guarantee, or a comparison — so that I can verify each one against our approved sources before this transcript moves to editing. Do not attempt to verify the claims yourself; you have no way to confirm they're accurate. Simply list them clearly, with the surrounding sentence for context, so I can check each one quickly. Here is the transcript: [paste transcript].

Prompt 11.2 — Drafting a candidate highlight cut list

I will paste a reviewed, accurate transcript with timestamps. Identify five to eight candidate highlight segments suitable for a short-form edit, each 15–45 seconds long. For each, give the start and end timestamp, a one-sentence description of what it covers, and a note on why it's a strong standalone moment. Do not combine or reorder segments from different parts of the transcript into a single candidate — each candidate must be one continuous, unedited span of the original recording, so a human reviewer can verify it represents what was actually said in that moment. Here is the transcript: [paste timestamped transcript].

Prompt 12.1 — Drafting an asset against an approved template

I will describe an approved template's structure and our minimum visual specification in words (layout, required elements, typography rules — not the actual visual file, which I will apply separately). Using that specification, draft the text content for a new asset of this type, ensuring every required element from the template structure is addressed. Do not suggest colours, fonts, or visual styling — describe content only; the approved template file itself will supply the visual treatment. Here is the template structure and the content needed: [describe template's required elements and the specific content for this instance].

Prompt 13.1 — Applying the qualification standard to an inbound lead

I will give you our six disqualifiers, our eight ICP criteria, our five fit questions, and the details of one inbound lead. Check the lead against the disqualifiers first, in order — if any disqualifier is confirmed true based only on the information I provide, stop and return DISQUALIFY with the specific disqualifier that applied. If no disqualifier applies, check the lead against the ICP criteria and the five fit questions, marking each fit question Yes, No, or Unknown based only on the information given — never assume an answer that wasn't stated. Return one of exactly three verdicts: QUALIFY, NURTURE (state whether it's an ICP-fit gap or a needs-more-information gap), or DISQUALIFY. Here are our criteria and the lead: [paste disqualifiers, ICP criteria, fit questions, and lead details].

Prompt 14.1 — Drafting a lifecycle-stage message

Using the BRIEF method, draft a message for the lifecycle stage I specify. Background: this message is for contacts in the [stage name] stage of our customer lifecycle, with [engagement level] engagement. Role: a relationship-focused communicator, never a purely promotional one. Instructions: write a message appropriate to this exact stage and engagement combination, checking first whether this combination appears in our suppressed-cell list — if it does, respond only with “SUPPRESSED — do not draft a message for this combination” and nothing else. Examples: [paste an approved example for this stage if available]. Format: [specify channel and length]. RULES: never draft a message for a suppressed stage/engagement combination; never use language from the prohibited list in Chapter 8. Here is the stage, engagement level, and our current suppressed-cell list: [specify].

Prompt 15.1 — Flagging stale commitments

I will paste our Commitment Log for one active deal, including every logged commitment with its due date. Identify any commitment whose due date has passed without a follow-up entry showing it was fulfilled, and any deal where the most recent logged commitment is more than three weeks old. For each one you flag, state exactly which commitment is stale and how many days past due or since logged it is. Do not speculate about why the commitment wasn't fulfilled — simply flag it factually for a human to review. Here is the log: [paste Commitment Log entries with dates].

Prompt 16.1 — Checking a contact against the never-automate list

I will give you our never-automate category list and one incoming customer message. Assess, conservatively, whether the message could reasonably match any category on the list, even if it doesn't use the exact language you'd expect — interpret the categories by substance and intent, not by keyword matching. If there is any reasonable ambiguity about whether it matches, treat it as a match and recommend immediate human routing. Only recommend an automated response path if you are confident no category applies. State your reasoning briefly. Here is our never-automate list and the message: [paste list and message].

Prompt 17.1 — Drafting a first-pass procedure from a transcript

I will paste a transcript of someone narrating how they perform a task. Convert it into a structured first-pass procedure: identify whether it reads as a fixed sequence (SOP), a set of independent checks (checklist), or a standing constraint (rule), and format it accordingly. Remove filler language and verbal habits from the transcript, but do not add any step, condition, or assumption that wasn't stated. Flag, separately, any point where the narration seems to skip a step or assumes context a new reader wouldn't have — list these as open questions for the original narrator to clarify, rather than guessing at what they meant. Here is the transcript: [paste narration transcript].

Prompt 18.1 — Flagging low-confidence agent decisions for escalation

I will describe a decision an agent needs to make and the information available to it at the time. Assess whether

this decision falls clearly within the defined authority limits I will provide, or whether it should be escalated to a person instead. Treat any case involving missing required information, an amount or scope outside the stated numeric limits, or a situation not clearly covered by the stated limits as requiring escalation — do not extend the limits by inference to cover an edge case they don't explicitly address. State your assessment and reasoning briefly. Here are the authority limits and the decision in question: [paste authority limits and decision details].

Prompt 19.1 — Compiling the weekly Business Brain review

I will paste this week's logged data from four sources: our artefact review-date tracker, our KPI Definition Sheet's period-over-period figures, our Escalation Standard's logged contacts, and our agent action log. Compile a single summary covering exactly four sections: artefacts overdue for review (name and how many days overdue), the KPI with the largest movement this week and its already-approved explanation, escalations logged this week (count and categories, no individual customer detail), and any agent actions that ran at Level 3 this week (count and category). Do not add commentary or recommendations beyond what the data directly shows. Here is this week's data: [paste tracker, KPI figures, escalation log, and agent action log].

Prompt 20.1 — Drafting a data-portability request

I will describe a tool we currently use and the categories of data and configuration we've built inside it. Draft a clear, specific written request to the vendor asking them to confirm: (1) whether a full structured export of this data exists, (2) in what file format, (3) whether the export includes configuration and workflow logic or only raw content, and (4) how long exported data remains accessible after a contract ends. Keep the request factual and specific — do not speculate about the vendor's likely answer or assume any capability they haven't confirmed. Here is the tool and the data categories: [describe tool and data/configuration held within it].